



Red Hat Integration 2023.q1

Camel Spring Boot Reference 3.14

Camel Spring Boot Reference

Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide describes the settings for Camel Spring Boot components.

Table of Contents

PREFACE	29
MAKING OPEN SOURCE MORE INCLUSIVE	29
CHAPTER 1. AWS CLOUDWATCH	30
1.1. URI FORMAT	30
1.2. CONFIGURING OPTIONS	30
1.2.1. Configuring Component Options	30
1.2.2. Configuring Endpoint Options	30
1.3. COMPONENT OPTIONS	31
1.4. ENDPOINT OPTIONS	32
1.4.1. Path Parameters (1 parameters)	33
1.4.2. Query Parameters (16 parameters)	33
1.5. USAGE	34
1.5.1. Static credentials vs Default Credential Provider	34
1.5.2. Message headers evaluated by the CW producer	35
1.5.3. Advanced CloudWatchClient configuration	35
1.6. DEPENDENCIES	35
1.7. EXAMPLES	36
1.7.1. Producer Example	36
1.8. SPRING BOOT AUTO-CONFIGURATION	36
CHAPTER 2. AWS DYNAMODB	39
2.1. URI FORMAT	39
2.2. CONFIGURING OPTIONS	39
2.2.1. Configuring Component Options	39
2.2.2. Configuring Endpoint Options	39
2.3. COMPONENT OPTIONS	39
2.4. ENDPOINT OPTIONS	42
2.4.1. Path Parameters (1 parameters)	42
2.4.2. Query Parameters (20 parameters)	42
2.5. USAGE	45
2.5.1. Static credentials vs Default Credential Provider	45
2.5.2. Message headers evaluated by the DDB producer	45
2.5.3. Message headers set during BatchGetItems operation	47
2.5.4. Message headers set during DeleteItem operation	47
2.5.5. Message headers set during DeleteTable operation	47
2.5.6. Message headers set during DescribeTable operation	48
2.5.7. Message headers set during GetItem operation	48
2.5.8. Message headers set during PutItem operation	49
2.5.9. Message headers set during Query operation	49
2.5.10. Message headers set during Scan operation	49
2.5.11. Message headers set during UpdateItem operation	50
2.5.12. Advanced AmazonDynamoDB configuration	50
2.6. SUPPORTED PRODUCER OPERATIONS	50
2.7. EXAMPLES	50
2.7.1. Producer Examples	50
2.8. SPRING BOOT AUTO-CONFIGURATION	51
CHAPTER 3. AWS KINESIS	57
3.1. URI FORMAT	57
3.2. CONFIGURING OPTIONS	57
3.2.1. Configuring Component Options	57

3.2.2. Configuring Endpoint Options	57
3.3. COMPONENT OPTIONS	58
3.4. ENDPOINT OPTIONS	60
3.4.1. Path Parameters (1 parameters)	61
3.4.2. Query Parameters (38 parameters)	61
3.5. BATCH CONSUMER	66
3.6. USAGE	66
3.6.1. Static credentials vs Default Credential Provider	66
3.6.2. Message headers set by the Kinesis consumer	66
3.6.3. AmazonKinesis configuration	67
3.6.4. Providing AWS Credentials	67
3.6.5. Message headers used by the Kinesis producer to write to Kinesis. The producer expects that the message body is a byte[].	67
3.6.6. Message headers set by the Kinesis producer on successful storage of a Record	67
3.7. DEPENDENCIES	68
3.8. SPRING BOOT AUTO-CONFIGURATION	68
CHAPTER 4. AWS 2 LAMBDA	74
4.1. URI FORMAT	74
4.2. CONFIGURING OPTIONS	74
4.2.1. Configuring Component Options	74
4.2.2. Configuring Endpoint Options	74
4.3. COMPONENT OPTIONS	75
4.4. ENDPOINT OPTIONS	77
4.4.1. Path Parameters (1 parameters)	78
4.4.2. Query Parameters (14 parameters)	78
4.5. USAGE	80
4.5.1. Static credentials vs Default Credential Provider	80
4.5.2. Message headers evaluated by the Lambda producer	81
4.6. LIST OF AVAILAIBLE OPERATIONS	84
4.7. EXAMPLES	85
4.7.1. Producer Example	85
4.7.2. Producer Examples	85
4.8. USING A POJO AS BODY	85
4.9. DEPENDENCIES	86
4.10. SPRING BOOT AUTO-CONFIGURATION	86
CHAPTER 5. AWS S3 STORAGE SERVICE	89
5.1. URI FORMAT	89
5.2. CONFIGURING OPTIONS	89
5.2.1. Configuring Component Options	89
5.2.2. Configuring Endpoint Options	89
5.3. COMPONENT OPTIONS	90
5.4. ENDPOINT OPTIONS	96
5.4.1. Path Parameters (1 parameters)	96
5.4.2. Query Parameters (68 parameters)	96
5.5. BATCH CONSUMER	104
5.6. USAGE	104
5.6.1. Message headers evaluated by the S3 producer	104
5.6.2. Message headers set by the S3 producer	106
5.6.3. Message headers set by the S3 consumer	106
5.6.4. S3 Producer operations	107
5.6.5. Advanced AmazonS3 configuration	108

5.6.6. Use KMS with the S3 component	108
5.6.7. Static credentials vs Default Credential Provider	108
5.6.8. S3 Producer Operation examples	109
5.7. STREAMING UPLOAD MODE	111
5.8. BUCKET AUTOCREATION	113
5.9. MOVING STUFF BETWEEN A BUCKET AND ANOTHER BUCKET	113
5.10. MOVEAFTERREAD CONSUMER OPTION	113
5.11. USING CUSTOMER KEY AS ENCRYPTION	114
5.12. USING A POJO AS BODY	114
5.13. CREATE S3 CLIENT AND ADD COMPONENT TO REGISTRY	114
5.14. DEPENDENCIES	115
5.15. SPRING BOOT AUTO-CONFIGURATION	115
CHAPTER 6. AWS SIMPLE NOTIFICATION SYSTEM (SNS)	122
6.1. URI FORMAT	122
6.2. URI OPTIONS	122
6.2.1. Configuring Options	122
6.2.1.1. Configuring Component Options	122
6.2.1.2. Configuring Endpoint Options	122
6.3. COMPONENT OPTIONS	123
6.4. ENDPOINT OPTIONS	125
6.4.1. Path Parameters (1 parameters)	126
6.4.2. Query Parameters (23 parameters)	126
6.5. USAGE	128
6.5.1. Static credentials vs Default Credential Provider	128
6.5.2. Message headers evaluated by the SNS producer	129
6.5.3. Message headers set by the SNS producer	129
6.5.4. Advanced AmazonSNS configuration	129
6.5.5. Create a subscription between an AWS SNS Topic and an AWS SQS Queue	129
6.6. TOPIC AUTOCREATION	130
6.7. SNS FIFO	130
6.7.1. SNS Fifo Topic Message group Id Strategy and message Deduplication Id Strategy	131
6.8. EXAMPLES	131
6.8.1. Producer Examples	131
6.9. DEPENDENCIES	131
6.10. SPRING BOOT AUTO-CONFIGURATION	131
CHAPTER 7. AWS SIMPLE QUEUE SERVICE (SQS)	135
7.1. URI FORMAT	135
7.2. CONFIGURING OPTIONS	135
7.2.1. Configuring Component Options	135
7.2.2. Configuring Endpoint Options	135
7.3. COMPONENT OPTIONS	136
7.4. ENDPOINT OPTIONS	140
7.4.1. Path Parameters (1 parameters)	141
7.4.2. Query Parameters (61 parameters)	141
7.5. BATCH CONSUMER	148
7.6. USAGE	148
7.6.1. Static credentials vs Default Credential Provider	148
7.6.2. Message headers set by the SQS producer	149
7.6.3. Message headers set by the SQS consumer	149
7.6.4. Advanced AmazonSQS configuration	149
7.6.5. Creating or updating an SQS Queue	149

7.6.6. DelayQueue VS Delay for Single message	150
7.6.7. Server Side Encryption	150
7.7. JMS-STYLE SELECTORS	150
7.8. AVAILABLE PRODUCER OPERATIONS	151
7.9. SEND MESSAGE	151
7.10. SEND BATCH MESSAGE	151
7.11. DELETE SINGLE MESSAGE	151
7.12. LIST QUEUES	152
7.13. PURGE QUEUE	152
7.14. QUEUE AUTOCREATION	152
7.15. SEND BATCH MESSAGE AND MESSAGE DEDUPLICATION STRATEGY	152
7.16. DEPENDENCIES	152
7.17. SPRING BOOT AUTO-CONFIGURATION	153
CHAPTER 8. AZURE STORAGE BLOB SERVICE	159
8.1. URI FORMAT	159
8.2. CONFIGURING OPTIONS	159
8.2.1. Configuring Component Options	159
8.2.2. Configuring Endpoint Options	160
8.3. COMPONENT OPTIONS	160
8.4. ENDPOINT OPTIONS	165
8.4.1. Path Parameters (2 parameters)	165
8.4.2. Query Parameters (48 parameters)	165
8.5. USAGE	173
8.5.1. Message headers evaluated by the component producer	173
8.5.2. Message headers set by either component producer or consumer	179
8.5.3. Advanced Azure Storage Blob configuration	182
8.5.4. Automatic detection of BlobServiceClient client in registry	182
8.5.5. Azure Storage Blob Producer operations	182
8.5.6. Consumer Examples	185
8.5.7. Producer Operations Examples	186
8.5.8. Development Notes (Important)	191
8.6. SPRING BOOT AUTO-CONFIGURATION	191
CHAPTER 9. AZURE STORAGE QUEUE SERVICE	196
9.1. URI FORMAT	196
9.2. CONFIGURING OPTIONS	196
9.2.1. Configuring Component Options	196
9.2.2. Configuring Endpoint Options	197
9.3. COMPONENT OPTIONS	197
9.4. ENDPOINT OPTIONS	199
9.4.1. Path Parameters (2 parameters)	199
9.4.2. Query Parameters (31 parameters)	200
9.5. USAGE	205
9.5.1. Message headers evaluated by the component producer	205
9.5.2. Message headers set by either component producer or consumer	207
9.5.3. Advanced Azure Storage Queue configuration	208
9.5.4. Automatic detection of QueueServiceClient client in registry	208
9.5.5. Azure Storage Queue Producer operations	208
9.5.6. Consumer Examples	209
9.5.7. Producer Operations Examples	209
9.5.8. Development Notes (Important)	212
9.6. SPRING BOOT AUTO-CONFIGURATION	212

CHAPTER 10. BEAN	216
10.1. URI FORMAT	216
10.2. CONFIGURING OPTIONS	216
10.2.1. Configuring Component Options	216
10.2.2. Configuring Endpoint Options	216
10.3. COMPONENT OPTIONS	216
10.4. ENDPOINT OPTIONS	218
10.4.1. Path Parameters (1 parameters)	218
10.4.2. Query Parameters (5 parameters)	218
10.5. USING	219
10.6. BEAN AS ENDPOINT	220
10.7. JAVA DSL BEAN SYNTAX	220
10.8. BEAN BINDING	220
10.9. SPRING BOOT AUTO-CONFIGURATION	221
CHAPTER 11. BEAN VALIDATOR	225
11.1. URI FORMAT	225
11.2. CONFIGURING OPTIONS	225
11.2.1. Configuring Component Options	225
11.2.2. Configuring Endpoint Options	225
11.3. COMPONENT OPTIONS	226
11.4. ENDPOINT OPTIONS	227
11.4.1. Path Parameters (1 parameters)	227
11.4.2. Query Parameters (8 parameters)	227
11.5. OSGI DEPLOYMENT	228
11.6. EXAMPLE	228
11.7. SPRING BOOT AUTO-CONFIGURATION	231
CHAPTER 12. BROWSE	233
12.1. URI FORMAT	233
12.2. CONFIGURING OPTIONS	233
12.2.1. Configuring Component Options	233
12.2.2. Configuring Endpoint Options	233
12.3. COMPONENT OPTIONS	233
12.4. ENDPOINT OPTIONS	234
12.4.1. Path Parameters (1 parameters)	234
12.4.2. Query Parameters (4 parameters)	234
12.5. SAMPLE	235
12.6. SPRING BOOT AUTO-CONFIGURATION	236
CHAPTER 13. CASSANDRA CQL	238
13.1. CONFIGURING OPTIONS	238
13.1.1. Configuring Component Options	238
13.1.2. Configuring Endpoint Options	238
13.2. COMPONENT OPTIONS	238
13.3. ENDPOINT OPTIONS	239
13.3.1. Path Parameters (4 parameters)	239
13.3.2. Query Parameters (30 parameters)	240
13.4. ENDPOINT CONNECTION SYNTAX	244
13.5. MESSAGES	244
13.5.1. Incoming Message	244
13.5.2. Outgoing Message	244
13.6. REPOSITORIES	245
13.7. IDEMPOTENT REPOSITORY	245

13.8. AGGREGATION REPOSITORY	245
13.9. EXAMPLES	246
13.10. SPRING BOOT AUTO-CONFIGURATION	247
CHAPTER 14. CONTROL BUS	249
14.1. COMMANDS	249
14.2. CONFIGURING OPTIONS	249
14.2.1. Configuring Component Options	249
14.2.2. Configuring Endpoint Options	250
14.3. COMPONENT OPTIONS	250
14.4. ENDPOINT OPTIONS	250
14.4.1. Path Parameters (2 parameters)	250
14.4.1.1. Query Parameters (6 parameters)	251
14.5. USING ROUTE COMMAND	253
14.6. GETTING PERFORMANCE STATISTICS	253
14.7. USING SIMPLE LANGUAGE	254
14.8. SPRING BOOT AUTO-CONFIGURATION	254
CHAPTER 15. CRON	256
15.1. CONFIGURING OPTIONS	256
15.1.1. Configuring Component Options	256
15.1.2. Configuring Endpoint Options	256
15.2. COMPONENT OPTIONS	257
15.3. ENDPOINT OPTIONS	257
15.3.1. Path Parameters (1 parameters)	257
15.3.2. Query Parameters (4 parameters)	258
15.4. USAGE	258
15.5. SPRING BOOT AUTO-CONFIGURATION	259
CHAPTER 16. CXF	261
16.1. URI FORMAT	261
16.2. CONFIGURING OPTIONS	261
16.2.1. Configuring Component Options	261
16.2.2. Configuring Endpoint Options	262
16.3. COMPONENT OPTIONS	262
16.4. ENDPOINT OPTIONS	263
16.4.1. Path Parameters (2 parameters)	263
16.4.2. Query Parameters (35 parameters)	263
16.4.3. Descriptions of the dataformats	267
16.4.4. How to enable CXF's LoggingOutInterceptor in RAW mode	268
16.4.5. Description of relayHeaders option	268
16.4.6. Available only in POJO mode	269
16.5. CONFIGURE THE CXF ENDPOINTS WITH SPRING	271
16.6. HOW TO MAKE THE CAMEL-CXF COMPONENT USE LOG4J INSTEAD OF JAVA.UTIL.LOGGING	274
16.7. HOW TO LET CAMEL-CXF RESPONSE START WITH XML PROCESSING INSTRUCTION	274
16.8. HOW TO OVERRIDE THE CXF PRODUCER ADDRESS FROM MESSAGE HEADER	275
16.9. HOW TO CONSUME A MESSAGE FROM A CAMEL-CXF ENDPOINT IN POJO DATA FORMAT	275
16.10. HOW TO PREPARE THE MESSAGE FOR THE CAMEL-CXF ENDPOINT IN POJO DATA FORMAT	276
16.11. HOW TO DEAL WITH THE MESSAGE FOR A CAMEL-CXF ENDPOINT IN PAYLOAD DATA FORMAT	276
16.12. HOW TO GET AND SET SOAP HEADERS IN POJO MODE	277
16.13. HOW TO GET AND SET SOAP HEADERS IN PAYLOAD MODE	279
16.14. SOAP HEADERS ARE NOT AVAILABLE IN RAW MODE	280
16.15. HOW TO THROW A SOAP FAULT FROM CAMEL	280
16.16. HOW TO PROPAGATE A CAMEL-CXF ENDPOINT'S REQUEST AND RESPONSE CONTEXT	281

16.17. ATTACHMENT SUPPORT	281
16.18. STREAMING SUPPORT IN PAYLOAD MODE	284
16.19. USING THE GENERIC CXF DISPATCH MODE	284
16.20. SPRING BOOT AUTO-CONFIGURATION	285
CHAPTER 17. DATA FORMAT	288
17.1. URI FORMAT	288
17.2. DATAFORMAT OPTIONS	288
17.2.1. Configuring Options	288
17.2.1.1. Configuring Component Options	288
17.2.1.2. Configuring Endpoint Options	288
17.3. COMPONENT OPTIONS	288
17.4. ENDPOINT OPTIONS	289
17.4.1. Path Parameters (2 parameters)	289
17.4.2. Query Parameters (1 parameters)	289
17.5. SAMPLES	290
17.6. SPRING BOOT AUTO-CONFIGURATION	290
CHAPTER 18. DATASET	292
18.1. URI FORMAT	292
18.2. CONFIGURING OPTIONS	292
18.2.1. Configuring Component Options	292
18.2.2. Configuring Endpoint Options	292
18.3. COMPONENT OPTIONS	293
18.4. ENDPOINT OPTIONS	294
18.4.1. Path Parameters (1 parameters)	294
18.4.2. Query Parameters (21 parameters)	294
18.5. CONFIGURING DATASET	298
18.6. EXAMPLE	298
18.7. DATASETSUPPORT (ABSTRACT CLASS)	299
18.7.1. Properties on DataSetSupport	299
18.8. SIMPLEDATASET	299
18.8.1. Additional Properties on SimpleDataSet	299
18.9. LISTDATASET	300
18.9.1. Additional Properties on ListDataSet	300
18.10. FILEDATASET	300
18.10.1. Additional Properties on FileDataSet	300
18.11. SPRING BOOT AUTO-CONFIGURATION	300
CHAPTER 19. DIRECT	303
19.1. URI FORMAT	303
19.2. CONFIGURING OPTIONS	303
19.2.1. Configuring Component Options	303
19.2.2. Configuring Endpoint Options	303
19.3. COMPONENT OPTIONS	304
19.4. ENDPOINT OPTIONS	304
19.4.1. Path Parameters (1 parameters)	305
19.4.2. Query Parameters (8 parameters)	305
19.5. SAMPLES	306
19.6. SPRING BOOT AUTO-CONFIGURATION	307
CHAPTER 20. FHIR	309
20.1. URI FORMAT	309
20.2. CONFIGURING OPTIONS	309

20.2.1. Configuring Component Options	310
20.2.2. Configuring Endpoint Options	310
20.3. COMPONENT OPTIONS	310
20.4. ENDPOINT OPTIONS	313
20.4.1. Path Parameters (2 parameters)	314
20.4.2. Query Parameters (44 parameters)	314
20.5. API PARAMETERS (13 APIS)	319
20.5.1. API: capabilities	320
20.5.1.1. Method ofType	321
20.5.2. API: create	321
20.5.2.1. Method resource	321
20.5.3. API: delete	322
20.5.3.1. Method resource	323
20.5.3.2. Method resourceById	323
20.5.3.3. Method resourceConditionalByUrl	324
20.5.4. API: history	324
20.5.4.1. Method onInstance	325
20.5.4.2. Method onServer	325
20.5.4.3. Method onType	326
20.5.5. API: load-page	327
20.5.5.1. Method byUrl	327
20.5.5.2. Method next	328
20.5.5.3. Method previous	328
20.5.6. API: meta	328
20.5.6.1. Method add	329
20.5.6.2. Method delete	329
20.5.6.3. Method getFromResource	330
20.5.6.4. Method getFromServer	330
20.5.6.5. Method getFromType	331
20.5.7. API: operation	331
20.5.7.1. Method onInstance	332
20.5.7.2. Method onInstanceVersion	332
20.5.7.3. Method onServer	333
20.5.7.4. Method onType	334
20.5.7.5. Method processMessage	335
20.5.8. API: patch	335
20.5.8.1. Method patchById	336
20.5.8.2. Method patchByUrl	336
20.5.9. API: read	337
20.5.9.1. Method resourceById	338
20.5.9.2. Method resourceByUrl	339
20.5.10. API: search	340
20.5.10.1. Method searchByUrl	340
20.5.11. API: transaction	341
20.5.11.1. Method withBundle	341
20.5.11.2. Method withResources	342
20.5.12. API: update	342
20.5.12.1. Method resource	343
20.5.12.2. Method resourceBySearchUrl	343
20.5.13. API: validate	344
20.5.13.1. Method resource	344
20.6. SPRING BOOT AUTO-CONFIGURATION	345

CHAPTER 21. FILE	353
21.1. URI FORMAT	353
21.2. CONFIGURING OPTIONS	353
21.2.1. Configuring Component Options	353
21.2.2. Configuring Endpoint Options	353
21.3. COMPONENT OPTIONS	354
21.4. ENDPOINT OPTIONS	354
21.4.1. Path Parameters (1 parameters)	355
21.4.2. Query Parameters (94 parameters)	355
21.5. MOVE AND DELETE OPERATIONS	371
21.6. FINE GRAINED CONTROL OVER MOVE AND PREMOVE OPTION	372
21.7. ABOUT MOVEFAILED	372
21.8. MESSAGE HEADERS	373
21.8.1. File producer only	373
21.8.2. File consumer only	373
21.9. BATCH CONSUMER	374
21.10. EXCHANGE PROPERTIES, FILE CONSUMER ONLY	374
21.11. USING CHARSET	374
21.12. COMMON GOTCHAS WITH FOLDER AND FILENAMES	375
21.13. FILENAME EXPRESSION	376
21.14. CONSUMING FILES FROM FOLDERS WHERE OTHERS DROP FILES DIRECTLY	376
21.15. USING DONE FILES	376
21.16. WRITING DONE FILES	377
21.17. SAMPLES	377
21.17.1. Read from a directory and write to another directory	377
21.17.2. Read from a directory and write to another directory using a overrule dynamic name	378
21.17.3. Reading recursively from a directory and writing to another	378
21.18. USING FLATTEN	378
21.19. READING FROM A DIRECTORY AND THE DEFAULT MOVE OPERATION	378
21.20. READ FROM A DIRECTORY AND PROCESS THE MESSAGE IN JAVA	379
21.21. WRITING TO FILES	379
21.21.1. Write to subdirectory using Exchange.FILE_NAME	379
21.21.2. Writing file through the temporary directory relative to the final destination	379
21.22. USING EXPRESSION FOR FILENAMES	380
21.23. AVOIDING READING THE SAME FILE MORE THAN ONCE (IDEMPOTENT CONSUMER)	380
21.24. USING A FILE BASED IDEMPOTENT REPOSITORY	380
21.25. USING A JPA BASED IDEMPOTENT REPOSITORY	381
21.26. FILTER USING ORG.APACHE.CAMEL.COMPONENT.FILE.GENERICFILEFILTER	381
21.27. FILTERING USING ANT PATH MATCHER	382
21.27.1. Sorting using Comparator	382
21.27.2. Sorting using sortBy	383
21.28. USING GENERICFILEPROCESSTRATEGY	384
21.29. USING FILTER	384
21.30. USING BRIDGEERRORHANDLER	384
21.31. DEBUG LOGGING	385
21.32. SPRING BOOT AUTO-CONFIGURATION	385
CHAPTER 22. FTP	387
22.1. URI FORMAT	387
22.2. CONFIGURING OPTIONS	387
22.2.1. Configuring Component Options	388
22.2.2. Configuring Endpoint Options	388
22.3. COMPONENT OPTIONS	388

22.4. ENDPOINT OPTIONS	389
22.4.1. Path Parameters (3 parameters)	389
22.4.2. Query Parameters (111 parameters)	389
22.5. FTPS COMPONENT DEFAULT TRUST STORE	408
22.6. EXAMPLES	409
22.7. CONCURRENCY	409
22.8. MORE INFORMATION	409
22.9. DEFAULT WHEN CONSUMING FILES	409
22.9.1. limitations	409
22.10. MESSAGE HEADERS	410
22.10.1. Exchange Properties	410
22.11. ABOUT TIMEOUTS	411
22.12. USING LOCAL WORK DIRECTORY	411
22.13. STEPWISE CHANGING DIRECTORIES	411
22.14. USING STEPWISE=TRUE (DEFAULT MODE)	412
22.15. USING STEPWISE=FALSE	413
22.16. SAMPLES	414
22.16.1. Consuming a remote FTPS server (implicit SSL) and client authentication	414
22.16.2. Consuming a remote FTPS server (explicit TLS) and a custom trust store configuration	415
22.17. CUSTOM FILTERING	415
22.18. FILTERING USING ANT PATH MATCHER	415
22.19. USING A PROXY WITH SFTP	415
22.20. SETTING PREFERRED SFTP AUTHENTICATION METHOD	416
22.21. CONSUMING A SINGLE FILE USING A FIXED NAME	416
22.22. DEBUG LOGGING	416
22.23. SPRING BOOT AUTO-CONFIGURATION	417
CHAPTER 23. HTTP	420
23.1. URI FORMAT	420
23.2. CONFIGURING OPTIONS	420
23.2.1. Configuring Component Options	420
23.2.2. Configuring Endpoint Options	420
23.3. COMPONENT OPTIONS	421
23.4. ENDPOINT OPTIONS	424
23.4.1. Path Parameters (1 parameters)	425
23.4.2. Query Parameters (51 parameters)	425
23.5. MESSAGE HEADERS	431
23.6. MESSAGE BODY	432
23.7. USING SYSTEM PROPERTIES	432
23.8. RESPONSE CODE	432
23.9. EXCEPTIONS	433
23.10. WHICH HTTP METHOD WILL BE USED	433
23.11. HOW TO GET ACCESS TO HTTPSERVLETREQUEST AND HTTPSERVLETRESPONSE	433
23.12. CONFIGURING URI TO CALL	433
23.13. CONFIGURING URI PARAMETERS	434
23.14. HOW TO SET THE HTTP METHOD (GET/PATCH/POST/PUT/DELETE/HEAD/OPTIONS/TRACE) TO THE HTTP PRODUCER	434
23.15. USING CLIENT TIMEOUT - SO_TIMEOUT	435
23.16. CONFIGURING A PROXY	435
23.16.1. Using proxy settings outside of URI	435
23.17. CONFIGURING CHARSET	435
23.17.1. Sample with scheduled poll	436
23.17.2. URI Parameters from the endpoint URI	436

23.17.3. URI Parameters from the Message	436
23.17.4. Getting the Response Code	436
23.18. DISABLING COOKIES	436
23.19. BASIC AUTH WITH THE STREAMING MESSAGE BODY	436
23.20. ADVANCED USAGE	437
23.20.1. Setting up SSL for HTTP Client	437
23.21. SPRING BOOT AUTO-CONFIGURATION	439
CHAPTER 24. INFINISPAN	445
24.1. URI FORMAT	445
24.2. CONFIGURING OPTIONS	445
24.2.1. Configuring Component Options	445
24.2.2. Configuring Endpoint Options	445
24.3. COMPONENT OPTIONS	446
24.4. ENDPOINT OPTIONS	450
24.4.1. Path Parameters (1 parameters)	450
24.4.2. Query Parameters (26 parameters)	450
24.5. CAMEL OPERATIONS	453
24.6. MESSAGE HEADERS	457
24.7. EXAMPLES	458
24.8. USING THE INFINISPAN BASED IDEMPOTENT REPOSITORY	459
24.9. USING THE INFINISPAN BASED AGGREGATION REPOSITORY	461
24.10. SPRING BOOT AUTO-CONFIGURATION	462
CHAPTER 25. JIRA	466
25.1. URI FORMAT	466
25.2. CONFIGURING OPTIONS	467
25.2.1. Configuring Component Options	467
25.2.2. Configuring Endpoint Options	467
25.3. COMPONENT OPTIONS	467
25.4. ENDPOINT OPTIONS	469
25.4.1. Path Parameters (1 parameters)	469
25.4.2. Query Parameters (16 parameters)	470
25.5. CLIENT FACTORY	472
25.6. AUTHENTICATION	472
25.6.1. Basic authentication requirements:	473
25.6.2. OAuth authentication requirements:	473
25.7. JQL	473
25.8. OPERATIONS	473
25.9. ADDISSUE	474
25.10. ADDCOMMENT	474
25.11. ATTACH	474
25.12. DELETEISSUE	474
25.13. TRANSITIONISSUE	474
25.14. UPDATEISSUE	475
25.15. WATCHER	475
25.16. WATCHUPDATES (CONSUMER)	475
25.17. SPRING BOOT AUTO-CONFIGURATION	475
CHAPTER 26. JMS	478
26.1. URI FORMAT	478
26.1.1. Using ActiveMQ	479
26.1.2. Transactions and Cache Levels	479
26.1.3. Durable Subscriptions	479

26.1.4. Message Header Mapping	479
26.2. CONFIGURING OPTIONS	480
26.2.1. Configuring Component Options	480
26.2.2. Configuring Endpoint Options	480
26.3. COMPONENT OPTIONS	481
26.4. ENDPOINT OPTIONS	499
26.4.1. Path Parameters (2 parameters)	500
26.4.2. Query Parameters (95 parameters)	500
26.5. SAMPLES	518
26.5.1. Receiving from JMS	518
26.5.2. Sending to JMS	519
26.5.3. Using Annotations	519
26.5.4. Spring DSL sample	519
26.5.5. Other samples	519
26.5.6. Using JMS as a Dead Letter Queue storing Exchange	519
26.5.7. Using JMS as a Dead Letter Channel storing error only	520
26.6. MESSAGE MAPPING BETWEEN JMS AND CAMEL	520
26.6.1. Disabling auto-mapping of JMS messages	521
26.6.2. Using a custom MessageConverter	521
26.6.3. Controlling the mapping strategy selected	521
26.7. MESSAGE FORMAT WHEN SENDING	522
26.8. MESSAGE FORMAT WHEN RECEIVING	522
26.9. ABOUT USING CAMEL TO SEND AND RECEIVE MESSAGES AND JMSREPLYTO	523
26.9.1. JmsProducer	523
26.9.2. JmsConsumer	524
26.10. REUSE ENDPOINT AND SEND TO DIFFERENT DESTINATIONS COMPUTED AT RUNTIME	525
26.11. CONFIGURING DIFFERENT JMS PROVIDERS	525
26.11.1. Using JNDI to find the ConnectionFactory	526
26.12. CONCURRENT CONSUMING	526
26.12.1. Concurrent Consuming with async consumer	526
26.13. REQUEST-REPLY OVER JMS	527
26.13.1. Request-reply over JMS and using a shared fixed reply queue	529
26.13.2. Request-reply over JMS and using an exclusive fixed reply queue	529
26.14. SYNCHRONIZING CLOCKS BETWEEN SENDERS AND RECEIVERS	530
26.15. ABOUT TIME TO LIVE	530
26.16. ENABLING TRANSACTED CONSUMPTION	531
26.17. USING JMSREPLYTO FOR LATE REPLIES	531
26.18. USING A REQUEST TIMEOUT	532
26.19. SENDING AN INONLY MESSAGE AND KEEPING THE JMSREPLYTO HEADER	532
26.20. SETTING JMS PROVIDER OPTIONS ON THE DESTINATION	532
26.21. SPRING BOOT AUTO-CONFIGURATION	533
CHAPTER 27. KAFKA	551
27.1. URI FORMAT	551
27.2. CONFIGURING OPTIONS	551
27.2.1. Configuring Component Options	551
27.2.2. Configuring Endpoint Options	551
27.3. COMPONENT OPTIONS	552
27.4. ENDPOINT OPTIONS	568
27.4.1. Path Parameters (1 parameters)	568
27.4.2. Query Parameters (102 parameters)	568
27.5. MESSAGE HEADERS	585
27.5.1. Consumer headers	585

27.5.2. Producer headers	586
27.6. CONSUMER ERROR HANDLING	587
27.7. SAMPLES	588
27.7.1. Consuming messages from Kafka	588
27.7.2. Producing messages to Kafka	589
27.8. SSL CONFIGURATION	589
27.9. USING THE KAFKA IDEMPOTENT REPOSITORY	590
27.10. USING MANUAL COMMIT WITH KAFKA CONSUMER	593
27.11. KAFKA HEADERS PROPAGATION	593
27.12. SPRING BOOT AUTO-CONFIGURATION	594
CHAPTER 28. KAMELET	611
28.1. URI FORMAT	611
28.2. CONFIGURING OPTIONS	611
28.2.1. Configuring Component Options	611
28.2.2. Configuring Endpoint Options	611
28.3. COMPONENT OPTIONS	611
28.4. ENDPOINT OPTIONS	613
28.4.1. Path Parameters (2 parameters)	613
28.4.2. Query Parameters (8 parameters)	613
28.5. DISCOVERY	614
28.6. SAMPLES	615
28.7. SPRING BOOT AUTO-CONFIGURATION	615
CHAPTER 29. LANGUAGE	618
29.1. URI FORMAT	618
29.2. CONFIGURING OPTIONS	618
29.2.1. Configuring Component Options	618
29.2.2. Configuring Endpoint Options	618
29.3. COMPONENT OPTIONS	619
29.4. ENDPOINT OPTIONS	619
29.4.1. Path Parameters (2 parameters)	619
29.4.2. Query Parameters (7 parameters)	620
29.5. MESSAGE HEADERS	621
29.6. EXAMPLES	622
29.7. LOADING SCRIPTS FROM RESOURCES	622
29.8. SPRING BOOT AUTO-CONFIGURATION	622
CHAPTER 30. LOG	624
30.1. URI FORMAT	624
30.2. CONFIGURING OPTIONS	624
30.2.1. Configuring Component Options	625
30.2.2. Configuring Endpoint Options	625
30.3. COMPONENT OPTIONS	625
30.4. ENDPOINT OPTIONS	626
30.4.1. Path Parameters (1 parameters)	626
30.4.2. Query Parameters (27 parameters)	626
30.5. REGULAR LOGGER SAMPLE	629
30.6. REGULAR LOGGER WITH FORMATTER SAMPLE	629
30.7. THROUGHPUT LOGGER WITH GROUPSIZE SAMPLE	629
30.8. THROUGHPUT LOGGER WITH GROUPINTERVAL SAMPLE	629
30.9. MASKING SENSITIVE INFORMATION LIKE PASSWORD	630
30.10. FULL CUSTOMIZATION OF THE LOGGING OUTPUT	630
30.10.1. Convention over configuration	631

30.11. SPRING BOOT AUTO-CONFIGURATION	631
CHAPTER 31. MAIL	633
31.1. URI FORMAT	633
31.2. CONFIGURING OPTIONS	633
31.2.1. Configuring Component Options	634
31.2.2. Configuring Endpoint Options	634
31.3. COMPONENT OPTIONS	634
31.4. ENDPOINT OPTIONS	639
31.4.1. Path Parameters (2 parameters)	639
31.4.2. Query Parameters (66 parameters)	640
31.4.3. Sample endpoints	647
31.4.4. Component alias names	648
31.4.5. Default ports	648
31.5. SSL SUPPORT	648
31.5.1. Using the JSSE Configuration Utility	648
31.5.2. Configuring JavaMail Directly	649
31.6. MAIL MESSAGE CONTENT	649
31.7. HEADERS TAKE PRECEDENCE OVER PRE-CONFIGURED RECIPIENTS	650
31.8. MULTIPLE RECIPIENTS FOR EASIER CONFIGURATION	650
31.9. SETTING SENDER NAME AND EMAIL	650
31.10. JAVAMAIL API (EX SUN JAVAMAIL)	650
31.11. SAMPLES	651
31.12. SENDING MAIL WITH ATTACHMENT SAMPLE	651
31.13. SSL SAMPLE	651
31.14. CONSUMING MAILS WITH ATTACHMENT SAMPLE	652
31.15. HOW TO SPLIT A MAIL MESSAGE WITH ATTACHMENTS	652
31.16. USING CUSTOM SEARCHTERM	653
31.17. POLLING OPTIMIZATION	654
31.18. USING HEADERS WITH ADDITIONAL JAVA MAIL SENDER PROPERTIES	654
31.19. SPRING BOOT AUTO-CONFIGURATION	655
CHAPTER 32. MASTER	661
32.1. USING THE MASTER ENDPOINT	661
32.2. URI FORMAT	661
32.3. CONFIGURING OPTIONS	661
32.3.1. Configuring Component Options	661
32.3.2. Configuring Endpoint Options	662
32.4. COMPONENT OPTIONS	662
32.5. ENDPOINT OPTIONS	662
32.5.1. Path Parameters (2 parameters)	663
32.5.2. Query Parameters (3 parameters)	663
32.6. EXAMPLE	663
32.7. IMPLEMENTATIONS	664
32.8. SPRING BOOT AUTO-CONFIGURATION	665
CHAPTER 33. MLLP	667
33.1. CONFIGURING OPTIONS	667
33.1.1. Configuring Component Options	667
33.1.2. Configuring Endpoint Options	668
33.2. COMPONENT OPTIONS	668
33.3. ENDPOINT OPTIONS	671
33.3.1. Path Parameters (2 parameters)	671
33.3.2. Query Parameters (26 parameters)	672

33.4. MLLP CONSUMER	675
33.4.1. Message Headers	675
33.4.2. Exchange Properties	676
33.5. MLLP PRODUCER	677
33.5.1. Message Headers	677
33.5.2. Exchange Properties	677
33.6. SPRING BOOT AUTO-CONFIGURATION	678
CHAPTER 34. MOCK	682
34.1. URI FORMAT	682
34.2. CONFIGURING OPTIONS	682
34.2.1. Configuring Component Options	683
34.2.2. Configuring Endpoint Options	683
34.3. COMPONENT OPTIONS	683
34.4. ENDPOINT OPTIONS	684
34.4.1. Path Parameters (1 parameters)	684
34.4.2. Query Parameters (12 parameters)	684
34.5. SIMPLE EXAMPLE	687
34.6. USING ASSERTPERIOD	687
34.7. SETTING EXPECTATIONS	688
34.8. ADDING EXPECTATIONS TO SPECIFIC MESSAGES	689
34.9. MOCKING EXISTING ENDPOINTS	689
34.10. MOCKING EXISTING ENDPOINTS USING THE CAMEL-TEST COMPONENT	691
34.11. MOCKING EXISTING ENDPOINTS WITH XML DSL	692
34.12. MOCKING ENDPOINTS AND SKIP SENDING TO ORIGINAL ENDPOINT	693
34.13. LIMITING THE NUMBER OF MESSAGES TO KEEP	694
34.14. TESTING WITH ARRIVAL TIMES	695
34.15. SPRING BOOT AUTO-CONFIGURATION	695
CHAPTER 35. MONGODB	697
35.1. URI FORMAT	697
35.2. CONFIGURING OPTIONS	697
35.2.1. Configuring Component Options	697
35.2.2. Configuring Endpoint Options	698
35.3. COMPONENT OPTIONS	698
35.4. ENDPOINT OPTIONS	699
35.4.1. Path Parameters (1 parameters)	699
35.4.2. Query Parameters (27 parameters)	699
35.5. CONFIGURATION OF DATABASE IN SPRING XML	704
35.6. SAMPLE ROUTE	705
35.7. MONGODB OPERATIONS - PRODUCER ENDPOINTS	705
35.7.1. Query operations	705
35.7.1.1. findById	705
35.7.1.2. findOneByQuery	705
35.7.1.3. Example without a query selector (returns the first document in a collection)	706
35.7.1.4. Example with a query selector (returns the first matching document in a collection):	706
35.7.1.5. findAll	706
35.7.1.5.1. Example without a query selector (returns all documents in a collection)	706
35.7.1.5.2. Example with a query selector (returns all matching documents in a collection)	706
35.7.1.5.3. Example with option outputType=Mongolterable and batch size	707
35.7.1.6. count	708
35.7.1.7. Specifying a fields filter (projection)	708
35.7.1.8. Specifying a sort clause	709

35.7.2. Create/update operations	709
35.7.2.1. insert	709
35.7.2.2. save	710
35.7.2.3. update	711
35.7.3. Delete operations	712
35.7.3.1. remove	712
35.7.4. Bulk Write Operations	713
35.7.4.1. bulkWrite	713
35.7.5. Other operations	713
35.7.5.1. aggregate	713
35.7.5.2. getDbStats	714
35.7.5.3. getColStats	715
35.7.5.4. command	716
35.7.6. Dynamic operations	716
35.8. CONSUMERS	716
35.8.1. Tailable Cursor Consumer	716
35.9. HOW THE TAILABLE CURSOR CONSUMER WORKS	717
35.10. PERSISTENT TAIL TRACKING	717
35.11. ENABLING PERSISTENT TAIL TRACKING	718
35.11.1. Change Streams Consumer	718
35.12. TYPE CONVERSIONS	719
35.13. SPRING BOOT AUTO-CONFIGURATION	720
CHAPTER 36. NETTY	722
36.1. URI FORMAT	722
36.2. CONFIGURING OPTIONS	722
36.2.1. Configuring Component Options	722
36.2.2. Configuring Endpoint Options	723
36.3. COMPONENT OPTIONS	723
36.4. ENDPOINT OPTIONS	732
36.4.1. Path Parameters (3 parameters)	732
36.4.2. Query Parameters (71 parameters)	732
36.5. REGISTRY BASED OPTIONS	741
36.5.1. Using non shareable encoders or decoders	742
36.6. SENDING MESSAGES TO/FROM A NETTY ENDPOINT	742
36.6.1. Netty Producer	742
36.6.2. Netty Consumer	743
36.7. EXAMPLES	743
36.7.1. A UDP Netty endpoint using Request-Reply and serialized object payload	743
36.7.2. A TCP based Netty consumer endpoint using One-way communication	743
36.7.3. An SSL/TCP based Netty consumer endpoint using Request-Reply communication	744
36.7.4. Using Multiple Codecs	745
36.8. CLOSING CHANNEL WHEN COMPLETE	747
36.9. CUSTOM PIPELINE	747
36.9.1. Using custom pipeline factory	748
36.10. REUSING NETTY BOSS AND WORKER THREAD POOLS	749
36.11. MULTIPLEXING CONCURRENT MESSAGES OVER A SINGLE CONNECTION WITH REQUEST/REPLY	750
36.12. SPRING BOOT AUTO-CONFIGURATION	750
CHAPTER 37. PAHO	760
37.1. URI FORMAT	760
37.2. CONFIGURING OPTIONS	760

37.2.1. Configuring Component Options	760
37.2.2. Configuring Endpoint Options	760
37.3. COMPONENT OPTIONS	761
37.4. ENDPOINT OPTIONS	766
37.4.1. Path Parameters (1 parameters)	766
37.4.2. Query Parameters (31 parameters)	767
37.5. HEADERS	772
37.6. DEFAULT PAYLOAD TYPE	773
37.7. SAMPLES	773
37.8. SPRING BOOT AUTO-CONFIGURATION	774
CHAPTER 38. PAHO MQTT 5	781
38.1. URI FORMAT	781
38.2. CONFIGURING OPTIONS	781
38.2.1. Configuring Component Options	781
38.2.2. Configuring Endpoint Options	781
38.3. COMPONENT OPTIONS	782
38.4. ENDPOINT OPTIONS	788
38.4.1. Path Parameters (1 parameters)	788
38.4.2. Query Parameters (32 parameters)	788
38.5. HEADERS	795
38.6. DEFAULT PAYLOAD TYPE	795
38.7. SAMPLES	796
38.8. SPRING BOOT AUTO-CONFIGURATION	796
CHAPTER 39. QUARTZ	805
39.1. URI FORMAT	805
39.2. CONFIGURING OPTIONS	805
39.2.1. Configuring Component Options	805
39.2.2. Configuring Endpoint Options	805
39.3. COMPONENT OPTIONS	806
39.4. ENDPOINT OPTIONS	807
39.4.1. Path Parameters (2 parameters)	807
39.4.2. Query Parameters (17 parameters)	808
39.4.3. Configuring quartz.properties file	810
39.5. ENABLING QUARTZ SCHEDULER IN JMX	810
39.6. STARTING THE QUARTZ SCHEDULER	810
39.7. CLUSTERING	810
39.8. MESSAGE HEADERS	811
39.9. USING CRON TRIGGERS	811
39.10. SPECIFYING TIME ZONE	811
39.11. CONFIGURING MISFIRE INSTRUCTIONS	811
39.11.1. SimpleTrigger.MISFIRE_INSTRUCTION_FIRE_NOW = 1 (default)	812
39.11.2. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NOW_WITH_EXISTING_REPEAT_COUNT = 2	812
39.11.3. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NOW_WITH_REMAINING_REPEAT_COUNT = 3	812
39.11.4. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NEXT_WITH_REMAINING_COUNT = 4	812
39.11.5. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NEXT_WITH_EXISTING_COUNT = 5	813
39.11.6. CronTrigger.MISFIRE_INSTRUCTION_FIRE_ONCE_NOW = 1 (default)	813
39.11.7. CronTrigger.MISFIRE_INSTRUCTION_DO_NOTHING = 2	813
39.12. USING QUARTZSCHEDULEDPOLLCONSUMERSCHEDULER	813
39.13. CRON COMPONENT SUPPORT	814

39.14. SPRING BOOT AUTO-CONFIGURATION	814
CHAPTER 40. REF	817
40.1. URI FORMAT	817
40.2. CONFIGURING OPTIONS	817
40.2.1. Configuring Component Options	817
40.2.2. Configuring Endpoint Options	817
40.3. COMPONENT OPTIONS	817
40.4. ENDPOINT OPTIONS	818
40.4.1. Path Parameters (1 parameters)	818
40.4.2. Query Parameters (4 parameters)	818
40.5. RUNTIME LOOKUP	819
40.6. SAMPLE	820
40.7. SPRING BOOT AUTO-CONFIGURATION	820
CHAPTER 41. REST	822
41.1. URI FORMAT	822
41.2. CONFIGURING OPTIONS	822
41.2.1. Configuring Component Options	822
41.2.2. Configuring Endpoint Options	822
41.3. COMPONENT OPTIONS	822
41.4. ENDPOINT OPTIONS	824
41.4.1. Path Parameters (3 parameters)	824
41.4.2. Query Parameters (16 parameters)	825
41.5. SUPPORTED REST COMPONENTS	827
41.6. PATH AND URITEMPLATE SYNTAX	828
41.7. REST PRODUCER EXAMPLES	828
41.8. REST PRODUCER BINDING	829
41.9. MORE EXAMPLES	830
41.10. SPRING BOOT AUTO-CONFIGURATION	830
CHAPTER 42. SAGA	833
42.1. URI FORMAT	833
42.2. CONFIGURING OPTIONS	833
42.2.1. Configuring Component Options	833
42.2.2. Configuring Endpoint Options	833
42.3. COMPONENT OPTIONS	833
42.4. ENDPOINT OPTIONS	834
42.4.1. Path Parameters (1 parameters)	834
42.4.2. Query Parameters (1 parameters)	834
42.5. SPRING BOOT AUTO-CONFIGURATION	835
CHAPTER 43. SALESFORCE	837
43.1. CONFIGURING OPTIONS	837
43.1.1. Configuring Component Options	837
43.1.2. Configuring Endpoint Options	837
43.2. COMPONENT OPTIONS	838
43.3. ENDPOINT OPTIONS	847
43.3.1. Path Parameters (2 parameters)	847
43.3.2. Query Parameters (57 parameters)	849
43.4. AUTHENTICATING TO SALESFORCE	855
43.5. URI FORMAT	856
43.6. PASSING IN SALESFORCE HEADERS AND FETCHING SALESFORCE RESPONSE HEADERS	856
43.7. SUPPORTED SALESFORCE APIS	857

43.7.1. Rest API	857
43.7.2. Bulk 2.0 API	858
43.7.3. Rest Bulk (original) API	859
43.7.4. Rest Streaming API	860
43.7.5. Platform events	861
43.7.6. Change data capture events	862
43.8. EXAMPLES	862
43.8.1. Uploading a document to a ContentWorkspace	862
43.9. USING SALESFORCE LIMITS API	863
43.10. WORKING WITH APPROVALS	863
43.11. USING SALESFORCE RECENT ITEMS API	864
43.12. USING SALESFORCE COMPOSITE API TO SUBMIT SUBJECT TREE	864
43.13. USING SALESFORCE COMPOSITE API TO SUBMIT MULTIPLE REQUESTS IN A BATCH	865
43.14. USING SALESFORCE COMPOSITE API TO SUBMIT MULTIPLE CHAINED REQUESTS	866
43.15. USING "RAW" SALESFORCE COMPOSITE	867
43.16. USING RAW OPERATION	868
43.16.1. Query example	869
43.16.2. SObject example	869
43.17. USING COMPOSITE SUBJECT COLLECTIONS	869
43.17.1. compositeRetrieveSObjectCollections	869
43.17.2. compositeCreateSObjectCollections	870
43.17.3. compositeUpdateSObjectCollections	870
43.17.4. compositeUpsertSObjectCollections	871
43.17.5. compositeDeleteSObjectCollections	871
43.18. SENDING NULL VALUES TO SALESFORCE	872
43.19. GENERATING SOQL QUERY STRINGS	872
43.20. CAMEL SALESFORCE MAVEN PLUGIN	872
43.21. SPRING BOOT AUTO-CONFIGURATION	872
CHAPTER 44. SCHEDULER	883
44.1. URI FORMAT	883
44.2. CONFIGURING OPTIONS	883
44.2.1. Configuring Component Options	883
44.2.2. Configuring Endpoint Options	883
44.3. COMPONENT OPTIONS	884
44.4. ENDPOINT OPTIONS	884
44.4.1. Path Parameters (1 parameters)	884
44.4.2. Query Parameters (21 parameters)	885
44.5. MORE INFORMATION	887
44.6. EXCHANGE PROPERTIES	887
44.7. SAMPLE	888
44.8. FORCING THE SCHEDULER TO TRIGGER IMMEDIATELY WHEN COMPLETED	888
44.9. FORCING THE SCHEDULER TO BE IDLE	888
44.10. SPRING BOOT AUTO-CONFIGURATION	888
CHAPTER 45. SEDA	890
45.1. URI FORMAT	890
45.2. CONFIGURING OPTIONS	890
45.2.1. Configuring Component Options	890
45.2.2. Configuring Endpoint Options	890
45.3. COMPONENT OPTIONS	891
45.4. ENDPOINT OPTIONS	892
45.4.1. Path Parameters (1 parameters)	892

45.4.2. Query Parameters (18 parameters)	893
45.5. CHOOSING BLOCKINGQUEUE IMPLEMENTATION	895
45.6. USE OF REQUEST REPLY	896
45.7. CONCURRENT CONSUMERS	896
45.8. THREAD POOLS	896
45.9. SAMPLE	897
45.10. USING MULTIPLECONSUMERS	897
45.11. EXTRACTING QUEUE INFORMATION	898
45.12. SPRING BOOT AUTO-CONFIGURATION	898
CHAPTER 46. SERVLET	901
46.1. URI FORMAT	901
46.2. CONFIGURING OPTIONS	901
46.2.1. Configuring Component Options	901
46.2.2. Configuring Endpoint Options	902
46.3. COMPONENT OPTIONS	902
46.4. ENDPOINT OPTIONS	903
46.4.1. Path Parameters (1 parameters)	903
46.4.2. Query Parameters (22 parameters)	904
46.5. MESSAGE HEADERS	906
46.6. USAGE	906
46.7. SPRING BOOT AUTO-CONFIGURATION	907
CHAPTER 47. SLACK	910
47.1. URI FORMAT	910
47.2. CONFIGURING OPTIONS	910
47.2.1. Configuring Component Options	910
47.2.2. Configuring Endpoint Options	910
47.3. COMPONENT OPTIONS	911
47.4. ENDPOINT OPTIONS	911
47.4.1. Path Parameters (1 parameters)	912
47.4.2. Query Parameters (29 parameters)	912
47.5. CONFIGURING IN SPRINT XML	916
47.6. EXAMPLE	916
47.7. PRODUCER	916
47.8. CONSUMER	917
47.9. SPRING BOOT AUTO-CONFIGURATION	917
CHAPTER 48. SQL	919
48.1. URI FORMAT	919
48.2. CONFIGURING OPTIONS	920
48.2.1. Configuring Component Options	920
48.2.2. Configuring Endpoint Options	920
48.3. COMPONENT OPTIONS	921
48.4. ENDPOINT OPTIONS	921
48.4.1. Path Parameters (1 parameters)	922
48.4.2. Query Parameters (45 parameters)	922
48.5. TREATMENT OF THE MESSAGE BODY	928
48.6. RESULT OF THE QUERY	928
48.7. USING STREAMLIST	929
48.8. HEADER VALUES	929
48.9. GENERATED KEYS	930
48.10. DATASOURCE	930
48.11. USING NAMED PARAMETERS	930

48.12. USING EXPRESSION PARAMETERS IN PRODUCERS	930
48.12.1. Using expression parameters in consumers	931
48.13. USING IN QUERIES WITH DYNAMIC VALUES	931
48.14. USING THE JDBC BASED IDEMPOTENT REPOSITORY	932
48.14.1. Customize the JDBC idempotency repository	933
48.14.2. Orphan Lock aware Jdbc IdempotentRepository	934
48.14.3. Caching Jdbc IdempotentRepository	934
48.15. USING THE JDBC BASED AGGREGATION REPOSITORY	934
48.15.1. Database	935
48.16. STORING BODY AND HEADERS AS TEXT	935
48.16.1. Codec (Serialization)	936
48.16.2. Transaction	936
48.16.2.1. Service (Start/Stop)	936
48.16.3. Aggregator configuration	936
48.16.4. Optimistic locking	937
48.16.5. Propagation behavior	938
48.16.6. PostgreSQL case	938
48.17. CAMEL SQL STARTER	938
48.18. SPRING BOOT AUTO-CONFIGURATION	939
CHAPTER 49. STUB	941
49.1. URI FORMAT	941
49.2. CONFIGURING OPTIONS	941
49.2.1. Configuring Component Options	941
49.2.1.1. Configuring Endpoint Options	941
49.3. COMPONENT OPTIONS	942
49.4. ENDPOINT OPTIONS	943
49.4.1. Path Parameters (1 parameters)	943
49.4.2. Query Parameters (18 parameters)	943
49.5. EXAMPLES	946
49.6. SPRING BOOT AUTO-CONFIGURATION	946
CHAPTER 50. TELEGRAM	949
50.1. URI FORMAT	949
50.2. CONFIGURING OPTIONS	949
50.2.1. Configuring Component Options	949
50.2.2. Configuring Endpoint Options	950
50.3. COMPONENT OPTIONS	950
50.4. ENDPOINT OPTIONS	951
50.4.1. Path Parameters (1 parameters)	951
50.4.2. Query Parameters (30 parameters)	951
50.4.3. Message Headers	955
50.5. USAGE	956
50.6. PRODUCER EXAMPLE	956
50.7. CONSUMER EXAMPLE	957
50.8. REACTIVE CHAT-BOT EXAMPLE	958
50.9. GETTING THE CHAT ID	959
50.10. CUSTOMIZING KEYBOARD	959
50.11. WEBHOOK MODE	960
50.12. SPRING BOOT AUTO-CONFIGURATION	961
CHAPTER 51. TIMER	963
51.1. URI FORMAT	963
51.2. CONFIGURING OPTIONS	963

51.2.1. Configuring Component Options	963
51.2.2. Configuring Endpoint Options	963
51.3. COMPONENT OPTIONS	964
51.4. ENDPOINT OPTIONS	964
51.4.1. Path Parameters (1 parameters)	964
51.4.2. Query Parameters (13 parameters)	964
51.5. EXCHANGE PROPERTIES	966
51.6. SAMPLE	966
51.7. FIRING AS SOON AS POSSIBLE	967
51.8. FIRING ONLY ONCE	967
51.9. SPRING BOOT AUTO-CONFIGURATION	967
CHAPTER 52. VALIDATOR	969
52.1. URI FORMAT	969
52.2. CONFIGURING OPTIONS	969
52.2.1. Configuring Component Options	969
52.2.2. Configuring Endpoint Options	970
52.3. COMPONENT OPTIONS	970
52.4. ENDPOINT OPTIONS	970
52.4.1. Path Parameters (1 parameters)	971
52.4.2. Query Parameters (10 parameters)	971
52.5. EXAMPLE	972
52.6. ADVANCED: JMX METHOD CLEARCACHEDSCHEMA	972
52.7. SPRING BOOT AUTO-CONFIGURATION	972
CHAPTER 53. WEBHOOK	974
53.1. URI FORMAT	974
53.2. CONFIGURING OPTIONS	974
53.2.1. Configuring Component Options	974
53.2.2. Configuring Endpoint Options	974
53.3. COMPONENT OPTIONS	975
53.4. ENDPOINT OPTIONS	976
53.4.1. Path Parameters (1 parameters)	976
53.4.2. Query Parameters (8 parameters)	976
53.5. EXAMPLES	977
53.6. SPRING BOOT AUTO-CONFIGURATION	977
CHAPTER 54. XSLT	979
54.1. URI FORMAT	979
54.2. CONFIGURING OPTIONS	979
54.2.1. Configuring Component Options	979
54.2.2. Configuring Endpoint Options	980
54.3. COMPONENT OPTIONS	980
54.4. ENDPOINT OPTIONS	981
54.4.1. Path Parameters (1 parameters)	981
54.4.2. Query Parameters (13 parameters)	981
54.5. USING XSLT ENDPOINTS	983
54.6. GETTING USEABLE PARAMETERS INTO THE XSLT	983
54.7. SPRING XML VERSIONS	984
54.8. USING XSL:INCLUDE	984
54.9. USING XSL:INCLUDE AND DEFAULT PREFIX	984
54.10. DYNAMIC STYLESHEETS	985
54.11. ACCESSING WARNINGS, ERRORS AND FATALERRORS FROM XSLT ERRORLISTENER	985
54.12. SPRING BOOT AUTO-CONFIGURATION	985

CHAPTER 55. AVRO	988
55.1. AVRO DATAFORMAT OPTIONS	988
55.2. AVRO DATA FORMAT USAGE	988
55.3. SPRING BOOT AUTO-CONFIGURATION	989
CHAPTER 56. AVRO JACKSON	990
56.1. CONFIGURING THE SCHEMARESOLVER	990
56.2. AVRO JACKSON OPTIONS	990
56.3. USING CUSTOM AVROMAPPER	992
56.4. DEPENDENCIES	992
56.5. SPRING BOOT AUTO-CONFIGURATION	992
CHAPTER 57. BINDY	996
57.1. OPTIONS	997
57.2. ANNOTATIONS	997
57.2.1. 1. CsvRecord	998
57.2.2. 2. Link	1002
57.2.3. 3. DataField	1003
57.2.4. 4. FixedLengthRecord	1009
57.2.5. 5. Message	1016
57.2.6. 6. KeyValuePairField	1018
57.2.7. 7. Section	1020
57.2.8. 8. OneToMany	1021
57.2.9. 9. BindyConverter	1023
57.2.10. 10. FormatFactories	1024
57.3. SUPPORTED DATATYPES	1025
57.4. USING THE JAVA DSL	1025
57.4.1. Setting locale	1026
57.4.2. Unmarshaling	1026
57.4.3. Marshaling	1027
57.5. USING SPRING XML	1027
57.6. DEPENDENCIES	1028
57.7. SPRING BOOT AUTO-CONFIGURATION	1029
CHAPTER 58. HL7	1031
58.1. HL7 MLLP PROTOCOL	1031
58.1.1. Exposing an HL7 listener using Mina	1032
58.1.2. Exposing an HL7 listener using Netty (available from Camel 2.15 onwards)	1032
58.2. HL7 MODEL USING JAVA.LANG.STRING OR BYTE[]	1033
58.3. HL7V2 MODEL USING HAPI	1033
58.4. HL7 DATAFORMAT	1033
58.4.1. Segment separators	1034
58.4.2. Charset	1034
58.5. MESSAGE HEADERS	1035
58.6. DEPENDENCIES	1036
58.7. SPRING BOOT AUTO-CONFIGURATION	1037
CHAPTER 59. JACKSONXML	1038
59.1. JACKSONXML OPTIONS	1038
59.1.1. Using Jackson XML in Spring DSL	1039
59.1.2. Excluding POJO fields from marshalling	1040
59.2. INCLUDE/EXCLUDE FIELDS USING THE JSONVIEW ATTRIBUTE WITH `JACKSONXML` DATAFORMAT	1040
59.3. SETTING SERIALIZATION INCLUDE OPTION	1040

59.4. UNMARSHALLING FROM XML TO POJO WITH DYNAMIC CLASS NAME	1041
59.5. UNMARSHALLING FROM XML TO LIST<MAP> OR LIST<POJO>	1041
59.6. USING CUSTOM JACKSON MODULES	1042
59.7. ENABLING OR DISABLE FEATURES USING JACKSON	1042
59.8. CONVERTING MAPS TO POJO USING JACKSON	1043
59.9. FORMATTED XML MARSHALLING (PRETTY-PRINTING)	1043
59.10. DEPENDENCIES	1043
59.11. SPRING BOOT AUTO-CONFIGURATION	1044
CHAPTER 60. JAXB	1047
60.1. OPTIONS	1047
60.2. USING THE JAVA DSL	1049
60.3. USING SPRING XML	1049
60.4. PARTIAL MARSHALLING/UNMARSHALLING	1050
60.5. FRAGMENT	1050
60.6. IGNORING THE NONXML CHARACTER	1050
60.7. WORKING WITH THE OBJECTFACTORY	1051
60.8. SETTING ENCODING	1051
60.9. CONTROLLING NAMESPACE PREFIX MAPPING	1051
60.10. SCHEMA VALIDATION	1052
60.11. SCHEMA LOCATION	1052
60.12. MARSHAL DATA THAT IS ALREADY XML	1052
60.13. DEPENDENCIES	1053
60.14. SPRING BOOT AUTO-CONFIGURATION	1053
CHAPTER 61. JSON GSON	1056
61.1. GSON OPTIONS	1056
61.2. DEPENDENCIES	1056
61.3. SPRING BOOT AUTO-CONFIGURATION	1056
CHAPTER 62. JSON JACKSON	1058
62.1. JACKSON OPTIONS	1058
62.2. USING CUSTOM OBJECTMAPPER	1063
62.3. USING JACKSON FOR AUTOMATIC TYPE CONVERSION	1063
62.4. DEPENDENCIES	1064
62.5. SPRING BOOT AUTO-CONFIGURATION	1064
CHAPTER 63. PROTOBUF JACKSON	1068
63.1. CONFIGURING THE SCHEMARESOLVER	1068
63.2. PROTOBUF JACKSON OPTIONS	1068
63.3. USING CUSTOM PROTOBUFMAPPER	1070
63.4. DEPENDENCIES	1070
63.5. SPRING BOOT AUTO-CONFIGURATION	1071
CHAPTER 64. SOAP	1074
64.1. SOAP OPTIONS	1074
64.2. ELEMENTNAMESTRATEGY	1075
64.3. USING THE JAVA DSL	1075
64.3.1. Using SOAP 1.2	1076
64.4. MULTI-PART MESSAGES	1076
64.4.1. Holder Object mapping	1077
64.5. EXAMPLES	1077
64.5.1. Webservice client	1077
64.5.2. Webservice Server	1077

64.6. DEPENDENCIES	1078
64.7. SPRING BOOT AUTO-CONFIGURATION	1078
CHAPTER 65. ZIP FILE	1080
65.1. ZIPFILE OPTIONS	1080
65.2. MARSHAL	1080
65.3. UNMARSHAL	1081
65.3.1. Aggregate	1081
65.4. DEPENDENCIES	1082
65.5. SPRING BOOT AUTO-CONFIGURATION	1082
CHAPTER 66. CONSTANT	1084
66.1. CONSTANT OPTIONS	1084
66.2. EXAMPLE	1084
66.2.1. Specifying type of value	1084
66.3. LOADING CONSTANT FROM EXTERNAL RESOURCE	1085
66.4. DEPENDENCIES	1085
66.5. SPRING BOOT AUTO-CONFIGURATION	1085
CHAPTER 67. CSIMPLE	1103
67.1. DIFFERENT BETWEEN CSIMPLE AND SIMPLE	1103
67.1.1. Additional CSimple functions	1103
67.2. COMPILATION	1104
67.2.1. Using camel-csimple-maven-plugin	1104
67.2.2. Using camel-csimple-joor	1105
67.3. CSIMPLE LANGUAGE OPTIONS	1106
67.4. LIMITATIONS	1106
67.5. AUTO IMPORTS	1106
67.6. CONFIGURATION FILE	1107
67.7. SEE ALSO	1107
67.8. SPRING BOOT AUTO-CONFIGURATION	1107
CHAPTER 68. EXCHANGEPROPERTY	1125
68.1. EXCHANGE PROPERTY OPTIONS	1125
68.2. EXAMPLE	1125
68.3. DEPENDENCIES	1125
68.4. SPRING BOOT AUTO-CONFIGURATION	1125
CHAPTER 69. FILE	1143
69.1. FILE LANGUAGE OPTIONS	1143
69.2. SYNTAX	1143
69.3. FILE TOKEN EXAMPLE	1145
69.3.1. Relative paths	1145
69.3.2. Absolute paths	1146
69.4. SAMPLES	1147
69.5. DEPENDENCIES	1147
69.6. SPRING BOOT AUTO-CONFIGURATION	1148
CHAPTER 70. HEADER	1165
70.1. HEADER OPTIONS	1165
70.2. EXAMPLE USAGE	1165
70.3. DEPENDENCIES	1165
70.4. SPRING BOOT AUTO-CONFIGURATION	1165
CHAPTER 71. JSONPATH	1183

71.1. JSONPATH OPTIONS	1183
71.2. EXAMPLES	1183
71.3. JSONPATH SYNTAX	1184
71.3.1. Easy JSONPath Syntax	1184
71.4. SUPPORTED MESSAGE BODY TYPES	1185
71.5. SUPPRESSING EXCEPTIONS	1185
71.6. INLINE SIMPLE EXPRESSIONS	1186
71.7. JSONPATH INJECTION	1187
71.8. ENCODING DETECTION	1187
71.9. SPLIT JSON DATA INTO SUB ROWS AS JSON	1187
71.10. USING HEADER AS INPUT	1187
71.11. SPRING BOOT AUTO-CONFIGURATION	1188
CHAPTER 72. REF	1190
72.1. REF LANGUAGE OPTIONS	1190
72.2. EXAMPLE USAGE	1190
72.3. DEPENDENCIES	1190
72.4. SPRING BOOT AUTO-CONFIGURATION	1190
CHAPTER 73. XQUERY	1208
73.1. XQUERY LANGUAGE OPTIONS	1208
73.2. VARIABLES	1208
73.3. EXAMPLE	1209
73.3.1. Using namespaces	1209
73.4. USING XQUERY AS TRANSFORMATION	1210
73.5. LOADING SCRIPT FROM EXTERNAL RESOURCE	1211
73.6. LEARNING XQUERY	1211
73.7. DEPENDENCIES	1211
73.8. SPRING BOOT AUTO-CONFIGURATION	1211
CHAPTER 74. SIMPLE	1214
74.1. SIMPLE LANGUAGE OPTIONS	1214
74.2. VARIABLES	1214
74.3. OGNL EXPRESSION SUPPORT	1218
74.4. OPERATOR SUPPORT	1220
74.4.1. Comparing with different types	1223
74.4.2. Using and / or	1224
74.5. EXAMPLES	1224
74.6. SETTING RESULT TYPE	1226
74.7. USING NEW LINES OR TABS IN XML DSLS	1226
74.8. LEADING AND TRAILING WHITESPACE HANDLING	1226
74.9. LOADING SCRIPT FROM EXTERNAL RESOURCE	1227
74.10. SPRING BOOT AUTO-CONFIGURATION	1227
CHAPTER 75. TOKENIZE	1245
75.1. TOKENIZE OPTIONS	1245
75.2. EXAMPLE	1246
75.3. SEE ALSO	1246
75.4. SPRING BOOT AUTO-CONFIGURATION	1246
CHAPTER 76. XML TOKENIZE	1264
76.1. XML TOKENIZER OPTIONS	1264
76.2. EXAMPLE	1265
76.3. SPRING BOOT AUTO-CONFIGURATION	1265

CHAPTER 77. XPATH	1266
77.1. XPATH LANGUAGE OPTIONS	1266
77.2. NAMESPACES	1267
77.3. VARIABLES	1267
77.3.1. Namespace given	1268
77.3.2. No namespace given	1268
77.4. FUNCTIONS	1268
77.4.1. Functions example	1269
77.5. STREAM BASED MESSAGE BODIES	1269
77.6. SETTING RESULT TYPE	1270
77.7. USING XPATH ON HEADERS	1270
77.8. EXAMPLE	1270
77.9. USING NAMESPACES	1271
77.10. USING @XPATH ANNOTATION FOR BEAN INTEGRATION	1272
77.11. USING XPATHBUILDER WITHOUT AN EXCHANGE	1272
77.12. USING SAXON WITH XPATHBUILDER	1273
77.12.1. Setting a custom XPathFactory using System Property	1273
77.12.2. Enabling Saxon from XML DSL	1273
77.13. NAMESPACE AUDITING TO AID DEBUGGING	1273
77.13.1. Logging the Namespace Context of your XPath expression/predicate	1274
77.13.2. Auditing namespaces	1274
77.14. LOADING SCRIPT FROM EXTERNAL RESOURCE	1275
77.15. DEPENDENCIES	1275
77.16. SPRING BOOT AUTO-CONFIGURATION	1275
CHAPTER 78. OPENAPI JAVA	1277
78.1. USING OPENAPI IN REST-DSL	1277
78.2. OPTIONS	1277
78.3. ADDING SECURITY DEFINITIONS IN API DOC	1278
78.4. JSON OR YAML	1279
78.5. USEXFORWARDHEADERS AND API URL RESOLUTION	1279
78.6. EXAMPLES	1280
78.7. SPRING BOOT AUTO-CONFIGURATION	1280

PREFACE

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

CHAPTER 1. AWS CLOUDWATCH

Only producer is supported

The AWS2 Cloudwatch component allows messages to be sent to an [Amazon CloudWatch](#) metrics. The implementation of the Amazon API is provided by the [AWS SDK](#).

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon CloudWatch. More information is available at [Amazon CloudWatch](#).

1.1. URI FORMAT

```
aws2-cw://namespace[?options]
```

The metrics will be created if they don't already exist. You can append query options to the URI in the following format, **?options=value&option2=value&...**

1.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

1.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

1.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

1.3. COMPONENT OPTIONS

The AWS CloudWatch component supports 18 options, which are listed below.

Name	Description	Default	Type
amazonCwClient (producer)	Autowired To use the AmazonCloudWatch as the client.		CloudWatchClient
configuration (producer)	The component configuration.		Cw2Configuration
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
name (producer)	The metric name.		String
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (producer)	To define a proxy host when instantiating the CW client.		String
proxyPort (producer)	To define a proxy port when instantiating the CW client.		Integer
proxyProtocol (producer)	To define a proxy protocol when instantiating the CW client. Enum values: <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol

Name	Description	Default	Type
region (producer)	The region in which CW client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
timestamp (producer)	The metric timestamp.		Instant
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
unit (producer)	The metric unit.		String
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
value (producer)	The metric value.		Double
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as <code>autowired</code>) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

1.4. ENDPOINT OPTIONS

The AWS CloudWatch endpoint is configured using URI syntax:

```
aws2-cw:namespace
```

with the following path and query parameters:

1.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
namespace (producer)	Required The metric namespace.		String

1.4.2. Query Parameters (16 parameters)

Name	Description	Default	Type
amazonCwClient (producer)	Autowired To use the AmazonCloudWatch as the client.		CloudWatchClient
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
name (producer)	The metric name.		String
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (producer)	To define a proxy host when instantiating the CW client.		String
proxyPort (producer)	To define a proxy port when instantiating the CW client.		Integer
proxyProtocol (producer)	To define a proxy protocol when instantiating the CW client. Enum values: <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol

Name	Description	Default	Type
region (producer)	The region in which CW client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
timestamp (producer)	The metric timestamp.		Instant
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
unit (producer)	The metric unit.		String
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
value (producer)	The metric value.		Double
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required CW component options

You have to provide the `amazonCwClient` in the Registry or your `accessKey` and `secretKey` to access the [Amazon's CloudWatch](#).

1.5. USAGE

1.5.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`
- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.
- Web Identity Token from AWS STS.
- The shared credentials and config files.

- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

1.5.2. Message headers evaluated by the CW producer

Header	Type	Description
CamelAwsCwMetricName	String	The Amazon CW metric name.
CamelAwsCwMetricValue	Double	The Amazon CW metric value.
CamelAwsCwMetricUnit	String	The Amazon CW metric unit.
CamelAwsCwMetricName space	String	The Amazon CW metric namespace.
CamelAwsCwMetricTimes tamp	Date	The Amazon CW metric timestamp.
CamelAwsCwMetricDime nsionName	String	The Amazon CW metric dimension name.
CamelAwsCwMetricDime nsionValue	String	The Amazon CW metric dimension value.
CamelAwsCwMetricDime nsions	Map<String, String>	A map of dimension names and dimension values.

1.5.3. Advanced CloudWatchClient configuration

If you need more control over the **CloudWatchClient** instance configuration you can create your own instance and refer to it from the URI:

```
from("direct:start")
.to("aws2-cw://namespace?amazonCwClient=#client");
```

The **#client** refers to a **CloudWatchClient** in the Registry.

1.6. DEPENDENCIES

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```
<dependency>
  <groupId>org.apache.camel</groupId>
```

```
<artifactId>camel-aws2-cw</artifactId>
<version>${camel-version}</version>
</dependency>
```

where **{camel-version}** must be replaced by the actual version of Camel.

1.7. EXAMPLES

1.7.1. Producer Example

```
from("direct:start")
.to("aws2-cw://http://camel.apache.org/aws-cw");
```

and sends something like

```
exchange.getIn().setHeader(Cw2Constants.METRIC_NAME, "ExchangesCompleted");
exchange.getIn().setHeader(Cw2Constants.METRIC_VALUE, "2.0");
exchange.getIn().setHeader(Cw2Constants.METRIC_UNIT, "Count");
```

1.8. SPRING BOOT AUTO-CONFIGURATION

When using `aws2-cw` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-aws2-cw-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 19 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.aws2-cw.access-key</code>	Amazon AWS Access Key.		String
<code>camel.component.aws2-cw.amazon-cw-client</code>	To use the AmazonCloudWatch as the client. The option is a <code>software.amazon.awssdk.services.cloudwatch.CloudWatchClient</code> type.		CloudWatchClient

Name	Description	Default	Type
<code>camel.component.aws2-cw.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-cw.configuration</code>	The component configuration. The option is a <code>org.apache.camel.component.aws2.cw.Cw2Configuration</code> type.		Cw2Configuration
<code>camel.component.aws2-cw.enabled</code>	Whether to enable auto configuration of the <code>aws2-cw</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-cw.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-cw.name</code>	The metric name.		String
<code>camel.component.aws2-cw.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-cw.proxy-host</code>	To define a proxy host when instantiating the CW client.		String
<code>camel.component.aws2-cw.proxy-port</code>	To define a proxy port when instantiating the CW client.		Integer
<code>camel.component.aws2-cw.proxy-protocol</code>	To define a proxy protocol when instantiating the CW client.		Protocol

Name	Description	Default	Type
<code>camel.component.aws2-cw.region</code>	The region in which CW client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-cw.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-cw.timestamp</code>	The metric timestamp. The option is a <code>java.time.Instant</code> type.		Instant
<code>camel.component.aws2-cw.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-cw.unit</code>	The metric unit.		String
<code>camel.component.aws2-cw.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-cw.use-default-credentials-provider</code>	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	Boolean
<code>camel.component.aws2-cw.value</code>	The metric value.		Double

CHAPTER 2. AWS DYNAMODB

Only producer is supported

The AWS2 DynamoDB component supports storing and retrieving data from/to service.

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon DynamoDB. More information is available at [Amazon DynamoDB](#).

2.1. URI FORMAT

```
aws2-ddb://domainName[?options]
```

You can append query options to the URI in the following format, **?options=value&option2=value&...**

2.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

2.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

2.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

2.3. COMPONENT OPTIONS

The AWS DynamoDB component supports 22 options, which are listed below.

Name	Description	Default	Type
amazonDDBClient (producer)	Autowired To use the AmazonDynamoDB as the client.		DynamoDbClient
configuration (producer)	The component configuration.		Ddb2Configuration
consistentRead (producer)	Determines whether or not strong consistency should be enforced when data is read.	false	boolean
enabledInitialDescribeTable (producer)	Set whether the initial Describe table operation in the DDB Endpoint must be done, or not.	true	boolean
keyAttributeName (producer)	Attribute name when creating table.		String
keyAttributeType (producer)	Attribute type when creating table.		String
keyScalarType (producer)	The key scalar type, it can be S (String), N (Number) and B (Bytes).		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>What operation to perform.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● BatchGetItems ● DeleteItem ● DeleteTable ● DescribeTable ● GetItem ● PutItem ● Query ● Scan ● UpdateItem ● UpdateTable 	PutItem	Ddb2Operations
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (producer)	To define a proxy host when instantiating the DDB client.		String
proxyPort (producer)	The region in which DynamoDB client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		Integer
proxyProtocol (producer)	<p>To define a proxy protocol when instantiating the DDB client.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
readCapacity (producer)	The provisioned throughput to reserve for reading resources from your table.		Long
region (producer)	The region in which DDB client needs to work.		String

Name	Description	Default	Type
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
writeCapacity (producer)	The provisioned throughput to reserved for writing resources to your table.		Long
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

2.4. ENDPOINT OPTIONS

The AWS DynamoDB endpoint is configured using URI syntax:

```
aws2-ddb:tableName
```

with the following path and query parameters:

2.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
tableName (producer)	Required The name of the table currently worked with.		String

2.4.2. Query Parameters (20 parameters)

Name	Description	Default	Type
amazonDDBClient (producer)	Autowired To use the AmazonDynamoDB as the client.		DynamoDbClient
consistentRead (producer)	Determines whether or not strong consistency should be enforced when data is read.	false	boolean
enabledInitialDescribeTable (producer)	Set whether the initial Describe table operation in the DDB Endpoint must be done, or not.	true	boolean
keyAttributeName (producer)	Attribute name when creating table.		String
keyAttributeType (producer)	Attribute type when creating table.		String
keyScalarType (producer)	The key scalar type, it can be S (String), N (Number) and B (Bytes).		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>What operation to perform.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● BatchGetItems ● DeleteItem ● DeleteTable ● DescribeTable ● GetItem ● PutItem ● Query ● Scan ● UpdateItem ● UpdateTable 	PutItem	Ddb2Operations
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (producer)	To define a proxy host when instantiating the DDB client.		String
proxyPort (producer)	The region in which DynamoDB client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		Integer
proxyProtocol (producer)	<p>To define a proxy protocol when instantiating the DDB client.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
readCapacity (producer)	The provisioned throughput to reserve for reading resources from your table.		Long
region (producer)	The region in which DDB client needs to work.		String

Name	Description	Default	Type
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
writeCapacity (producer)	The provisioned throughput to reserved for writing resources to your table.		Long
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required DDB component options

You have to provide the `amazonDDBClient` in the Registry or your `accessKey` and `secretKey` to access the [Amazon's DynamoDB](#).

2.5. USAGE

2.5.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`
- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.
- Web Identity Token from AWS STS.
- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

2.5.2. Message headers evaluated by the DDB producer

Header	Type	Description
CamelAwsDdbBatchItems	Map<String, KeysAndAttributes>	A map of the table name and corresponding items to get by primary key.
CamelAwsDdbTableName	String	Table Name for this operation.
CamelAwsDdbKey	Key	The primary key that uniquely identifies each item in a table.
CamelAwsDdbReturnValues	String	Use this parameter if you want to get the attribute name-value pairs before or after they are modified(NONE, ALL_OLD, UPDATED_OLD, ALL_NEW, UPDATED_NEW).
CamelAwsDdbUpdateCondition	Map<String, ExpectedAttributeValue>	Designates an attribute for a conditional modification.
CamelAwsDdbAttributeNames	Collection<String>	If attribute names are not specified then all attributes will be returned.
CamelAwsDdbConsistentRead	Boolean	If set to true, then a consistent read is issued, otherwise eventually consistent is used.
CamelAwsDdbIndexName	String	If set will be used as Secondary Index for Query operation.
CamelAwsDdbItem	Map<String, AttributeValue>	A map of the attributes for the item, and must include the primary key values that define the item.
CamelAwsDdbExactCount	Boolean	If set to true, Amazon DynamoDB returns a total number of items that match the query parameters, instead of a list of the matching items and their attributes.
CamelAwsDdbKeyConditions	Map<String, Condition>	This header specify the selection criteria for the query, and merge together the two old headers CamelAwsDdbHashKeyValue and CamelAwsDdbScanRangeKeyCondition
CamelAwsDdbStartKey	Key	Primary key of the item from which to continue an earlier query.
CamelAwsDdbHashKeyValue	AttributeValue	Value of the hash component of the composite primary key.
CamelAwsDdbLimit	Integer	The maximum number of items to return.

Header	Type	Description
CamelAwsDdbScanRangeKeyCondition	Condition	A container for the attribute values and comparison operators to use for the query.
CamelAwsDdbScanIndexForward	Boolean	Specifies forward or backward traversal of the index.
CamelAwsDdbScanFilter	Map<String, Condition>	Evaluates the scan results and returns only the desired values.
CamelAwsDdbUpdateValues	Map<String, AttributeValueUpdate>	Map of attribute name to the new value and action for the update.

2.5.3. Message headers set during BatchGetItems operation

Header	Type	Description
CamelAwsDdbBatchResponse	Map<String, BatchResponse>	Table names and the respective item attributes from the tables.
CamelAwsDdbUnprocessedKeys	Map<String, KeysAndAttributes>	Contains a map of tables and their respective keys that were not processed with the current response.

2.5.4. Message headers set during DeleteItem operation

Header	Type	Description
CamelAwsDdbAttributes	Map<String, AttributeValue>	The list of attributes returned by the operation.

2.5.5. Message headers set during DeleteTable operation

Header	Type	Description
CamelAwsDdbProvisionedThroughput		
ProvisionedThroughputDescription		The value of the ProvisionedThroughput property for this table
CamelAwsDdbCreationDate	Date	Creation DateTime of this table.

Header	Type	Description
CamelAwsDdbTableItemC ount	Long	Item count for this table.
CamelAwsDdbKeySchem a	KeySchema	The KeySchema that identifies the primary key for this table. From Camel 2.16.0 the type of this header is List<KeySchemaElement> and not KeySchema
CamelAwsDdbTableName	String	The table name.
CamelAwsDdbTableSize	Long	The table size in bytes.
CamelAwsDdbTableStatu s	String	The status of the table: CREATING, UPDATING, DELETING, ACTIVE

2.5.6. Message headers set during DescribeTable operation

Header	Type	Description
CamelAwsDdbProvisionedThroug hput	<code>\ {{ProvisionedThroug hputDescription }}</code>	The value of the ProvisionedThroughput property for this table
CamelAwsDdbCreationDate	Date	Creation DateTime of this table.
CamelAwsDdbTableItemC ount	Long	Item count for this table.
CamelAwsDdbKeySchema	<code>\{{KeySchema}}</code>	The KeySchema that identifies the primary key for this table.
CamelAwsDdbTableName	String	The table name.
CamelAwsDdbTableSize	Long	The table size in bytes.
CamelAwsDdbTableStatus	String	The status of the table: CREATING, UPDATING, DELETING, ACTIVE
CamelAwsDdbReadCapacity	Long	ReadCapacityUnits property of this table.
CamelAwsDdbWriteCapacity	Long	WriteCapacityUnits property of this table.

2.5.7. Message headers set during GetItem operation

Header	Type	Description
CamelAwsDdbAttributes	Map<String, AttributeValue>	The list of attributes returned by the operation.

2.5.8. Message headers set during PutItem operation

Header	Type	Description
CamelAwsDdbAttributes	Map<String, AttributeValue>	The list of attributes returned by the operation.

2.5.9. Message headers set during Query operation

Header	Type	Description
CamelAwsDdbItems	List<java.util.Map<String, AttributeValue>>	The list of attributes returned by the operation.
CamelAwsDdbLastEvaluatedKey	Key	Primary key of the item where the query operation stopped, inclusive of the previous result set.
CamelAwsDdbConsumedCapacity	Double	The number of Capacity Units of the provisioned throughput of the table consumed during the operation.
CamelAwsDdbCount	Integer	Number of items in the response.

2.5.10. Message headers set during Scan operation

Header	Type	Description
CamelAwsDdbItems	List<java.util.Map<String, AttributeValue>>	The list of attributes returned by the operation.
CamelAwsDdbLastEvaluatedKey	Key	Primary key of the item where the query operation stopped, inclusive of the previous result set.
CamelAwsDdbConsumedCapacity	Double	The number of Capacity Units of the provisioned throughput of the table consumed during the operation.
CamelAwsDdbCount	Integer	Number of items in the response.

Header	Type	Description
CamelAwsDdbScannedCount	Integer	Number of items in the complete scan before any filters are applied.

2.5.11. Message headers set during UpdateItem operation

Header	Type	Description
CamelAwsDdbAttributes	Map<String, AttributeValue>	The list of attributes returned by the operation.

2.5.12. Advanced AmazonDynamoDB configuration

If you need more control over the **AmazonDynamoDB** instance configuration you can create your own instance and refer to it from the URI:

```
from("direct:start")
.to("aws2-ddb://domainName?amazonDDBClient=#client");
```

The **#client** refers to a **DynamoDbClient** in the Registry.

2.6. SUPPORTED PRODUCER OPERATIONS

- BatchGetItems
- DeleteItem
- DeleteTable
- DescribeTable
- GetItem
- PutItem
- Query
- Scan
- UpdateItem
- UpdateTable

2.7. EXAMPLES

2.7.1. Producer Examples

- PutItem: this operation will create an entry into DynamoDB

-

```

from("direct:start")
  .setHeader(Ddb2Constants.OPERATION, Ddb2Operations.PutItem)
  .setHeader(Ddb2Constants.CONSISTENT_READ, "true")
  .setHeader(Ddb2Constants.RETURN_VALUES, "ALL_OLD")
  .setHeader(Ddb2Constants.ITEM, attributeMap)
  .setHeader(Ddb2Constants.ATTRIBUTE_NAMES, attributeMap.keySet());
.to("aws2-ddb://" + tableName + "?keyAttributeName=" + attributeName + "&keyAttributeType=" +
KeyType.HASH
  + "&keyScalarType=" + ScalarAttributeType.S
  + "&readCapacity=1&writeCapacity=1");

```

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-ddb</artifactId>
  <version>${camel-version}</version>
</dependency>

```

where **3.14.2** must be replaced by the actual version of Camel.

2.8. SPRING BOOT AUTO-CONFIGURATION

When using aws2-ddb with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-ddb-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 40 options, which are listed below.

Name	Description	Default	Type
camel.component.aws2-ddb.access-key	Amazon AWS Access Key.		String
camel.component.aws2-ddb.amazon-d-d-b-client	To use the AmazonDynamoDB as the client. The option is a software.amazon.awssdk.services.dynamodb.DynamoDbClient type.		DynamoDbClient

Name	Description	Default	Type
<code>camel.component.aws2-ddb.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-ddb.configuration</code>	The component configuration. The option is a <code>org.apache.camel.component.aws2.ddb.Ddb2Configuration</code> type.		Ddb2Configuration
<code>camel.component.aws2-ddb.consistent-read</code>	Determines whether or not strong consistency should be enforced when data is read.	false	Boolean
<code>camel.component.aws2-ddb.enabled</code>	Whether to enable auto configuration of the <code>aws2-ddb</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-ddb.enabled-initial-describe-table</code>	Set whether the initial Describe table operation in the DDB Endpoint must be done, or not.	true	Boolean
<code>camel.component.aws2-ddb.key-attribute-name</code>	Attribute name when creating table.		String
<code>camel.component.aws2-ddb.key-attribute-type</code>	Attribute type when creating table.		String
<code>camel.component.aws2-ddb.key-scalar-type</code>	The key scalar type, it can be S (String), N (Number) and B (Bytes).		String

Name	Description	Default	Type
<code>camel.component.aws2-ddb.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-ddb.operation</code>	What operation to perform.		Ddb2Operations
<code>camel.component.aws2-ddb.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-ddb.proxy-host</code>	To define a proxy host when instantiating the DDB client.		String
<code>camel.component.aws2-ddb.proxy-port</code>	The region in which DynamoDB client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code>) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		Integer
<code>camel.component.aws2-ddb.proxy-protocol</code>	To define a proxy protocol when instantiating the DDB client.		Protocol
<code>camel.component.aws2-ddb.read-capacity</code>	The provisioned throughput to reserve for reading resources from your table.		Long
<code>camel.component.aws2-ddb.region</code>	The region in which DDB client needs to work.		String
<code>camel.component.aws2-ddb.secret-key</code>	Amazon AWS Secret Key.		String

Name	Description	Default	Type
<code>camel.component.aws2-ddb.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-ddb.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-ddb.use-default-credentials-provider</code>	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	Boolean
<code>camel.component.aws2-ddb.write-capacity</code>	The provisioned throughput to reserved for writing resources to your table.		Long
<code>camel.component.aws2-ddbstream.access-key</code>	Amazon AWS Access Key.		String
<code>camel.component.aws2-ddbstream.amazon-dynamo-db-streams-client</code>	Amazon DynamoDB client to use for all requests for this endpoint. The option is a <code>software.amazon.awssdk.services.dynamodb.streams.DynamoDbStreamsClient</code> type.		<code>DynamoDbStreamsClient</code>
<code>camel.component.aws2-ddbstream.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as <code>autowired</code>) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-ddbstream.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean

Name	Description	Default	Type
<code>camel.component.aws2-ddbstream.configuration</code>	The component configuration. The option is a <code>org.apache.camel.component.aws2.ddbstream.Ddb2StreamConfiguration</code> type.		Ddb2StreamConfiguration
<code>camel.component.aws2-ddbstream.enabled</code>	Whether to enable auto configuration of the <code>aws2-ddbstream</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-ddbstream.max-results-per-request</code>	Maximum number of records that will be fetched in each poll.		Integer
<code>camel.component.aws2-ddbstream.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-ddbstream.proxy-host</code>	To define a proxy host when instantiating the DDBStreams client.		String
<code>camel.component.aws2-ddbstream.proxy-port</code>	To define a proxy port when instantiating the DDBStreams client.		Integer
<code>camel.component.aws2-ddbstream.proxy-protocol</code>	To define a proxy protocol when instantiating the DDBStreams client.		Protocol
<code>camel.component.aws2-ddbstream.region</code>	The region in which DDBStreams client needs to work.		String
<code>camel.component.aws2-ddbstream.secret-key</code>	Amazon AWS Secret Key.		String

Name	Description	Default	Type
<code>camel.component.aws2-ddbstream.stream-iterator-type</code>	Defines where in the DynamoDB stream to start getting records. Note that using <code>FROM_START</code> can cause a significant delay before the stream has caught up to real-time.		<code>Ddb2StreamConfiguration\$StreamIteratorType</code>
<code>camel.component.aws2-ddbstream.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	<code>false</code>	<code>Boolean</code>
<code>camel.component.aws2-ddbstream.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		<code>String</code>
<code>camel.component.aws2-ddbstream.use-default-credentials-provider</code>	Set whether the DynamoDB Streams client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	<code>false</code>	<code>Boolean</code>

CHAPTER 3. AWS KINESIS

Both producer and consumer are supported

The AWS2 Kinesis component supports receiving messages from and sending messages to Amazon Kinesis (no Batch supported) service.

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon Kinesis. More information are available at [AWS Kinesis](#).

3.1. URI FORMAT

```
aws2-kinesis://stream-name[?options]
```

The stream needs to be created prior to it being used. You can append query options to the URI in the following format, **?options=value&option2=value&...**

3.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

3.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

3.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

3.3. COMPONENT OPTIONS

The AWS Kinesis component supports 22 options, which are listed below.

Name	Description	Default	Type
amazonKinesisClient (common)	Autowired Amazon Kinesis client to use for all requests for this endpoint.		KinesisClient
cborEnabled (common)	This option will set the CBOR_ENABLED property during the execution.	true	boolean
configuration (common)	Component configuration.		Kinesis2Configuration
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (common)	To define a proxy host when instantiating the Kinesis client.		String
proxyPort (common)	To define a proxy port when instantiating the Kinesis client.		Integer
proxyProtocol (common)	To define a proxy protocol when instantiating the Kinesis client. Enum values: <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
region (common)	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		String
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with overrideEndpoint option.		String
useDefaultCredentialsProvider (common)	Set whether the Kinesis client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
iteratorType (consumer)	Defines where in the Kinesis stream to start getting records. Enum values: <ul style="list-style-type: none"> • AT_SEQUENCE_NUMBER • AFTER_SEQUENCE_NUMBER • TRIM_HORIZON • LATEST • AT_TIMESTAMP • null 	TRIM_HORIZON	ShardIteratorType
maxResultsPerRequest (consumer)	Maximum number of records that will be fetched in each poll.	1	int
resumeStrategy (consumer)	Defines a resume strategy for AWS Kinesis. The default strategy reads the <code>sequenceNumber</code> if provided.	KinesisUserConfigurationResumeStrategy	KinesisResumeStrategy
sequenceNumber (consumer)	The sequence number to start polling from. Required if <code>iteratorType</code> is set to <code>AFTER_SEQUENCE_NUMBER</code> or <code>AT_SEQUENCE_NUMBER</code> .		String

Name	Description	Default	Type
shardClosed (consumer)	<p>Define what will be the behavior in case of shard closed. Possible value are ignore, silent and fail. In case of ignore a message will be logged and the consumer will restart from the beginning, in case of silent there will be no logging and the consumer will start from the beginning, in case of fail a <code>ReachedClosedStateException</code> will be raised.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● ignore ● fail ● silent 	ignore	Kinesis2ShardClosedStrategyEnum
shardId (consumer)	Defines which shardId in the Kinesis stream to get records from.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

3.4. ENDPOINT OPTIONS

The AWS Kinesis endpoint is configured using URI syntax:

aws2-kinesis:streamName

with the following path and query parameters:

3.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
streamName (common)	Required Name of the stream.		String

3.4.2. Query Parameters (38 parameters)

Name	Description	Default	Type
amazonKinesisClient (common)	Autowired Amazon Kinesis client to use for all requests for this endpoint.		KinesisClient
cborEnabled (common)	This option will set the CBOR_ENABLED property during the execution.	true	boolean
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
proxyHost (common)	To define a proxy host when instantiating the Kinesis client.		String
proxyPort (common)	To define a proxy port when instantiating the Kinesis client.		Integer
proxyProtocol (common)	To define a proxy protocol when instantiating the Kinesis client. Enum values: <ul style="list-style-type: none"> • HTTP • HTTPS 	HTTPS	Protocol
region (common)	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		String

Name	Description	Default	Type
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (common)	Set whether the Kinesis client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
iteratorType (consumer)	Defines where in the Kinesis stream to start getting records. Enum values: <ul style="list-style-type: none"> ● AT_SEQUENCE_NUMBER ● AFTER_SEQUENCE_NUMBER ● TRIM_HORIZON ● LATEST ● AT_TIMESTAMP ● null 	TRIM_HORIZON	ShardIteratorType
maxResultsPerRequest (consumer)	Maximum number of records that will be fetched in each poll.	1	int
resumeStrategy (consumer)	Defines a resume strategy for AWS Kinesis. The default strategy reads the <code>sequenceNumber</code> if provided.	KinesisUserConfigurationResumeStrategy	KinesisResumeStrategy

Name	Description	Default	Type
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
sequenceNumber (consumer)	The sequence number to start polling from. Required if iteratorType is set to AFTER_SEQUENCE_NUMBER or AT_SEQUENCE_NUMBER.		String
shardClosed (consumer)	Define what will be the behavior in case of shard closed. Possible value are ignore, silent and fail. In case of ignore a message will be logged and the consumer will restart from the beginning, in case of silent there will be no logging and the consumer will start from the beginning, in case of fail a ReachedClosedStateException will be raised. Enum values: <ul style="list-style-type: none"> ● ignore ● fail ● silent 	ignore	Kinesis2ShardClosedStrategyEnum
shardId (consumer)	Defines which shardId in the Kinesis stream to get records from.		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long

Name	Description	Default	Type
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	<p>Time unit for initialDelay and delay options.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit

Name	Description	Default	Type
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required Kinesis component options

You have to provide the KinesisClient in the Registry with proxies and relevant credentials configured.

3.5. BATCH CONSUMER

This component implements the Batch Consumer.

This allows you for instance to know how many messages exists in this batch and for instance let the Aggregator aggregate this number of messages.

3.6. USAGE

3.6.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`
- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.
- Web Identity Token from AWS STS.
- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

3.6.2. Message headers set by the Kinesis consumer

Header	Type	Description
CamelAwsKinesisSequenceNumber	String	The sequence number of the record. This is represented as a String as its size is not defined by the API. If it is to be used as a numerical type then use
CamelAwsKinesisApproximateArrivalTimestamp	String	The time AWS assigned as the arrival time of the record.
CamelAwsKinesisPartitionKey	String	Identifies which shard in the stream the data record is assigned to.

3.6.3. AmazonKinesis configuration

You then have to reference the KinesisClient in the **amazonKinesisClient** URI option.

```
from("aws2-kinesis://mykinesisstream?amazonKinesisClient=#kinesisClient")
    .to("log:out?showAll=true");
```

3.6.4. Providing AWS Credentials

It is recommended that the credentials are obtained by using the [DefaultAWSCredentialsProviderChain](#) that is the default when creating a new ClientConfiguration instance, however, a different [AWSCredentialsProvider](#) can be specified when calling `createClient(...)`.

3.6.5. Message headers used by the Kinesis producer to write to Kinesis. The producer expects that the message body is a `byte[]`.

Header	Type	Description
CamelAwsKinesisPartitionKey	String	The PartitionKey to pass to Kinesis to store this record.
CamelAwsKinesisSequenceNumber	String	Optional parameter to indicate the sequence number of this record.

3.6.6. Message headers set by the Kinesis producer on successful storage of a Record

Header	Type	Description
CamelAwsKinesisSequenceNumber	String	The sequence number of the record, as defined in Response Syntax
CamelAwsKinesisShardId	String	The shard ID of where the Record was stored

3.7. DEPENDENCIES

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-kinesis</artifactId>
  <version>${camel-version}</version>
</dependency>
```

where **3.14.2** must be replaced by the actual version of Camel.

3.8. SPRING BOOT AUTO-CONFIGURATION

When using aws2-kinesis with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-kinesis-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use your Camel Spring Boot version -->
</dependency>
```

The component supports 40 options, which are listed below.

Name	Description	Default	Type
camel.component.aws2-kinesis-firehose.access-key	Amazon AWS Access Key.		String
camel.component.aws2-kinesis-firehose.amazon-kinesis-firehose-client	Amazon Kinesis Firehose client to use for all requests for this endpoint. The option is a software.amazon.awssdk.services.firehose.FirehoseClient type.		FirehoseClient
camel.component.aws2-kinesis-firehose.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.aws2-kinesis-firehose.cbor-enabled</code>	This option will set the <code>CBOR_ENABLED</code> property during the execution.	true	Boolean
<code>camel.component.aws2-kinesis-firehose.configuration</code>	Component configuration. The option is a <code>org.apache.camel.component.aws2.firehose.KinesisFirehose2Configuration</code> type.		KinesisFirehose2Configuration
<code>camel.component.aws2-kinesis-firehose.enabled</code>	Whether to enable auto configuration of the <code>aws2-kinesis-firehose</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-kinesis-firehose.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-kinesis-firehose.operation</code>	The operation to do in case the user don't want to send only a record.		KinesisFirehose2Operations
<code>camel.component.aws2-kinesis-firehose.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-kinesis-firehose.proxy-host</code>	To define a proxy host when instantiating the Kinesis Firehose client.		String
<code>camel.component.aws2-kinesis-firehose.proxy-port</code>	To define a proxy port when instantiating the Kinesis Firehose client.		Integer

Name	Description	Default	Type
<code>camel.component.aws2-kinesis-firehose.proxy-protocol</code>	To define a proxy protocol when instantiating the Kinesis Firehose client.		Protocol
<code>camel.component.aws2-kinesis-firehose.region</code>	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-kinesis-firehose.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-kinesis-firehose.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-kinesis-firehose.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-kinesis-firehose.use-default-credentials-provider</code>	Set whether the Kinesis Firehose client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	Boolean
<code>camel.component.aws2-kinesis.access-key</code>	Amazon AWS Access Key.		String
<code>camel.component.aws2-kinesis.amazon-kinesis-client</code>	Amazon Kinesis client to use for all requests for this endpoint. The option is a <code>software.amazon.awssdk.services.kinesis.KinesisClient</code> type.		KinesisClient

Name	Description	Default	Type
<code>camel.component.aws2-kinesis.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-kinesis.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.aws2-kinesis.cbor-enabled</code>	This option will set the <code>CBOR_ENABLED</code> property during the execution.	true	Boolean
<code>camel.component.aws2-kinesis.configuration</code>	Component configuration. The option is a <code>org.apache.camel.component.aws2.kinesis.Kinesis2C</code> onfiguration type.		Kinesis2Configurat ion
<code>camel.component.aws2-kinesis.enabled</code>	Whether to enable auto configuration of the <code>aws2-kinesis</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-kinesis.iterator-type</code>	Defines where in the Kinesis stream to start getting records.		ShardIteratorType

Name	Description	Default	Type
<code>camel.component.aws2-kinesis.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-kinesis.max-results-per-request</code>	Maximum number of records that will be fetched in each poll.	1	Integer
<code>camel.component.aws2-kinesis.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-kinesis.proxy-host</code>	To define a proxy host when instantiating the Kinesis client.		String
<code>camel.component.aws2-kinesis.proxy-port</code>	To define a proxy port when instantiating the Kinesis client.		Integer
<code>camel.component.aws2-kinesis.proxy-protocol</code>	To define a proxy protocol when instantiating the Kinesis client.		Protocol
<code>camel.component.aws2-kinesis.region</code>	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code>) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-kinesis.resume-strategy</code>	Defines a resume strategy for AWS Kinesis. The default strategy reads the <code>sequenceNumber</code> if provided. The option is a <code>org.apache.camel.component.aws2.kinesis.consumer.KinesisResumeStrategy</code> type.		<code>KinesisResumeStrategy</code>

Name	Description	Default	Type
<code>camel.component.aws2-kinesis.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-kinesis.sequence-number</code>	The sequence number to start polling from. Required if <code>iteratorType</code> is set to <code>AFTER_SEQUENCE_NUMBER</code> or <code>AT_SEQUENCE_NUMBER</code> .		String
<code>camel.component.aws2-kinesis.shard-closed</code>	Define what will be the behavior in case of shard closed. Possible value are <code>ignore</code> , <code>silent</code> and <code>fail</code> . In case of <code>ignore</code> a message will be logged and the consumer will restart from the beginning, in case of <code>silent</code> there will be no logging and the consumer will start from the beginning, in case of <code>fail</code> a <code>ReachedClosedStateException</code> will be raised.		<code>Kinesis2ShardClosedStrategyEnum</code>
<code>camel.component.aws2-kinesis.shard-id</code>	Defines which <code>shardId</code> in the Kinesis stream to get records from.		String
<code>camel.component.aws2-kinesis.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	<code>false</code>	Boolean
<code>camel.component.aws2-kinesis.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-kinesis.use-default-credentials-provider</code>	Set whether the Kinesis client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	<code>false</code>	Boolean

CHAPTER 4. AWS 2 LAMBDA

Only producer is supported

The AWS2 Lambda component supports create, get, list, delete and invoke [AWS Lambda](#) functions.

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon Lambda. More information is available at [AWS Lambda](#).

When creating a Lambda function, you need to specify a IAM role which has at least the `AWSLambdaBasicExecuteRole` policy attached.

4.1. URI FORMAT

```
aws2-lambda://functionName[?options]
```

You can append query options to the URI in the following format, **options=value&option2=value&...**

4.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

4.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

4.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

4.3. COMPONENT OPTIONS

The AWS Lambda component supports 16 options, which are listed below.

Name	Description	Default	Type
configuration (producer)	Component configuration.		Lambda2Configuration
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>The operation to perform. It can be listFunctions, getFunction, createFunction, deleteFunction or invokeFunction.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● listFunctions ● getFunction ● createAlias ● deleteAlias ● getAlias ● listAliases ● createFunction ● deleteFunction ● invokeFunction ● updateFunction ● createEventSourceMapping ● deleteEventSourceMapping ● listEventSourceMapping ● listTags ● tagResource ● untagResource ● publishVersion ● listVersions 	invokeFunction	Lambda2Operations
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
pojoRequest (producer)	If we want to use a POJO request as body or not.	false	boolean
region (producer)	The region in which Lambda client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		String

Name	Description	Default	Type
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the Lambda client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
awsLambdaClient (advanced)	Autowired To use a existing configured <code>AwsLambdaClient</code> as client.		<code>LambdaClient</code>
proxyHost (proxy)	To define a proxy host when instantiating the Lambda client.		String
proxyPort (proxy)	To define a proxy port when instantiating the Lambda client.		Integer
proxyProtocol (proxy)	To define a proxy protocol when instantiating the Lambda client. Enum values: <ul style="list-style-type: none"> • HTTP • HTTPS 	HTTPS	Protocol
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

4.4. ENDPOINT OPTIONS

The AWS Lambda endpoint is configured using URI syntax:

```
aws2-lambda:function
```

with the following path and query parameters:

4.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
function (producer)	Required Name of the Lambda function.		String

4.4.2. Query Parameters (14 parameters)

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>The operation to perform. It can be listFunctions, getFunction, createFunction, deleteFunction or invokeFunction.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● listFunctions ● getFunction ● createAlias ● deleteAlias ● getAlias ● listAliases ● createFunction ● deleteFunction ● invokeFunction ● updateFunction ● createEventSourceMapping ● deleteEventSourceMapping ● listEventSourceMapping ● listTags ● tagResource ● untagResource ● publishVersion ● listVersions 	invokeFunction	Lambda2Operations
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
pojoRequest (producer)	If we want to use a POJO request as body or not.	false	boolean
region (producer)	The region in which Lambda client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id().		String

Name	Description	Default	Type
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the Lambda client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
awsLambdaClient (advanced)	Autowired To use an existing configured <code>AwsLambdaClient</code> as client.		<code>LambdaClient</code>
proxyHost (proxy)	To define a proxy host when instantiating the Lambda client.		String
proxyPort (proxy)	To define a proxy port when instantiating the Lambda client.		Integer
proxyProtocol (proxy)	To define a proxy protocol when instantiating the Lambda client. Enum values: <ul style="list-style-type: none"> • HTTP • HTTPS 	HTTPS	Protocol
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required Lambda component options

You have to provide the `awsLambdaClient` in the Registry or your `accessKey` and `secretKey` to access the [Amazon Lambda](#) service..

4.5. USAGE

4.5.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`

- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.
- Web Identity Token from AWS STS.
- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

4.5.2. Message headers evaluated by the Lambda producer

Operation	Header	Type	Description	Required
All	CamelAwsLambdaOperation	String	The operation we want to perform. Override operation passed as query parameter	Yes
createFunction	CamelAwsLambdaS3Bucket	String	Amazon S3 bucket name where the .zip file containing your deployment package is stored. This bucket must reside in the same AWS region where you are creating the Lambda function.	No
createFunction	CamelAwsLambdaS3Key	String	The Amazon S3 object (the deployment package) key name you want to upload.	No
createFunction	CamelAwsLambdaS3ObjectVersion	String	The Amazon S3 object (the deployment package) version you want to upload.	No
createFunction	CamelAwsLambdaZipFile	String	The local path of the zip file (the deployment package). Content of zip file can also be put in Message body.	No

Operation	Header	Type	Description	Required
createFunction	CamelAwsLambdaRole	String	The Amazon Resource Name (ARN) of the IAM role that Lambda assumes when it executes your function to access any other Amazon Web Services (AWS) resources.	Yes
createFunction	CamelAwsLambdaRuntime	String	The runtime environment for the Lambda function you are uploading. (nodejs, nodejs4.3, nodejs6.10, java8, python2.7, python3.6, dotnetcore1.0, odejs4.3-edge)	Yes
createFunction	CamelAwsLambdaHandler	String	The function within your code that Lambda calls to begin execution. For Node.js, it is the module-name.export value in your function. For Java, it can be package.class-name::handler or package.class-name.	Yes
createFunction	CamelAwsLambdaDescription	String	The user-provided description.	No
createFunction	CamelAwsLambdaTargetArn	String	The parent object that contains the target ARN (Amazon Resource Name) of an Amazon SQS queue or Amazon SNS topic.	No
createFunction	CamelAwsLambdaMemorySize	Integer	The memory size, in MB, you configured for the function. Must be a multiple of 64 MB.	No

Operation	Header	Type	Description	Required
createFunction	CamelAwsLambdaKMSKeyArn	String	The Amazon Resource Name (ARN) of the KMS key used to encrypt your function's environment variables. If not provided, AWS Lambda will use a default service key.	No
createFunction	CamelAwsLambdaPublish	Boolean	This boolean parameter can be used to request AWS Lambda to create the Lambda function and publish a version as an atomic operation.	No
createFunction	CamelAwsLambdaTimeout	Integer	The function execution time at which Lambda should terminate the function. The default is 3 seconds.	No
createFunction	CamelAwsLambdaTracingConfig	String	Your function's tracing settings (Active or PassThrough).	No
createFunction	CamelAwsLambdaEnvironmentVariables	Map<String, String>	The key-value pairs that represent your environment's configuration settings.	No
createFunction	CamelAwsLambdaEnvironmentTags	Map<String, String>	The list of tags (key-value pairs) assigned to the new function.	No
createFunction	CamelAwsLambdaSecurityGroupIds	List<String>	If your Lambda function accesses resources in a VPC, a list of one or more security groups IDs in your VPC.	No
createFunction	CamelAwsLambdaSubnetIds	List<String>	If your Lambda function accesses resources in a VPC, a list of one or more subnet IDs in your VPC.	No
createAlias	CamelAwsLambdaFunctionVersion	String	The function version to set in the alias	Yes

Operation	Header	Type	Description	Required
createAlias	CamelAwsLambdaAliasFunctionName	String	The function name to set in the alias	Yes
createAlias	CamelAwsLambdaAliasFunctionDescription	String	The function description to set in the alias	No
deleteAlias	CamelAwsLambdaAliasFunctionName	String	The function name of the alias	Yes
getAlias	CamelAwsLambdaAliasFunctionName	String	The function name of the alias	Yes
listAliases	CamelAwsLambdaFunctionVersion	String	The function version to set in the alias	No

4.6. LIST OF AVAILABLE OPERATIONS

- listFunctions
- getFunction
- createFunction
- deleteFunction
- invokeFunction
- updateFunction
- createEventSourceMapping
- deleteEventSourceMapping
- listEventSourceMapping
- listTags
- tagResource
- untagResource
- publishVersion
- listVersions
- createAlias
- deleteAlias
- getAlias

- listAliases

4.7. EXAMPLES

4.7.1. Producer Example

To have a full understanding of how the component works, you may have a look at these [integration tests](#).

4.7.2. Producer Examples

- CreateFunction: this operation will create a function for you in AWS Lambda

```
from("direct:createFunction").to("aws2-lambda://GetHelloWithName?
operation=createFunction").to("mock:result");
```

and by sending

```
template.send("direct:createFunction", ExchangePattern.InOut, new Processor() {
    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(Lambda2Constants.RUNTIME, "nodejs6.10");
        exchange.getIn().setHeader(Lambda2Constants.HANDLER, "GetHelloWithName.handler");
        exchange.getIn().setHeader(Lambda2Constants.DESCRPTION, "Hello with node.js on
Lambda");
        exchange.getIn().setHeader(Lambda2Constants.ROLE,
            "arn:aws:iam::643534317684:role/lambda-execution-role");
        ClassLoader classLoader = getClass().getClassLoader();
        File file = new File(
            classLoader

.getResource("org/apache/camel/component/aws2/lambda/function/node/GetHelloWithName.zip")
                .getFile());
        FileInputStream inputStream = new FileInputStream(file);
        exchange.getIn().setBody(inputStream);
    }
});
```

4.8. USING A POJO AS BODY

Sometimes build an AWS Request can be complex, because of multiple options. We introduce the possibility to use a POJO as body. In AWS Lambda there are multiple operations you can submit, as an example for Get Function request, you can do something like:

```
from("direct:getFunction")
    .setBody(GetFunctionRequest.builder().functionName("test").build())
    .to("aws2-lambda://GetHelloWithName?
awsLambdaClient=#awsLambdaClient&operation=getFunction&pojoRequest=true")
```

In this way you'll pass the request directly without the need of passing headers and options specifically related to this operation.

4.9. DEPENDENCIES

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-lambda</artifactId>
  <version>${camel-version}</version>
</dependency>
```

where **3.14.2** must be replaced by the actual version of Camel.

4.10. SPRING BOOT AUTO-CONFIGURATION

When using aws2-lambda with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-lambda-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use your Camel Spring Boot version -->
</dependency>
```

The component supports 17 options, which are listed below.

Name	Description	Default	Type
camel.component.aws2-lambda.access-key	Amazon AWS Access Key.		String
camel.component.aws2-lambda.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.aws2-lambda.aws-lambda-client	To use a existing configured AwsLambdaClient as client. The option is a software.amazon.awssdk.services.lambda.LambdaClient type.		LambdaClient

Name	Description	Default	Type
<code>camel.component.aws2-lambda.configuration</code>	Component configuration. The option is a <code>org.apache.camel.component.aws2.lambda.Lambda2Configuration</code> type.		<code>Lambda2Configuration</code>
<code>camel.component.aws2-lambda.enabled</code>	Whether to enable auto configuration of the <code>aws2-lambda</code> component. This is enabled by default.		<code>Boolean</code>
<code>camel.component.aws2-lambda.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	<code>false</code>	<code>Boolean</code>
<code>camel.component.aws2-lambda.operation</code>	The operation to perform. It can be <code>listFunctions</code> , <code>getFunction</code> , <code>createFunction</code> , <code>deleteFunction</code> or <code>invokeFunction</code> .		<code>Lambda2Operations</code>
<code>camel.component.aws2-lambda.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	<code>false</code>	<code>Boolean</code>
<code>camel.component.aws2-lambda.pojo-request</code>	If we want to use a POJO request as body or not.	<code>false</code>	<code>Boolean</code>
<code>camel.component.aws2-lambda.proxy-host</code>	To define a proxy host when instantiating the Lambda client.		<code>String</code>
<code>camel.component.aws2-lambda.proxy-port</code>	To define a proxy port when instantiating the Lambda client.		<code>Integer</code>
<code>camel.component.aws2-lambda.proxy-protocol</code>	To define a proxy protocol when instantiating the Lambda client.		<code>Protocol</code>

Name	Description	Default	Type
<code>camel.component.aws2-lambda.region</code>	The region in which Lambda client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-lambda.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-lambda.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-lambda.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-lambda.use-default-credentials-provider</code>	Set whether the Lambda client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	Boolean

CHAPTER 5. AWS S3 STORAGE SERVICE

Both producer and consumer are supported

The AWS2 S3 component supports storing and retrieving objects from/to [Amazon's S3](#) service.

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon S3. More information is available at link:<https://aws.amazon.com/s3> [Amazon S3].

5.1. URI FORMAT

```
aws2-s3://bucketNameOrArn[?options]
```

The bucket will be created if it don't already exists. You can append query options to the URI in the following format,

options=value&option2=value&...

5.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

5.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

5.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

5.3. COMPONENT OPTIONS

The AWS S3 Storage Service component supports 50 options, which are listed below.

Name	Description	Default	Type
amazonS3Client (common)	Autowired Reference to a <code>com.amazonaws.services.s3.AmazonS3</code> in the registry.		S3Client
amazonS3Presigner (common)	Autowired An S3 Presigner for Request, used mainly in <code>createDownloadLink</code> operation.		S3Presigner
autoCreateBucket (common)	Setting the autocreation of the S3 bucket <code>bucketName</code> . This will apply also in case of <code>moveAfterRead</code> option enabled and it will create the <code>destinationBucket</code> if it doesn't exist already.	false	boolean
configuration (common)	The component configuration.		AWS2S3Configuration
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	boolean
pojoRequest (common)	If we want to use a POJO request as body or not.	false	boolean
policy (common)	The policy for this queue to set in the <code>com.amazonaws.services.s3.AmazonS3#setBucketPolicy()</code> method.		String
proxyHost (common)	To define a proxy host when instantiating the SQS client.		String
proxyPort (common)	Specify a proxy port to be used inside the client definition.		Integer
proxyProtocol (common)	To define a proxy protocol when instantiating the S3 client. Enum values: <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol

Name	Description	Default	Type
region (common)	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (common)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
customerAlgorithm (common (advanced))	Define the customer algorithm to use in case <code>CustomerKey</code> is enabled.		String
customerKeyId (common (advanced))	Define the id of Customer key to use in case <code>CustomerKey</code> is enabled.		String
customerKeyMD5 (common (advanced))	Define the MD5 of Customer key to use in case <code>CustomerKey</code> is enabled.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
deleteAfterRead (consumer)	Delete objects from S3 after they have been retrieved. The delete is only performed if the Exchange is committed. If a rollback occurs, the object is not deleted. If this option is false, then the same objects will be retrieve over and over again on the polls. Therefore you need to use the Idempotent Consumer EIP in the route to filter out duplicates. You can filter using the <code>AWS2S3Constants#BUCKET_NAME</code> and <code>AWS2S3Constants#KEY</code> headers, or only the <code>AWS2S3Constants#KEY</code> header.	true	boolean
delimiter (consumer)	The delimiter which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.		String
destinationBucket (consumer)	Define the destination bucket where an object must be moved when <code>moveAfterRead</code> is set to true.		String
destinationBucketPrefix (consumer)	Define the destination bucket prefix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
destinationBucketSuffix (consumer)	Define the destination bucket suffix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
doneFileName (consumer)	If provided, Camel will only consume files if a done file exists.		String
fileName (consumer)	To get the object from the bucket with the given file name.		String
ignoreBody (consumer)	If it is true, the S3 Object Body will be ignored completely, if it is set to false the S3 Object will be put in the body. Setting this to true, will override any behavior defined by <code>includeBody</code> option.	false	boolean

Name	Description	Default	Type
includeBody (consumer)	If it is true, the S3Object exchange will be consumed and put into the body and closed. If false the S3Object stream will be put raw into the body and the headers will be set with the S3 object metadata. This option is strongly related to autocloseBody option. In case of setting includeBody to true because the S3Object stream will be consumed then it will also be closed, while in case of includeBody false then it will be up to the caller to close the S3Object stream. However setting autocloseBody to true when includeBody is false it will schedule to close the S3Object stream automatically on exchange completion.	true	boolean
includeFolders (consumer)	If it is true, the folders/directories will be consumed. If it is false, they will be ignored, and Exchanges will not be created for those.	true	boolean
moveAfterRead (consumer)	Move objects from S3 bucket to a different bucket after they have been retrieved. To accomplish the operation the destinationBucket option must be set. The copy bucket operation is only performed if the Exchange is committed. If a rollback occurs, the object is not moved.	false	boolean
prefix (consumer)	The prefix which is used in the com.amazonaws.services.s3.model.ListObjectsRequest to only consume objects we are interested in.		String
autocloseBody (consumer (advanced))	If this option is true and includeBody is false, then the S3Object.close() method will be called on exchange completion. This option is strongly related to includeBody option. In case of setting includeBody to false and autocloseBody to false, it will be up to the caller to close the S3Object stream. Setting autocloseBody to true, will close the S3Object stream automatically.	true	boolean
batchMessageNumber (producer)	The number of messages composing a batch in streaming upload mode.	10	int
batchSize (producer)	The batch size (in bytes) in streaming upload mode.	10000 00	int
deleteAfterWrite (producer)	Delete file object after the S3 file has been uploaded.	false	boolean
keyName (producer)	Setting the key name for an element in the bucket through endpoint parameter.		String

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
multiPartUpload (producer)	If it is true, camel will upload the file with multi part format, the part size is decided by the option of partSize.	false	boolean
namingStrategy (producer)	The naming strategy to use in streaming upload mode. Enum values: <ul style="list-style-type: none"> ● progressive ● random 	progressive	AWS3NamingStrategyEnum
operation (producer)	The operation to do in case the user don't want to do only an upload. Enum values: <ul style="list-style-type: none"> ● copyObject ● listObjects ● deleteObject ● deleteBucket ● listBuckets ● getObject ● getObjectRange ● createDownloadLink 		AWS2S3Operations
partSize (producer)	Setup the partSize which is used in multi part upload, the default size is 25M.	26214400	long

Name	Description	Default	Type
restartingPolicy (producer)	The restarting policy to use in streaming upload mode. Enum values: <ul style="list-style-type: none"> ● <code>override</code> ● <code>lastPart</code> 	<code>override</code>	<code>AWSS3RestartingPolicyEnum</code>
storageClass (producer)	The storage class to set in the <code>com.amazonaws.services.s3.model.PutObjectRequest</code> request.		<code>String</code>
streamingUploadMode (producer)	When stream mode is true the upload to bucket will be done in streaming.	<code>false</code>	<code>boolean</code>
streamingUploadTimeout (producer)	While streaming upload mode is true, this option set the timeout to complete upload.		<code>long</code>
awsKMSKeyId (producer (advanced))	Define the id of KMS key to use in case KMS is enabled.		<code>String</code>
useAwsKMS (producer (advanced))	Define if KMS must be used or not.	<code>false</code>	<code>boolean</code>
useCustomerKey (producer (advanced))	Define if Customer Key must be used or not.	<code>false</code>	<code>boolean</code>
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	<code>true</code>	<code>boolean</code>
accessKey (security)	Amazon AWS Access Key.		<code>String</code>
secretKey (security)	Amazon AWS Secret Key.		<code>String</code>

5.4. ENDPOINT OPTIONS

The AWS S3 Storage Service endpoint is configured using URI syntax:

```
aws2-s3://bucketNameOrArn
```

with the following path and query parameters:

5.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
bucketNameOrArn (common)	Required Bucket name or ARN.		String

5.4.2. Query Parameters (68 parameters)

Name	Description	Default	Type
amazonS3Client (common)	Autowired Reference to a <code>com.amazonaws.services.s3.AmazonS3</code> in the registry.		S3Client
amazonS3Presigner (common)	Autowired An S3 Presigner for Request, used mainly in <code>createDownloadLink</code> operation.		S3Presigner
autoCreateBucket (common)	Setting the autocreation of the S3 bucket <code>bucketName</code> . This will apply also in case of <code>moveAfterRead</code> option enabled and it will create the <code>destinationBucket</code> if it doesn't exist already.	false	boolean
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	boolean
pojoRequest (common)	If we want to use a POJO request as body or not.	false	boolean
policy (common)	The policy for this queue to set in the <code>com.amazonaws.services.s3.AmazonS3#setBucketPolicy()</code> method.		String
proxyHost (common)	To define a proxy host when instantiating the SQS client.		String
proxyPort (common)	Specify a proxy port to be used inside the client definition.		Integer

Name	Description	Default	Type
proxyProtocol (common)	To define a proxy protocol when instantiating the S3 client. Enum values: <ul style="list-style-type: none"> • HTTP • HTTPS 	HTTPS	Protocol
region (common)	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (common)	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	boolean
customerAlgorithm (common (advanced))	Define the customer algorithm to use in case <code>CustomerKey</code> is enabled.		String
customerKeyId (common (advanced))	Define the id of Customer key to use in case <code>CustomerKey</code> is enabled.		String
customerKeyMD5 (common (advanced))	Define the MD5 of Customer key to use in case <code>CustomerKey</code> is enabled.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
deleteAfterRead (consumer)	Delete objects from S3 after they have been retrieved. The delete is only performed if the Exchange is committed. If a rollback occurs, the object is not deleted. If this option is false, then the same objects will be retrieve over and over again on the polls. Therefore you need to use the Idempotent Consumer EIP in the route to filter out duplicates. You can filter using the <code>AWS2S3Constants#BUCKET_NAME</code> and <code>AWS2S3Constants#KEY</code> headers, or only the <code>AWS2S3Constants#KEY</code> header.	true	boolean
delimiter (consumer)	The delimiter which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.		String
destinationBucket (consumer)	Define the destination bucket where an object must be moved when <code>moveAfterRead</code> is set to true.		String
destinationBucketPrefix (consumer)	Define the destination bucket prefix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
destinationBucketSuffix (consumer)	Define the destination bucket suffix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
doneFileName (consumer)	If provided, Camel will only consume files if a done file exists.		String
fileName (consumer)	To get the object from the bucket with the given file name.		String
ignoreBody (consumer)	If it is true, the S3 Object Body will be ignored completely, if it is set to false the S3 Object will be put in the body. Setting this to true, will override any behavior defined by <code>includeBody</code> option.	false	boolean

Name	Description	Default	Type
includeBody (consumer)	If it is true, the S3Object exchange will be consumed and put into the body and closed. If false the S3Object stream will be put raw into the body and the headers will be set with the S3 object metadata. This option is strongly related to autocloseBody option. In case of setting includeBody to true because the S3Object stream will be consumed then it will also be closed, while in case of includeBody false then it will be up to the caller to close the S3Object stream. However setting autocloseBody to true when includeBody is false it will schedule to close the S3Object stream automatically on exchange completion.	true	boolean
includeFolders (consumer)	If it is true, the folders/directories will be consumed. If it is false, they will be ignored, and Exchanges will not be created for those.	true	boolean
maxConnections (consumer)	Set the maxConnections parameter in the S3 client configuration.	60	int
maxMessagesPer Poll (consumer)	Gets the maximum number of messages as a limit to poll at each polling. Gets the maximum number of messages as a limit to poll at each polling. The default value is 10. Use 0 or a negative number to set it as unlimited.	10	int
moveAfterRead (consumer)	Move objects from S3 bucket to a different bucket after they have been retrieved. To accomplish the operation the destinationBucket option must be set. The copy bucket operation is only performed if the Exchange is committed. If a rollback occurs, the object is not moved.	false	boolean
prefix (consumer)	The prefix which is used in the com.amazonaws.services.s3.model.ListObjectsRequest to only consume objects we are interested in.		String
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean

Name	Description	Default	Type
autocloseBody (consumer (advanced))	If this option is true and includeBody is false, then the <code>S3Object.close()</code> method will be called on exchange completion. This option is strongly related to includeBody option. In case of setting includeBody to false and autocloseBody to false, it will be up to the caller to close the <code>S3Object</code> stream. Setting autocloseBody to true, will close the <code>S3Object</code> stream automatically.	true	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at <code>WARN</code> or <code>ERROR</code> level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • <code>InOnly</code> • <code>InOut</code> • <code>InOptionalOut</code> 		<code>ExchangePattern</code>
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>
batchMessageNumber (producer)	The number of messages composing a batch in streaming upload mode.	10	int
batchSize (producer)	The batch size (in bytes) in streaming upload mode.	1000000	int
deleteAfterWrite (producer)	Delete file object after the S3 file has been uploaded.	false	boolean
keyName (producer)	Setting the key name for an element in the bucket through endpoint parameter.		String

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
multiPartUpload (producer)	If it is true, camel will upload the file with multi part format, the part size is decided by the option of partSize.	false	boolean
namingStrategy (producer)	The naming strategy to use in streaming upload mode. Enum values: <ul style="list-style-type: none"> ● progressive ● random 	progressive	AWSS3NamingStrategyEnum
operation (producer)	The operation to do in case the user don't want to do only an upload. Enum values: <ul style="list-style-type: none"> ● copyObject ● listObjects ● deleteObject ● deleteBucket ● listBuckets ● getObject ● getObjectRange ● createDownloadLink 		AWS2S3Operations
partSize (producer)	Setup the partSize which is used in multi part upload, the default size is 25M.	26214400	long

Name	Description	Default	Type
restartingPolicy (producer)	The restarting policy to use in streaming upload mode. Enum values: <ul style="list-style-type: none"> ● <code>override</code> ● <code>lastPart</code> 	<code>override</code>	<code>AWSS3RestartingPolicyEnum</code>
storageClass (producer)	The storage class to set in the <code>com.amazonaws.services.s3.model.PutObjectRequest</code> request.		<code>String</code>
streamingUploadMode (producer)	When stream mode is true the upload to bucket will be done in streaming.	<code>false</code>	<code>boolean</code>
streamingUploadTimeout (producer)	While streaming upload mode is true, this option set the timeout to complete upload.		<code>long</code>
awsKMSKeyId (producer (advanced))	Define the id of KMS key to use in case KMS is enabled.		<code>String</code>
useAwsKMS (producer (advanced))	Define if KMS must be used or not.	<code>false</code>	<code>boolean</code>
useCustomerKey (producer (advanced))	Define if Customer Key must be used or not.	<code>false</code>	<code>boolean</code>
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		<code>int</code>
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.		<code>int</code>
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		<code>int</code>
delay (scheduler)	Milliseconds before the next poll.	<code>500</code>	<code>long</code>

Name	Description	Default	Type
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean

Name	Description	Default	Type
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> • NANOSECONDS • MICROSECONDS • MILLISECONDS • SECONDS • MINUTES • HOURS • DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required S3 component options

You have to provide the `amazonS3Client` in the Registry or your `accessKey` and `secretKey` to access the [Amazon's S3](#).

5.5. BATCH CONSUMER

This component implements the Batch Consumer.

This allows you for instance to know how many messages exists in this batch and for instance let the Aggregator aggregate this number of messages.

5.6. USAGE

For example in order to read file **hello.txt** from bucket **helloBucket**, use the following snippet:

```
from("aws2-s3://helloBucket?
accessKey=yourAccessKey&secretKey=yourSecretKey&prefix=hello.txt")
.to("file:/var/downloaded");
```

5.6.1. Message headers evaluated by the S3 producer

Header	Type	Description
CamelAwsS3BucketName	String	The bucket Name which this object will be stored or which will be used for the current operation
CamelAwsS3BucketDestinationName	String	The bucket Destination Name which will be used for the current operation
CamelAwsS3ContentLength	Long	The content length of this object.
CamelAwsS3ContentType	String	The content type of this object.
CamelAwsS3ContentControl	String	The content control of this object.
CamelAwsS3ContentDisposition	String	The content disposition of this object.
CamelAwsS3ContentEncoding	String	The content encoding of this object.
CamelAwsS3ContentMD5	String	The md5 checksum of this object.
CamelAwsS3DestinationKey	String	The Destination key which will be used for the current operation
CamelAwsS3Key	String	The key under which this object will be stored or which will be used for the current operation
CamelAwsS3LastModified	java.util.Date	The last modified timestamp of this object.
CamelAwsS3Operation	String	The operation to perform. Permitted values are copyObject, deleteObject, listBuckets, deleteBucket, listObjects
CamelAwsS3StorageClass	String	The storage class of this object.
CamelAwsS3CannedAcl	String	The canned acl that will be applied to the object. see software.amazon.awssdk.services.s3.model.ObjectCannedACL for allowed values.
CamelAwsS3Acl	software.amazon.awssdk.services.s3.model.BucketCannedACL	A well constructed Amazon S3 Access Control List object. see software.amazon.awssdk.services.s3.model.BucketCannedACL for more details

Header	Type	Description
CamelAwsS3ServerSideEncryption	String	Sets the server-side encryption algorithm when encrypting the object using AWS-managed keys. For example use AES256.
CamelAwsS3VersionId	String	The version Id of the object to be stored or returned from the current operation
CamelAwsS3Metadata	Map<String, String>	A map of metadata to be stored with the object in S3. More details about metadata .

5.6.2. Message headers set by the S3 producer

Header	Type	Description
CamelAwsS3ETag	String	The ETag value for the newly uploaded object.
CamelAwsS3VersionId	String	The optional version ID of the newly uploaded object.

5.6.3. Message headers set by the S3 consumer

Header	Type	Description
CamelAwsS3Key	String	The key under which this object is stored.
CamelAwsS3BucketName	String	The name of the bucket in which this object is contained.
CamelAwsS3ETag	String	The hex encoded 128-bit MD5 digest of the associated object according to RFC 1864. This data is used as an integrity check to verify that the data received by the caller is the same data that was sent by Amazon S3.
CamelAwsS3LastModified	Date	The value of the Last-Modified header, indicating the date and time at which Amazon S3 last recorded a modification to the associated object.
CamelAwsS3VersionId	String	The version ID of the associated Amazon S3 object if available. Version IDs are only assigned to objects when an object is uploaded to an Amazon S3 bucket that has object versioning enabled.

Header	Type	Description
CamelAwsS3ContentType	String	The Content-Type HTTP header, which indicates the type of content stored in the associated object. The value of this header is a standard MIME type.
CamelAwsS3ContentMD5	String	The base64 encoded 128-bit MD5 digest of the associated object (content - not including headers) according to RFC 1864. This data is used as a message integrity check to verify that the data received by Amazon S3 is the same data that the caller sent.
CamelAwsS3ContentLength	Long	The Content-Length HTTP header indicating the size of the associated object in bytes.
CamelAwsS3ContentEncoding	String	The optional Content-Encoding HTTP header specifying what content encodings have been applied to the object and what decoding mechanisms must be applied in order to obtain the media-type referenced by the Content-Type field.
CamelAwsS3ContentDisposition	String	The optional Content-Disposition HTTP header, which specifies presentational information such as the recommended filename for the object to be saved as.
CamelAwsS3ContentControl	String	The optional Cache-Control HTTP header which allows the user to specify caching behavior along the HTTP request/reply chain.
CamelAwsS3ServerSideEncryption	String	The server-side encryption algorithm when encrypting the object using AWS-managed keys.
CamelAwsS3Metadata	Map<String, String>	A map of metadata stored with the object in S3. More details about metadata .

5.6.4. S3 Producer operations

Camel-AWS2-S3 component provides the following operation on the producer side:

- copyObject
- deleteObject
- listBuckets
- deleteBucket

- listObjects
- getObject (this will return an S3Object instance)
- getObjectRange (this will return an S3Object instance)
- createDownloadLink

If you don't specify an operation explicitly the producer will do: - a single file upload - a multipart upload if multiPartUpload option is enabled.

5.6.5. Advanced AmazonS3 configuration

If your Camel Application is running behind a firewall or if you need to have more control over the **S3Client** instance configuration, you can create your own instance and refer to it in your Camel aws2-s3 component configuration:

```
from("aws2-s3://MyBucket?amazonS3Client=#client&delay=5000&maxMessagesPerPoll=5")
.to("mock:result");
```

5.6.6. Use KMS with the S3 component

To use AWS KMS to encrypt/decrypt data by using AWS infrastructure you can use the options introduced in 2.21.x like in the following example

```
from("file:tmp/test?fileName=test.txt")
.setHeader(S3Constants.KEY, constant("testFile"))
.to("aws2-s3://mybucket?amazonS3Client=#client&useAwsKMS=true&awsKMSKeyId=3f0637ad-296a-3dfe-a796-e60654fb128c");
```

In this way you'll ask to S3, to use the KMS key 3f0637ad-296a-3dfe-a796-e60654fb128c, to encrypt the file test.txt. When you'll ask to download this file, the decryption will be done directly before the download.

5.6.7. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the useDefaultCredentialsProvider option and set it to true.

- Java system properties - aws.accessKeyId and aws.secretKey
- Environment variables - AWS_ACCESS_KEY_ID and AWS_SECRET_ACCESS_KEY.
- Web Identity Token from AWS STS.
- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable AWS_CONTAINER_CREDENTIALS_RELATIVE_URI is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

5.6.8. S3 Producer Operation examples

- Single Upload: This operation will upload a file to S3 based on the body content

```
from("direct:start").process(new Processor() {

    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.KEY, "camel.txt");
        exchange.getIn().setBody("Camel rocks!");
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client")
.to("mock:result");
```

This operation will upload the file camel.txt with the content "Camel rocks!" in the mycamelbucket bucket

- Multipart Upload: This operation will perform a multipart upload of a file to S3 based on the body content

```
from("direct:start").process(new Processor() {

    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(AWS2S3Constants.KEY, "empty.txt");
        exchange.getIn().setBody(new File("src/empty.txt"));
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&multiPartUpload=true&autoCreateBucket=true&partSize=1048576")
.to("mock:result");
```

This operation will perform a multipart upload of the file empty.txt with based on the content the file src/empty.txt in the mycamelbucket bucket

- CopyObject: this operation copy an object from one bucket to a different one

```
from("direct:start").process(new Processor() {

    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.BUCKET_DESTINATION_NAME,
"camelDestinationBucket");
        exchange.getIn().setHeader(S3Constants.KEY, "camelKey");
        exchange.getIn().setHeader(S3Constants.DESTINATION_KEY, "camelDestinationKey");
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=copyObject")
.to("mock:result");
```

This operation will copy the object with the name expressed in the header camelDestinationKey to the camelDestinationBucket bucket, from the bucket mycamelbucket.

- DeleteObject: this operation deletes an object from a bucket

```
from("direct:start").process(new Processor() {
    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.KEY, "camelKey");
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=deleteObject")
.to("mock:result");
```

This operation will delete the object camelKey from the bucket mycamelbucket.

- ListBuckets: this operation list the buckets for this account in this region

```
from("direct:start")
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=listBuckets")
.to("mock:result");
```

This operation will list the buckets for this account

- DeleteBucket: this operation delete the bucket specified as URI parameter or header

```
from("direct:start")
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=deleteBucket")
.to("mock:result");
```

This operation will delete the bucket mycamelbucket

- ListObjects: this operation list object in a specific bucket

```
from("direct:start")
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=listObjects")
.to("mock:result");
```

This operation will list the objects in the mycamelbucket bucket

- GetObject: this operation get a single object in a specific bucket

```
from("direct:start").process(new Processor() {
    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.KEY, "camelKey");
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=getObject")
.to("mock:result");
```

This operation will return an S3Object instance related to the camelKey object in mycamelbucket bucket.

- GetObjectRange: this operation get a single object range in a specific bucket

```

from("direct:start").process(new Processor() {

    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.KEY, "camelKey");
        exchange.getIn().setHeader(S3Constants.RANGE_START, "0");
        exchange.getIn().setHeader(S3Constants.RANGE_END, "9");
    }
})
.to("aws2-s3://mycamelbucket?amazonS3Client=#amazonS3Client&operation=getObjectRange")
.to("mock:result");

```

This operation will return an `S3Object` instance related to the `camelKey` object in `mycamelbucket` bucket, containing a the bytes from 0 to 9.

- `CreateDownloadLink`: this operation will return a download link through S3 Presigner

```

from("direct:start").process(new Processor() {

    @Override
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(S3Constants.KEY, "camelKey");
    }
})
.to("aws2-s3://mycamelbucket?
accessKey=xxx&secretKey=yyy&region=region&operation=createDownloadLink")
.to("mock:result");

```

This operation will return a download link url for the file `camel-key` in the bucket `mycamelbucket` and region `region`

5.7. STREAMING UPLOAD MODE

With the stream mode enabled users will be able to upload data to S3 without knowing ahead of time the dimension of the data, by leveraging multipart upload. The upload will be completed when: the `batchSize` has been completed or the `batchMessageNumber` has been reached. There are two possible naming strategy:

- `progressive`
With the progressive strategy each file will have the name composed by `keyName` option and a progressive counter, and eventually the file extension (if any)
- `random`.
With the random strategy a UUID will be added after `keyName` and eventually the file extension will appended.

As an example:

```

from(kafka("topic1").brokers("localhost:9092"))
    .log("Kafka Message is: ${body}")
    .to(aws2S3("camel-
bucket").streamingUploadMode(true).batchMessageNumber(25).namingStrategy(AWS2S3EndpointBu
ilderFactory.AWSS3NamingStrategyEnum.progressive).keyName("
{{kafkaTopic1}}/{{kafkaTopic1}}.txt"));

```

```

from(kafka("topic2").brokers("localhost:9092"))
    .log("Kafka Message is: ${body}")
    .to(aws2S3("camel-
bucket").streamingUploadMode(true).batchMessageNumber(25).namingStrategy(AWS2S3EndpointBu
ilderFactory.AWSS3NamingStrategyEnum.progressive).keyName("
{{kafkaTopic2}}/{{kafkaTopic2}}.txt"));

```

The default size for a batch is 1 Mb, but you can adjust it according to your requirements.

When you'll stop your producer route, the producer will take care of flushing the remaining buffered messaged and complete the upload.

In Streaming upload you'll be able restart the producer from the point where it left. It's important to note that this feature is critical only when using the progressive naming strategy.

By setting the `restartingPolicy` to `lastPart`, you will restart uploading files and contents from the last part number the producer left.

Example

1. Start the route with progressive naming strategy and keyname equals to camel.txt, with `batchMessageNumber` equals to 20, and `restartingPolicy` equals to `lastPart` - Send 70 messages.
2. Stop the route
3. On your S3 bucket you should now see 4 files: * camel.txt
 - camel-1.txt
 - camel-2.txt
 - camel-3.txt

The first three will have 20 messages, while the last one only 10.
4. Restart the route.
5. Send 25 messages.
6. Stop the route.
7. You'll now have 2 other files in your bucket: camel-5.txt and camel-6.txt, the first with 20 messages and second with 5 messages.
8. Go ahead

This won't be needed when using the random naming strategy.

On the opposite you can specify the override `restartingPolicy`. In that case you'll be able to override whatever you written before (for that particular `keyName`) on your bucket.



NOTE

In Streaming upload mode the only `keyName` option that will be taken into account is the `endpoint` option. Using the header will throw an `NPE` and this is done by design. Setting the header means potentially change the file name on each exchange and this is against the aim of the streaming upload producer. The `keyName` needs to be fixed and static. The selected naming strategy will do the rest of the of the work.

Another possibility is specifying a `streamingUploadTimeout` with `batchMessageNumber` and `batchSize` options. With this option the user will be able to complete the upload of a file after a certain time passed. In this way the upload completion will be passed on three tiers: the timeout, the number of messages and the batch size.

As an example:

```
from(kafka("topic1").brokers("localhost:9092"))
    .log("Kafka Message is: ${body}")
    .to(aws2S3("camel-
bucket").streamingUploadMode(true).batchMessageNumber(25).streamingUploadTimeout(10000).na
mingStrategy(AWS2S3EndpointBuilderFactory.AWSS3NamingStrategyEnum.progressive).keyName(
"{{kafkaTopic1}}/{{kafkaTopic1}}.txt"));
```

In this case the upload will be completed after 10 seconds.

5.8. BUCKET AUTOCREATION

With the option **autoCreateBucket** users are able to avoid the autocreation of an S3 Bucket in case it doesn't exist. The default for this option is **true**. If set to false any operation on a not-existent bucket in AWS won't be successful and an error will be returned.

5.9. MOVING STUFF BETWEEN A BUCKET AND ANOTHER BUCKET

Some users like to consume stuff from a bucket and move the content in a different one without using the `copyObject` feature of this component. If this is case for you, don't forget to remove the `bucketName` header from the incoming exchange of the consumer, otherwise the file will be always overwritten on the same original bucket.

5.10. MOVEAFTERREAD CONSUMER OPTION

In addition to `deleteAfterRead` it has been added another option, `moveAfterRead`. With this option enabled the consumed object will be moved to a target `destinationBucket` instead of being only deleted. This will require specifying the `destinationBucket` option. As example:

```
from("aws2-s3://mycamelbucket?
amazonS3Client=#amazonS3Client&moveAfterRead=true&destinationBucket=myothercamelbucket")
    .to("mock:result");
```

In this case the objects consumed will be moved to `myothercamelbucket` bucket and deleted from the original one (because of `deleteAfterRead` set to true as default).

You have also the possibility of using a key prefix/suffix while moving the file to a different bucket. The options are `destinationBucketPrefix` and `destinationBucketSuffix`.

Taking the above example, you could do something like:

```
from("aws2-s3://mycamelbucket?
amazonS3Client=#amazonS3Client&moveAfterRead=true&destinationBucket=myothercamelbucket&de
stinationBucketPrefix=RAW(pre-)&destinationBucketSuffix=RAW(-suff)")
.to("mock:result");
```

In this case the objects consumed will be moved to myothercamelbucket bucket and deleted from the original one (because of deleteAfterRead set to true as default).

So if the file name is test, in the myothercamelbucket you should see a file called pre-test-suff.

5.11. USING CUSTOMER KEY AS ENCRYPTION

We introduced also the customer key support (an alternative of using KMS). The following code shows an example.

```
String key = UUID.randomUUID().toString();
byte[] secretKey = generateSecretKey();
String b64Key = Base64.getEncoder().encodeToString(secretKey);
String b64KeyMd5 = Md5Utils.md5AsBase64(secretKey);

String awsEndpoint = "aws2-s3://mycamel?
autoCreateBucket=false&useCustomerKey=true&customerKeyId=RAW(" + b64Key +
"&customerKeyMD5=RAW(" + b64KeyMd5 + ")&customerAlgorithm=" + AES256.name());

from("direct:putObject")
.setHeader(AWS2S3Constants.KEY, constant("test.txt"))
.setBody(constant("Test"))
.to(awsEndpoint);
```

5.12. USING A POJO AS BODY

Sometimes build an AWS Request can be complex, because of multiple options. We introduce the possibility to use a POJO as body. In AWS S3 there are multiple operations you can submit, as an example for List brokers request, you can do something like:

```
from("direct:aws2-s3")
.setBody(ListObjectsRequest.builder().bucket(bucketName).build())
.to("aws2-s3://test?
amazonS3Client=#amazonS3Client&operation=listObjects&pojoRequest=true")
```

In this way you'll pass the request directly without the need of passing headers and options specifically related to this operation.

5.13. CREATE S3 CLIENT AND ADD COMPONENT TO REGISTRY

Sometimes you would want to perform some advanced configuration using AWS2S3Configuration which also allows to set the S3 client. You can create and set the S3 client in the component configuration as shown in the following example

```
String awsBucketAccessKey = "your_access_key";
String awsBucketSecretKey = "your_secret_key";
```

```

S3Client s3Client =
S3Client.builder().credentialsProvider(StaticCredentialsProvider.create(AwsBasicCredentials.create(aws
BucketAccessKey, awsBucketSecretKey)))
    .region(Region.US_EAST_1).build();

AWS2S3Configuration configuration = new AWS2S3Configuration();
configuration.setAmazonS3Client(s3Client);
configuration.setAutoDiscoverClient(true);
configuration.setBucketName("s3bucket2020");
configuration.setRegion("us-east-1");

```

Now you can configure the S3 component (using the configuration object created above) and add it to the registry in the configure method before initialization of routes.

```

AWS2S3Component s3Component = new AWS2S3Component(getContext());
s3Component.setConfiguration(configuration);
s3Component.setLazyStartProducer(true);
camelContext.addComponent("aws2-s3", s3Component);

```

Now your component will be used for all the operations implemented in camel routes.

5.14. DEPENDENCIES

Maven users will need to add the following dependency to their **pom.xml**.

pom.xml

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-s3</artifactId>
  <version>${camel-version}</version>
</dependency>

```

where **3.14.2** must be replaced by the actual version of Camel.

5.15. SPRING BOOT AUTO-CONFIGURATION

When using `aws2-s3` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-s3-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 51 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.aws2-s3.access-key</code>	Amazon AWS Access Key.		String
<code>camel.component.aws2-s3.amazon-s3-client</code>	Reference to a <code>com.amazonaws.services.s3.AmazonS3</code> in the registry. The option is a <code>software.amazon.awssdk.services.s3.S3Client</code> type.		S3Client
<code>camel.component.aws2-s3.amazon-s3-presigner</code>	An S3 Presigner for Request, used mainly in <code>createDownloadLink</code> operation. The option is a <code>software.amazon.awssdk.services.s3.presigner.S3Presigner</code> type.		S3Presigner
<code>camel.component.aws2-s3.auto-create-bucket</code>	Setting the autocreation of the S3 bucket <code>bucketName</code> . This will apply also in case of <code>moveAfterRead</code> option enabled and it will create the <code>destinationBucket</code> if it doesn't exist already.	false	Boolean
<code>camel.component.aws2-s3.autoclose-body</code>	If this option is true and <code>includeBody</code> is false, then the <code>S3Object.close()</code> method will be called on exchange completion. This option is strongly related to <code>includeBody</code> option. In case of setting <code>includeBody</code> to false and <code>autocloseBody</code> to false, it will be up to the caller to close the <code>S3Object</code> stream. Setting <code>autocloseBody</code> to true, will close the <code>S3Object</code> stream automatically.	true	Boolean
<code>camel.component.aws2-s3.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as <code>autowired</code>) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-s3.aws-kms-key-id</code>	Define the id of KMS key to use in case KMS is enabled.		String
<code>camel.component.aws2-s3.batch-message-number</code>	The number of messages composing a batch in streaming upload mode.	10	Integer
<code>camel.component.aws2-s3.batch-size</code>	The batch size (in bytes) in streaming upload mode.	1000000	Integer

Name	Description	Default	Type
<code>camel.component.aws2-s3.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.aws2-s3.configuration</code>	The component configuration. The option is a <code>org.apache.camel.component.aws2.s3.AWS2S3Configuration</code> type.		AWS2S3Configuration
<code>camel.component.aws2-s3.customer-algorithm</code>	Define the customer algorithm to use in case CustomerKey is enabled.		String
<code>camel.component.aws2-s3.customer-key-id</code>	Define the id of Customer key to use in case CustomerKey is enabled.		String
<code>camel.component.aws2-s3.customer-key-md5</code>	Define the MD5 of Customer key to use in case CustomerKey is enabled.		String
<code>camel.component.aws2-s3.delete-after-read</code>	Delete objects from S3 after they have been retrieved. The delete is only performed if the Exchange is committed. If a rollback occurs, the object is not deleted. If this option is false, then the same objects will be retrieve over and over again on the polls. Therefore you need to use the Idempotent Consumer EIP in the route to filter out duplicates. You can filter using the <code>AWS2S3Constants#BUCKET_NAME</code> and <code>AWS2S3Constants#KEY</code> headers, or only the <code>AWS2S3Constants#KEY</code> header.	true	Boolean
<code>camel.component.aws2-s3.delete-after-write</code>	Delete file object after the S3 file has been uploaded.	false	Boolean
<code>camel.component.aws2-s3.delimiter</code>	The delimiter which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.		String

Name	Description	Default	Type
<code>camel.component.aws2-s3.destination-bucket</code>	Define the destination bucket where an object must be moved when <code>moveAfterRead</code> is set to true.		String
<code>camel.component.aws2-s3.destination-bucket-prefix</code>	Define the destination bucket prefix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
<code>camel.component.aws2-s3.destination-bucket-suffix</code>	Define the destination bucket suffix to use when an object must be moved and <code>moveAfterRead</code> is set to true.		String
<code>camel.component.aws2-s3.done-file-name</code>	If provided, Camel will only consume files if a done file exists.		String
<code>camel.component.aws2-s3.enabled</code>	Whether to enable auto configuration of the <code>aws2-s3</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-s3.file-name</code>	To get the object from the bucket with the given file name.		String
<code>camel.component.aws2-s3.ignore-body</code>	If it is true, the S3 Object Body will be ignored completely, if it is set to false the S3 Object will be put in the body. Setting this to true, will override any behavior defined by <code>includeBody</code> option.	false	Boolean
<code>camel.component.aws2-s3.include-body</code>	If it is true, the S3Object exchange will be consumed and put into the body and closed. If false the S3Object stream will be put raw into the body and the headers will be set with the S3 object metadata. This option is strongly related to <code>autocloseBody</code> option. In case of setting <code>includeBody</code> to true because the S3Object stream will be consumed then it will also be closed, while in case of <code>includeBody</code> false then it will be up to the caller to close the S3Object stream. However setting <code>autocloseBody</code> to true when <code>includeBody</code> is false it will schedule to close the S3Object stream automatically on exchange completion.	true	Boolean
<code>camel.component.aws2-s3.include-folders</code>	If it is true, the folders/directories will be consumed. If it is false, they will be ignored, and Exchanges will not be created for those.	true	Boolean

Name	Description	Default	Type
<code>camel.component.aws2-s3.key-name</code>	Setting the key name for an element in the bucket through endpoint parameter.		String
<code>camel.component.aws2-s3.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-s3.move-after-read</code>	Move objects from S3 bucket to a different bucket after they have been retrieved. To accomplish the operation the destinationBucket option must be set. The copy bucket operation is only performed if the Exchange is committed. If a rollback occurs, the object is not moved.	false	Boolean
<code>camel.component.aws2-s3.multi-part-upload</code>	If it is true, camel will upload the file with multi part format, the part size is decided by the option of partSize.	false	Boolean
<code>camel.component.aws2-s3.naming-strategy</code>	The naming strategy to use in streaming upload mode.		AWSS3NamingStrategyEnum
<code>camel.component.aws2-s3.operation</code>	The operation to do in case the user don't want to do only an upload.		AWS2S3Operations
<code>camel.component.aws2-s3.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	Boolean
<code>camel.component.aws2-s3.part-size</code>	Setup the partSize which is used in multi part upload, the default size is 25M.	26214400	Long
<code>camel.component.aws2-s3.pojo-request</code>	If we want to use a POJO request as body or not.	false	Boolean

Name	Description	Default	Type
<code>camel.component.aws2-s3.policy</code>	The policy for this queue to set in the <code>com.amazonaws.services.s3.AmazonS3#setBucketPolicy()</code> method.		String
<code>camel.component.aws2-s3.prefix</code>	The prefix which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.		String
<code>camel.component.aws2-s3.proxy-host</code>	To define a proxy host when instantiating the SQS client.		String
<code>camel.component.aws2-s3.proxy-port</code>	Specify a proxy port to be used inside the client definition.		Integer
<code>camel.component.aws2-s3.proxy-protocol</code>	To define a proxy protocol when instantiating the S3 client.		Protocol
<code>camel.component.aws2-s3.region</code>	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code>) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-s3.restarting-policy</code>	The restarting policy to use in streaming upload mode.		<code>AWSS3RestartingPolicyEnum</code>
<code>camel.component.aws2-s3.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-s3.storage-class</code>	The storage class to set in the <code>com.amazonaws.services.s3.model.PutObjectRequest</code> request.		String
<code>camel.component.aws2-s3.streaming-upload-mode</code>	When stream mode is true the upload to bucket will be done in streaming.	false	Boolean
<code>camel.component.aws2-s3.streaming-upload-timeout</code>	While streaming upload mode is true, this option set the timeout to complete upload.		Long

Name	Description	Default	Type
<code>camel.component.aws2-s3.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-s3.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-s3.use-aws-k-m-s</code>	Define if KMS must be used or not.	false	Boolean
<code>camel.component.aws2-s3.use-customer-key</code>	Define if Customer Key must be used or not.	false	Boolean
<code>camel.component.aws2-s3.use-default-credentials-provider</code>	Set whether the S3 client should expect to load credentials through a default credentials provider or to expect static credentials to be passed in.	false	Boolean

CHAPTER 6. AWS SIMPLE NOTIFICATION SYSTEM (SNS)

Only producer is supported

The AWS2 SNS component allows messages to be sent to an [Amazon Simple Notification Topic](#). The implementation of the Amazon API is provided by the [AWS SDK](#).

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon SNS. More information is available at [Amazon SNS](#).

6.1. URI FORMAT

```
aws2-sns://topicNameOrArn[?options]
```

The topic will be created if they don't already exists. You can append query options to the URI in the following format, **?options=value&option2=value&...**

6.2. URI OPTIONS

6.2.1. Configuring Options

Camel components are configured on two separate levels:

- component level
- endpoint level

6.2.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

6.2.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

6.3. COMPONENT OPTIONS

The AWS Simple Notification System (SNS) component supports 24 options, which are listed below.

Name	Description	Default	Type
amazonSNSClient (producer)	Autowired To use the AmazonSNS as the client.		SnsClient
autoCreateTopic (producer)	Setting the autocreation of the topic.	false	boolean
configuration (producer)	Component configuration.		Sns2Configuration
kmsMasterKeyId (producer)	The ID of an AWS-managed customer master key (CMK) for Amazon SNS or a custom CMK.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
messageDeduplicationIdStrategy (producer)	Only for FIFO Topic. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message. Enum values: <ul style="list-style-type: none"> ● useExchangeId ● useContentBasedDeduplication 	useExchangeId	String

Name	Description	Default	Type
messageGroupIdStrategy (producer)	<p>Only for FIFO Topic. Strategy for setting the messageGroupId on the message. Can be one of the following options: useConstant, useExchangeId, usePropertyValue. For the usePropertyValue option, the value of property CamelAwsMessageGroupId will be used.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useConstant ● useExchangeId ● usePropertyValue 		String
messageStructure (producer)	The message structure to use such as json.		String
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
policy (producer)	The policy for this topic. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
proxyHost (producer)	To define a proxy host when instantiating the SNS client.		String
proxyPort (producer)	To define a proxy port when instantiating the SNS client.		Integer
proxyProtocol (producer)	<p>To define a proxy protocol when instantiating the SNS client.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
queueUrl (producer)	The queueUrl to subscribe to.		String

Name	Description	Default	Type
region (producer)	The region in which SNS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
serverSideEncryptionEnabled (producer)	Define if Server Side Encryption is enabled or not on the topic.	false	boolean
subject (producer)	The subject which is used if the message header 'CamelAwsSnsSubject' is not present.		String
subscribeSNSstoSQS (producer)	Define if the subscription between SNS Topic and SQS must be done or not.	false	boolean
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the SNS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

6.4. ENDPOINT OPTIONS

The AWS Simple Notification System (SNS) endpoint is configured using URI syntax:

```
aws2-sns:topicNameOrArn
```

with the following path and query parameters:

6.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
topicNameOrArn (producer)	Required Topic name or ARN.		String

6.4.2. Query Parameters (23 parameters)

Name	Description	Default	Type
amazonSNSClient (producer)	Autowired To use the AmazonSNS as the client.		SnsClient
autoCreateTopic (producer)	Setting the autocreation of the topic.	false	boolean
headerFilterStrategy (producer)	To use a custom HeaderFilterStrategy to map headers to/from Camel.		HeaderFilterStrategy
kmsMasterKeyId (producer)	The ID of an AWS-managed customer master key (CMK) for Amazon SNS or a custom CMK.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
messageDeduplicationIdStrategy (producer)	Only for FIFO Topic. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message. Enum values: <ul style="list-style-type: none"> • useExchangeId • useContentBasedDeduplication 	useExchangeId	String

Name	Description	Default	Type
messageGroupIdStrategy (producer)	<p>Only for FIFO Topic. Strategy for setting the messageGroupId on the message. Can be one of the following options: useConstant, useExchangeId, usePropertyValue. For the usePropertyValue option, the value of property CamelAwsMessageGroupId will be used.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useConstant ● useExchangeId ● usePropertyValue 		String
messageStructure (producer)	The message structure to use such as json.		String
overrideEndpoint (producer)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
policy (producer)	The policy for this topic. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
proxyHost (producer)	To define a proxy host when instantiating the SNS client.		String
proxyPort (producer)	To define a proxy port when instantiating the SNS client.		Integer
proxyProtocol (producer)	<p>To define a proxy protocol when instantiating the SNS client.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
queueUrl (producer)	The queueUrl to subscribe to.		String

Name	Description	Default	Type
region (producer)	The region in which SNS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
serverSideEncryptionEnabled (producer)	Define if Server Side Encryption is enabled or not on the topic.	false	boolean
subject (producer)	The subject which is used if the message header 'CamelAwsSnsSubject' is not present.		String
subscribeSNSstoSQS (producer)	Define if the subscription between SNS Topic and SQS must be done or not.	false	boolean
trustAllCertificates (producer)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (producer)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (producer)	Set whether the SNS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required SNS component options

You have to provide the `amazonSNSClient` in the Registry or your `accessKey` and `secretKey` to access the [Amazon's SNS](#).

6.5. USAGE

6.5.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`
- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.

- Web Identity Token from AWS STS.
- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#).

6.5.2. Message headers evaluated by the SNS producer

Header	Type	Description
CamelAwsSnsSubject	String	The Amazon SNS message subject. If not set, the subject from the SnsConfiguration is used.

6.5.3. Message headers set by the SNS producer

Header	Type	Description
CamelAwsSnsMessageId	String	The Amazon SNS message ID.

6.5.4. Advanced AmazonSNS configuration

If you need more control over the **SnsClient** instance configuration you can create your own instance and refer to it from the URI:

```
from("direct:start")
.to("aws2-sns://MyTopic?amazonSNSClient=#client");
```

The **#client** refers to a **AmazonSNS** in the Registry.

6.5.5. Create a subscription between an AWS SNS Topic and an AWS SQS Queue

You can create a subscription of an SQS Queue to an SNS Topic in this way:

```
from("direct:start")
.to("aws2-sns://test-camel-sns1?
amazonSNSClient=#amazonSNSClient&subscribeSNSstoSQS=true&queueUrl=https://sqs.eu-central-1.amazonaws.com/780410022472/test-camel");
```

The **#amazonSNSClient** refers to a **SnsClient** in the Registry. By specifying **subscribeSNSstoSQS** to true and a **queueUrl** of an existing SQS Queue, you'll be able to subscribe your SQS Queue to your SNS Topic.

At this point you can consume messages coming from SNS Topic through your SQS Queue

```
from("aws2-sqs://test-camel?
amazonSQSClient=#amazonSQSClient&delay=50&maxMessagesPerPoll=5")
.to(...);
```

6.6. TOPIC AUTOCREATION

With the option **autoCreateTopic** users are able to avoid the autocreation of an SNS Topic in case it doesn't exist. The default for this option is **true**. If set to false any operation on a not-existent topic in AWS won't be successful and an error will be returned.

6.7. SNS FIFO

SNS FIFO are supported. While creating the SQS queue you will subscribe to the SNS topic there is an important point to remember, you'll need to make possible for the SNS Topic to send message to the SQS Queue.

Example

Suppose you created an SNS FIFO Topic called **Order.fifo** and an SQS Queue called **QueueSub.fifo**.

In the access Policy of the **QueueSub.fifo** you should submit something like this:

```
{
  "Version": "2008-10-17",
  "Id": "__default_policy_ID",
  "Statement": [
    {
      "Sid": "__owner_statement",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::780560123482:root"
      },
      "Action": "SQS:*",
      "Resource": "arn:aws:sqs:eu-west-1:780560123482:QueueSub.fifo"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "sns.amazonaws.com"
      },
      "Action": "SQS:SendMessage",
      "Resource": "arn:aws:sqs:eu-west-1:780560123482:QueueSub.fifo",
      "Condition": {
        "ArnLike": {
          "aws:SourceArn": "arn:aws:sns:eu-west-1:780410022472:Order.fifo"
        }
      }
    }
  ]
}
```

This is a critical step to make the subscription work correctly.

6.7.1. SNS Fifo Topic Message group Id Strategy and message Deduplication Id Strategy

When sending something to the FIFO topic you'll need to always set up a message group Id strategy.

If the content-based message deduplication has been enabled on the SNS Fifo topic, there won't be the need of setting a message deduplication id strategy, otherwise you'll have to set it.

6.8. EXAMPLES

6.8.1. Producer Examples

Sending to a topic

```
from("direct:start")
  .to("aws2-sns://camel-topic?subject=The+subject+message&autoCreateTopic=true");
```

6.9. DEPENDENCIES

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-sns</artifactId>
  <version>${camel-version}</version>
</dependency>
```

where **3.14.2** must be replaced by the actual version of Camel.

6.10. SPRING BOOT AUTO-CONFIGURATION

When using aws2-sns with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-sns-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 25 options, which are listed below.

Name	Description	Default	Type
camel.component.aws2-sns.access-key	Amazon AWS Access Key.		String

Name	Description	Default	Type
<code>camel.component.aws2-sns.amazon-sns-client</code>	To use the AmazonSNS as the client. The option is a <code>software.amazon.awssdk.services.sns.SnsClient</code> type.		SnsClient
<code>camel.component.aws2-sns.auto-create-topic</code>	Setting the autocreation of the topic.	false	Boolean
<code>camel.component.aws2-sns.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.aws2-sns.configuration</code>	Component configuration. The option is a <code>org.apache.camel.component.aws2.sns.Sns2Configuration</code> type.		Sns2Configuration
<code>camel.component.aws2-sns.enabled</code>	Whether to enable auto configuration of the <code>aws2-sns</code> component. This is enabled by default.		Boolean
<code>camel.component.aws2-sns.kms-master-key-id</code>	The ID of an AWS-managed customer master key (CMK) for Amazon SNS or a custom CMK.		String
<code>camel.component.aws2-sns.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-sns.message-deduplication-id-strategy</code>	Only for FIFO Topic. Strategy for setting the <code>messageDeduplicationId</code> on the message. Can be one of the following options: <code>useExchangeId</code> , <code>useContentBasedDeduplication</code> . For the <code>useContentBasedDeduplication</code> option, no <code>messageDeduplicationId</code> will be set on the message.	<code>useExchangeId</code>	String

Name	Description	Default	Type
<code>camel.component.aws2-sns.message-group-id-strategy</code>	Only for FIFO Topic. Strategy for setting the <code>messageGroupId</code> on the message. Can be one of the following options: <code>useConstant</code> , <code>useExchangeId</code> , <code>usePropertyValue</code> . For the <code>usePropertyValue</code> option, the value of property <code>CamelAwsMessageGroupId</code> will be used.		String
<code>camel.component.aws2-sns.message-structure</code>	The message structure to use such as <code>json</code> .		String
<code>camel.component.aws2-sns.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-sns.policy</code>	The policy for this topic. Is loaded by default from classpath, but you can prefix with <code>classpath:</code> , <code>file:</code> , or <code>http:</code> to load the resource from different systems.		String
<code>camel.component.aws2-sns.proxy-host</code>	To define a proxy host when instantiating the SNS client.		String
<code>camel.component.aws2-sns.proxy-port</code>	To define a proxy port when instantiating the SNS client.		Integer
<code>camel.component.aws2-sns.proxy-protocol</code>	To define a proxy protocol when instantiating the SNS client.		Protocol
<code>camel.component.aws2-sns.queue-url</code>	The <code>queueUrl</code> to subscribe to.		String
<code>camel.component.aws2-sns.region</code>	The region in which SNS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code>) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-sns.secret-key</code>	Amazon AWS Secret Key.		String

Name	Description	Default	Type
<code>camel.component.aws2-sns.server-side-encryption-enabled</code>	Define if Server Side Encryption is enabled or not on the topic.	false	Boolean
<code>camel.component.aws2-sns.subject</code>	The subject which is used if the message header 'CamelAwsSnsSubject' is not present.		String
<code>camel.component.aws2-sns.subscribe-s-n-sto-s-q-s</code>	Define if the subscription between SNS Topic and SQS must be done or not.	false	Boolean
<code>camel.component.aws2-sns.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-sns.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-sns.use-default-credentials-provider</code>	Set whether the SNS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	Boolean

CHAPTER 7. AWS SIMPLE QUEUE SERVICE (SQS)

Both producer and consumer are supported

The AWS2 SQS component supports sending and receiving messages to [Amazon's SQS service](#).

Prerequisites

You must have a valid Amazon Web Services developer account, and be signed up to use Amazon SQS. More information is available at [Amazon SQS](#).

7.1. URI FORMAT

```
aws2-sqs://queueNameOrArn[?options]
```

The queue will be created if they don't already exists. You can append query options to the URI in the following format,

```
?options=value&option2=value&...
```

7.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

7.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

7.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

7.3. COMPONENT OPTIONS

The AWS Simple Queue Service (SQS) component supports 43 options, which are listed below.

Name	Description	Default	Type
amazonAWSHost (common)	The hostname of the Amazon AWS cloud.	amazonaws.com	String
amazonSQSClient (common)	Autowired To use the AmazonSQS as client.		SqsClient
autoCreateQueue (common)	Setting the autocreation of the queue.	false	boolean
configuration (common)	The AWS SQS default configuration.		Sqs2Configuration
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	boolean
protocol (common)	The underlying protocol used to communicate with SQS.	https	String
proxyProtocol (common)	To define a proxy protocol when instantiating the SQS client. Enum values: <ul style="list-style-type: none"> • HTTP • HTTPS 	HTTPS	Protocol
queueOwnerAWSAccountid (common)	Specify the queue owner aws account id when you need to connect the queue with different account owner.		String
region (common)	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean

Name	Description	Default	Type
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (common)	Set whether the SQS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	boolean
attributeNames (consumer)	A list of attribute names to receive when consuming. Multiple names can be separated by comma.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Allows you to use multiple threads to poll the sqs queue to increase throughput.	1	int
defaultVisibilityTimeout (consumer)	The default visibility timeout (in seconds).		Integer
deleteAfterRead (consumer)	Delete message from SQS after it has been read.	true	boolean
deleteIfFiltered (consumer)	Whether or not to send the <code>DeleteMessage</code> to the SQS queue if the exchange has property with key <code>Sqs2Constants#SQS_DELETE_FILTERED</code> (<code>CamelAwsSqsDeleteFiltered</code>) set to true.	true	boolean
extendMessageVisibility (consumer)	If enabled then a scheduled background task will keep extending the message visibility on SQS. This is needed if it takes a long time to process the message. If set to true <code>defaultVisibilityTimeout</code> must be set.	false	boolean
kmsDataKeyReusePeriodSeconds (consumer)	The length of time, in seconds, for which Amazon SQS can reuse a data key to encrypt or decrypt messages before calling AWS KMS again. An integer representing seconds, between 60 seconds (1 minute) and 86,400 seconds (24 hours). Default: 300 (5 minutes).		Integer

Name	Description	Default	Type
kmsMasterKeyId (consumer)	The ID of an AWS-managed customer master key (CMK) for Amazon SQS or a custom CMK.		String
messageAttributeName (consumer)	A list of message attribute names to receive when consuming. Multiple names can be separated by comma.		String
serverSideEncryptionEnabled (consumer)	Define if Server Side Encryption is enabled or not on the queue.	false	boolean
visibilityTimeout (consumer)	The duration (in seconds) that the received messages are hidden from subsequent retrieve requests after being retrieved by a <code>ReceiveMessage</code> request to set in the <code>com.amazonaws.services.sqs.model.SetQueueAttributesRequest</code> . This only make sense if its different from <code>defaultVisibilityTimeout</code> . It changes the queue visibility timeout attribute permanently.		Integer
waitTimeSeconds (consumer)	Duration in seconds (0 to 20) that the <code>ReceiveMessage</code> action call will wait until a message is in the queue to include in the response.		Integer
batchSeparator (producer)	Set the separator when passing a String to send batch message operation.	,	String
delaySeconds (producer)	Delay sending messages for a number of seconds.		Integer
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
messageDeduplicationIdStrategy (producer)	<p>Only for FIFO queues. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useExchangeId ● useContentBasedDeduplication 	useExchangeId	String
messageGroupIdStrategy (producer)	<p>Only for FIFO queues. Strategy for setting the messageGroupId on the message. Can be one of the following options: useConstant, useExchangeId, usePropertyValue. For the usePropertyValue option, the value of property CamelAwsMessageGroupId will be used.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useConstant ● useExchangeId ● usePropertyValue 		String
operation (producer)	<p>The operation to do in case the user don't want to send only a message.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● sendBatchMessage ● deleteMessage ● listQueues ● purgeQueue ● deleteQueue 		Sqs2Operations
autowiredEnabled (advanced)	<p>Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.</p>	true	boolean

Name	Description	Default	Type
delayQueue (advanced)	Define if you want to apply <code>delaySeconds</code> option to the queue or on single messages.	false	boolean
queueUrl (advanced)	To define the <code>queueUrl</code> explicitly. All other parameters, which would influence the <code>queueUrl</code> , are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.		String
proxyHost (proxy)	To define a proxy host when instantiating the SQS client.		String
proxyPort (proxy)	To define a proxy port when instantiating the SQS client.		Integer
maximumMessageSize (queue)	The <code>maximumMessageSize</code> (in bytes) an SQS message can contain for this queue.		Integer
messageRetentionPeriod (queue)	The <code>messageRetentionPeriod</code> (in seconds) a message will be retained by SQS for this queue.		Integer
policy (queue)	The policy for this queue. It can be loaded by default from classpath, but you can prefix with <code>classpath:</code> , <code>file:</code> , or <code>http:</code> to load the resource from different systems.		String
receiveMessageWaitTimeSeconds (queue)	If you do not specify <code>WaitTimeSeconds</code> in the request, the queue attribute <code>ReceiveMessageWaitTimeSeconds</code> is used to determine how long to wait.		Integer
redrivePolicy (queue)	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.		String
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

7.4. ENDPOINT OPTIONS

The AWS Simple Queue Service (SQS) endpoint is configured using URI syntax:

```
aws2-sqs:queueNameOrArn
```


with the following path and query parameters:

7.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
queueNameOrArn (common)	Required Queue name or ARN.		String

7.4.2. Query Parameters (61 parameters)

Name	Description	Default	Type
amazonAWSHost (common)	The hostname of the Amazon AWS cloud.	amazonaws.com	String
amazonSQSClient (common)	Autowired To use the AmazonSQS as client.		SqsClient
autoCreateQueue (common)	Setting the autocreation of the queue.	false	boolean
headerFilterStrategy (common)	To use a custom HeaderFilterStrategy to map headers to/from Camel.		HeaderFilterStrategy
overrideEndpoint (common)	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option.	false	boolean
protocol (common)	The underlying protocol used to communicate with SQS.	https	String
proxyProtocol (common)	To define a proxy protocol when instantiating the SQS client. Enum values: <ul style="list-style-type: none"> ● HTTP ● HTTPS 	HTTPS	Protocol
queueOwnerAWSAccountId (common)	Specify the queue owner aws account id when you need to connect the queue with different account owner.		String

Name	Description	Default	Type
region (common)	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
trustAllCertificates (common)	If we want to trust all certificates in case of overriding the endpoint.	false	boolean
uriEndpointOverride (common)	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
useDefaultCredentialsProvider (common)	Set whether the SQS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	boolean
attributeNames (consumer)	A list of attribute names to receive when consuming. Multiple names can be separated by comma.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Allows you to use multiple threads to poll the sqs queue to increase throughput.	1	int
defaultVisibilityTimeout (consumer)	The default visibility timeout (in seconds).		Integer
deleteAfterRead (consumer)	Delete message from SQS after it has been read.	true	boolean
deleteIfFiltered (consumer)	Whether or not to send the <code>DeleteMessage</code> to the SQS queue if the exchange has property with key <code>Sqs2Constants#SQS_DELETE_FILTERED</code> (<code>CamelAwsSqsDeleteFiltered</code>) set to true.	true	boolean

Name	Description	Default	Type
extendMessageVisibility (consumer)	If enabled then a scheduled background task will keep extending the message visibility on SQS. This is needed if it takes a long time to process the message. If set to true defaultVisibilityTimeout must be set. See details at Amazon docs.	false	boolean
kmsDataKeyReusePeriodSeconds (consumer)	The length of time, in seconds, for which Amazon SQS can reuse a data key to encrypt or decrypt messages before calling AWS KMS again. An integer representing seconds, between 60 seconds (1 minute) and 86,400 seconds (24 hours). Default: 300 (5 minutes).		Integer
kmsMasterKeyId (consumer)	The ID of an AWS-managed customer master key (CMK) for Amazon SQS or a custom CMK.		String
maxMessagesPerPoll (consumer)	Gets the maximum number of messages as a limit to poll at each polling. Is default unlimited, but use 0 or negative number to disable it as unlimited.		int
messageAttributeNames (consumer)	A list of message attribute names to receive when consuming. Multiple names can be separated by comma.		String
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
serverSideEncryptionEnabled (consumer)	Define if Server Side Encryption is enabled or not on the queue.	false	boolean
visibilityTimeout (consumer)	The duration (in seconds) that the received messages are hidden from subsequent retrieve requests after being retrieved by a ReceiveMessage request to set in the <code>com.amazonaws.services.sqs.model.SetQueueAttributesRequest</code> . This only make sense if its different from defaultVisibilityTimeout. It changes the queue visibility timeout attribute permanently.		Integer
waitTimeSeconds (consumer)	Duration in seconds (0 to 20) that the ReceiveMessage action call will wait until a message is in the queue to include in the response.		Integer

Name	Description	Default	Type
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
pollStrategy (consumer (advanced))	A pluggable org.apache.camel.PollingConsumerPollingStrategy allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
batchSeparator (producer)	Set the separator when passing a String to send batch message operation.	,	String
delaySeconds (producer)	Delay sending messages for a number of seconds.		Integer
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
messageDeduplicationIdStrategy (producer)	<p>Only for FIFO queues. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useExchangeId ● useContentBasedDeduplication 	useExchangeId	String
messageGroupIdStrategy (producer)	<p>Only for FIFO queues. Strategy for setting the messageGroupId on the message. Can be one of the following options: useConstant, useExchangeId, usePropertyValue. For the usePropertyValue option, the value of property CamelAwsMessageGroupId will be used.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● useConstant ● useExchangeId ● usePropertyValue 		String
operation (producer)	<p>The operation to do in case the user don't want to send only a message.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● sendBatchMessage ● deleteMessage ● listQueues ● purgeQueue ● deleteQueue 		Sqs2Operations
delayQueue (advanced)	Define if you want to apply delaySeconds option to the queue or on single messages.	false	boolean
queueUrl (advanced)	To define the queueUrl explicitly. All other parameters, which would influence the queueUrl, are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.		String

Name	Description	Default	Type
proxyHost (proxy)	To define a proxy host when instantiating the SQS client.		String
proxyPort (proxy)	To define a proxy port when instantiating the SQS client.		Integer
maximumMessageSize (queue)	The maximumMessageSize (in bytes) an SQS message can contain for this queue.		Integer
messageRetentionPeriod (queue)	The messageRetentionPeriod (in seconds) a message will be retained by SQS for this queue.		Integer
policy (queue)	The policy for this queue. It can be loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
receiveMessageWaitTimeSeconds (queue)	If you do not specify WaitTimeSeconds in the request, the queue attribute ReceiveMessageWaitTimeSeconds is used to determine how long to wait.		Integer
redrivePolicy (queue)	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.		String
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int

Name	Description	Default	Type
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean

Name	Description	Default	Type
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> • NANOSECONDS • MICROSECONDS • MILLISECONDS • SECONDS • MINUTES • HOURS • DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
accessKey (security)	Amazon AWS Access Key.		String
secretKey (security)	Amazon AWS Secret Key.		String

Required SQS component options

You have to provide the `amazonSQSClient` in the Registry or your `accessKey` and `secretKey` to access the [Amazon's SQS](#).

7.5. BATCH CONSUMER

This component implements the Batch Consumer.

This allows you for instance to know how many messages exists in this batch and for instance let the Aggregator aggregate this number of messages.

7.6. USAGE

7.6.1. Static credentials vs Default Credential Provider

You have the possibility of avoiding the usage of explicit static credentials, by specifying the `useDefaultCredentialsProvider` option and set it to true.

- Java system properties - `aws.accessKeyId` and `aws.secretKey`
- Environment variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`.
- Web Identity Token from AWS STS.

- The shared credentials and config files.
- Amazon ECS container credentials - loaded from the Amazon ECS if the environment variable **AWS_CONTAINER_CREDENTIALS_RELATIVE_URI** is set.
- Amazon EC2 Instance profile credentials.

For more information about this you can look at [AWS credentials documentation](#)

7.6.2. Message headers set by the SQS producer

Header	Type	Description
CamelAwsSqsMD5OfBody	String	The MD5 checksum of the Amazon SQS message.
CamelAwsSqsMessageId	String	The Amazon SQS message ID.
CamelAwsSqsDelaySeconds	Integer	The delay seconds that the Amazon SQS message can be see by others.

7.6.3. Message headers set by the SQS consumer

Header	Type	Description
CamelAwsSqsMD5OfBody	String	The MD5 checksum of the Amazon SQS message.
CamelAwsSqsMessageId	String	The Amazon SQS message ID.
CamelAwsSqsReceiptHandle	String	The Amazon SQS message receipt handle.
CamelAwsSqsMessageAttributes	Map<String, String>	The Amazon SQS message attributes.

7.6.4. Advanced AmazonSQS configuration

If your Camel Application is running behind a firewall or if you need to have more control over the **SqsClient** instance configuration, you can create your own instance:

```
from("aws2-sqs://MyQueue?amazonSQSClient=#client&delay=5000&maxMessagesPerPoll=5")
.to("mock:result");
```

7.6.5. Creating or updating an SQS Queue

In the SQS Component, when an endpoint is started, a check is executed to obtain information about the existence of the queue or not. You're able to customize the creation through the **QueueAttributeName** mapping with the **SQSConfiguration** option.

```
from("aws2-sqs://MyQueue?amazonSQSClient=#client&delay=5000&maxMessagesPerPoll=5")
.to("mock:result");
```

In this example if the **MyQueue** queue is not already created on AWS (and the **autoCreateQueue** option is set to true), it will be created with default parameters from the SQS configuration. If it's already up on AWS, the SQS configuration options will be used to override the existent AWS configuration.

7.6.6. DelayQueue VS Delay for Single message

When the option **delayQueue** is set to true, the SQS Queue will be a **DelayQueue** with the **DelaySeconds** option as delay. For more information about **DelayQueue** you can read the [AWS SQS documentation](#). One important information to take into account is the following:

- For standard queues, the per-queue delay setting is not retroactive—changing the setting doesn't affect the delay of messages already in the queue.
- For FIFO queues, the per-queue delay setting is retroactive—changing the setting affects the delay of messages already in the queue.

as stated in the official documentation. If you want to specify a delay on single messages, you can ignore the **delayQueue** option, while you can set this option to true, if you need to add a fixed delay to all messages enqueued.

7.6.7. Server Side Encryption

There is a set of Server Side Encryption attributes for a queue. The related option are **serverSideEncryptionEnabled**, **keyMasterKeyId** and **kmsDataKeyReusePeriod**. The SSE is disabled by default. You need to explicitly set the option to true and set the related parameters as queue attributes.

7.7. JMS-STYLE SELECTORS

SQS does not allow selectors, but you can effectively achieve this by using the Camel Filter EIP and setting an appropriate **visibilityTimeout**. When SQS dispatches a message, it will wait up to the visibility timeout before it will try to dispatch the message to a different consumer unless a **DeleteMessage** is received. By default, Camel will always send the **DeleteMessage** at the end of the route, unless the route ended in failure. To achieve appropriate filtering and not send the **DeleteMessage** even on successful completion of the route, use a Filter:

```
from("aws2-sqs://MyQueue?
amazonSQSClient=#client&defaultVisibilityTimeout=5000&deleteIfFiltered=false&deleteAfterRead=false
)
.filter("${header.login} == true")
.setProperty(Sqs2Constants.SQS_DELETE_FILTERED, constant(true))
.to("mock:filter");
```

In the above code, if an exchange doesn't have an appropriate header, it will not make it through the filter AND also not be deleted from the SQS queue. After 5000 milliseconds, the message will become visible to other consumers.

Note we must set the property **Sqs2Constants.SQS_DELETE_FILTERED** to **true** to instruct Camel to send the **DeleteMessage**, if being filtered.

7.8. AVAILABLE PRODUCER OPERATIONS

- single message (default)
- `sendBatchMessage`
- `deleteMessage`
- `listQueues`

7.9. SEND MESSAGE

You can set a **SendMessageBatchRequest** or an **Iterable**

```
from("direct:start")
  .setBody(constant("Camel rocks!"))
  .to("aws2-sqs://camel-1?accessKey=RAW(xxx)&secretKey=RAW(xxx)&region=eu-west-1");
```

7.10. SEND BATCH MESSAGE

You can set a **SendMessageBatchRequest** or an **Iterable**

```
from("direct:start")
  .setHeader(SqsConstants.SQS_OPERATION, constant("sendBatchMessage"))
  .process(new Processor() {
    @Override
    public void process(Exchange exchange) throws Exception {
      Collection c = new ArrayList();
      c.add("team1");
      c.add("team2");
      c.add("team3");
      c.add("team4");
      exchange.getIn().setBody(c);
    }
  })
  .to("aws2-sqs://camel-1?accessKey=RAW(xxx)&secretKey=RAW(xxx)&region=eu-west-1");
```

As result you'll get an exchange containing a **SendMessageBatchResponse** instance, that you can examine to check what messages were successfull and what not. The id set on each message of the batch will be a Random UUID.

7.11. DELETE SINGLE MESSAGE

Use **deleteMessage** operation to delete a single message. You'll need to set a receipt handle header for the message you want to delete.

```
from("direct:start")
  .setHeader(SqsConstants.SQS_OPERATION, constant("deleteMessage"))
  .setHeader(SqsConstants.RECEIPT_HANDLE, constant("123456"))
  .to("aws2-sqs://camel-1?accessKey=RAW(xxx)&secretKey=RAW(xxx)&region=eu-west-1");
```

As result you'll get an exchange containing a **DeleteMessageResponse** instance, that you can use to check if the message was deleted or not.

7.12. LIST QUEUES

Use **listQueues** operation to list queues.

```
from("direct:start")
  .setHeader(SqsConstants.SQS_OPERATION, constant("listQueues"))
  .to("aws2-sqs://camel-1?accessKey=RAW(xxx)&secretKey=RAW(xxx)&region=eu-west-1");
```

As result you'll get an exchange containing a **ListQueuesResponse** instance, that you can examine to check the actual queues.

7.13. PURGE QUEUE

Use **purgeQueue** operation to purge queue.

```
from("direct:start")
  .setHeader(SqsConstants.SQS_OPERATION, constant("purgeQueue"))
  .to("aws2-sqs://camel-1?accessKey=RAW(xxx)&secretKey=RAW(xxx)&region=eu-west-1");
```

As result you'll get an exchange containing a **PurgeQueueResponse** instance.

7.14. QUEUE AUTOCREATION

With the option **autoCreateQueue** users are able to avoid the autocreation of an SQS Queue in case it doesn't exist. The default for this option is **true**. If set to false any operation on a not-existent queue in AWS won't be successful and an error will be returned.

7.15. SEND BATCH MESSAGE AND MESSAGE DEDUPLICATION STRATEGY

In case you're using a **SendBatchMessage** Operation, you can set two different kind of Message Deduplication Strategy: - useExchangeId - useContentBasedDeduplication

The first one will use a **ExchangeIdMessageDeduplicationIdStrategy**, that will use the Exchange ID as parameter. The other one will use a **NullMessageDeduplicationIdStrategy**, that will use the body as deduplication element.

In case of send batch message operation, you'll need to use the **useContentBasedDeduplication** and on the Queue you're pointing you'll need to enable the **content based deduplication** option.

7.16. DEPENDENCIES

Maven users will need to add the following dependency to their pom.xml.

pom.xml

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-aws2-sqs</artifactId>
  <version>${camel-version}</version>
</dependency>
```

where **3.14.2** must be replaced by the actual version of Camel.

7.17. SPRING BOOT AUTO-CONFIGURATION

When using `aws2-sqs` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-aws2-sqs-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 44 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.aws2-sqs.access-key</code>	Amazon AWS Access Key.		String
<code>camel.component.aws2-sqs.amazon-aws-host</code>	The hostname of the Amazon AWS cloud.	<code>amazonaws.com</code>	String
<code>camel.component.aws2-sqs.amazon-sqs-client</code>	To use the AmazonSQS as client. The option is a <code>software.amazon.awssdk.services.sqs.SqsClient</code> type.		SqsClient
<code>camel.component.aws2-sqs.attribute-names</code>	A list of attribute names to receive when consuming. Multiple names can be separated by comma.		String
<code>camel.component.aws2-sqs.auto-create-queue</code>	Setting the autocreation of the queue.	<code>false</code>	Boolean
<code>camel.component.aws2-sqs.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	<code>true</code>	Boolean

Name	Description	Default	Type
<code>camel.component.aws2-sqs.batch-separator</code>	Set the separator when passing a String to send batch message operation.	,	String
<code>camel.component.aws2-sqs.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.aws2-sqs.concurrent-consumers</code>	Allows you to use multiple threads to poll the sqs queue to increase throughput.	1	Integer
<code>camel.component.aws2-sqs.configuration</code>	The AWS SQS default configuration. The option is a <code>org.apache.camel.component.aws2.sqs.Sqs2Configuration</code> type.		Sqs2Configuration
<code>camel.component.aws2-sqs.default-visibility-timeout</code>	The default visibility timeout (in seconds).		Integer
<code>camel.component.aws2-sqs.delay-queue</code>	Define if you want to apply <code>delaySeconds</code> option to the queue or on single messages.	false	Boolean
<code>camel.component.aws2-sqs.delay-seconds</code>	Delay sending messages for a number of seconds.		Integer
<code>camel.component.aws2-sqs.delete-after-read</code>	Delete message from SQS after it has been read.	true	Boolean
<code>camel.component.aws2-sqs.delete-if-filtered</code>	Whether or not to send the <code>DeleteMessage</code> to the SQS queue if the exchange has property with key <code>Sqs2Constants#SQS_DELETE_FILTERED</code> (<code>CamelAwsSqsDeleteFiltered</code>) set to true.	true	Boolean
<code>camel.component.aws2-sqs.enabled</code>	Whether to enable auto configuration of the <code>aws2-sqs</code> component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.aws2-sqs.extend-message-visibility</code>	If enabled then a scheduled background task will keep extending the message visibility on SQS. This is needed if it takes a long time to process the message. If set to true defaultVisibilityTimeout must be set. See details at Amazon docs.	false	Boolean
<code>camel.component.aws2-sqs.kms-data-key-reuse-period-seconds</code>	The length of time, in seconds, for which Amazon SQS can reuse a data key to encrypt or decrypt messages before calling AWS KMS again. An integer representing seconds, between 60 seconds (1 minute) and 86,400 seconds (24 hours). Default: 300 (5 minutes).		Integer
<code>camel.component.aws2-sqs.kms-master-key-id</code>	The ID of an AWS-managed customer master key (CMK) for Amazon SQS or a custom CMK.		String
<code>camel.component.aws2-sqs.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.aws2-sqs.maximum-message-size</code>	The maximumMessageSize (in bytes) an SQS message can contain for this queue.		Integer
<code>camel.component.aws2-sqs.message-attribute-names</code>	A list of message attribute names to receive when consuming. Multiple names can be separated by comma.		String
<code>camel.component.aws2-sqs.message-deduplication-id-strategy</code>	Only for FIFO queues. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message.	useExchangeId	String

Name	Description	Default	Type
<code>camel.component.aws2-sqs.message-group-id-strategy</code>	Only for FIFO queues. Strategy for setting the <code>messageGroupId</code> on the message. Can be one of the following options: <code>useConstant</code> , <code>useExchangeId</code> , <code>usePropertyValue</code> . For the <code>usePropertyValue</code> option, the value of property <code>CamelAwsMessageGroupId</code> will be used.		String
<code>camel.component.aws2-sqs.message-retention-period</code>	The <code>messageRetentionPeriod</code> (in seconds) a message will be retained by SQS for this queue.		Integer
<code>camel.component.aws2-sqs.operation</code>	The operation to do in case the user don't want to send only a message.		Sqs2Operations
<code>camel.component.aws2-sqs.override-endpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option.	false	Boolean
<code>camel.component.aws2-sqs.policy</code>	The policy for this queue. It can be loaded by default from classpath, but you can prefix with <code>classpath:</code> , <code>file:</code> , or <code>http:</code> to load the resource from different systems.		String
<code>camel.component.aws2-sqs.protocol</code>	The underlying protocol used to communicate with SQS.	https	String
<code>camel.component.aws2-sqs.proxy-host</code>	To define a proxy host when instantiating the SQS client.		String
<code>camel.component.aws2-sqs.proxy-port</code>	To define a proxy port when instantiating the SQS client.		Integer
<code>camel.component.aws2-sqs.proxy-protocol</code>	To define a proxy protocol when instantiating the SQS client.		Protocol
<code>camel.component.aws2-sqs.queue-owner-aws-account-id</code>	Specify the queue owner aws account id when you need to connect the queue with different account owner.		String

Name	Description	Default	Type
<code>camel.component.aws2-sqs.queue-url</code>	To define the <code>queueUrl</code> explicitly. All other parameters, which would influence the <code>queueUrl</code> , are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.		String
<code>camel.component.aws2-sqs.receive-message-wait-time-seconds</code>	If you do not specify <code>WaitTimeSeconds</code> in the request, the queue attribute <code>ReceiveMessageWaitTimeSeconds</code> is used to determine how long to wait.		Integer
<code>camel.component.aws2-sqs.redrive-policy</code>	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.		String
<code>camel.component.aws2-sqs.region</code>	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code>) You'll need to use the name <code>Region.EU_WEST_1.id()</code> .		String
<code>camel.component.aws2-sqs.secret-key</code>	Amazon AWS Secret Key.		String
<code>camel.component.aws2-sqs.server-side-encryption-enabled</code>	Define if Server Side Encryption is enabled or not on the queue.	false	Boolean
<code>camel.component.aws2-sqs.trust-all-certificates</code>	If we want to trust all certificates in case of overriding the endpoint.	false	Boolean
<code>camel.component.aws2-sqs.uri-endpoint-override</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option.		String
<code>camel.component.aws2-sqs.use-default-credentials-provider</code>	Set whether the SQS client should expect to load credentials on an AWS infra instance or to expect static credentials to be passed in.	false	Boolean

Name	Description	Default	Type
camel.component .aws2- sqs.visibility- timeout	The duration (in seconds) that the received messages are hidden from subsequent retrieve requests after being retrieved by a ReceiveMessage request to set in the <code>com.amazonaws.services.sqs.model.SetQueueAttributesRequest</code> . This only make sense if its different from <code>defaultVisibilityTimeout</code> . It changes the queue visibility timeout attribute permanently.		Integer
camel.component .aws2-sqs.wait- time-seconds	Duration in seconds (0 to 20) that the ReceiveMessage action call will wait until a message is in the queue to include in the response.		Integer

CHAPTER 8. AZURE STORAGE BLOB SERVICE

Both producer and consumer are supported

The Azure Storage Blob component is used for storing and retrieving blobs from [Azure Storage Blob Service](#) using **Azure APIs v12**. However in case of versions above v12, we will see if this component can adopt these changes depending on how much breaking changes can result.

Prerequisites

You must have a valid Windows Azure Storage account. More information is available at [Azure Documentation Portal](#).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-azure-storage-blob</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

8.1. URI FORMAT

```
azure-storage-blob://accountName[/containerName][?options]
```

In case of consumer, **accountName**, **containerName** are required. In case of producer, it depends on the operation that being requested, for example if operation is on a container level, for example, **createContainer**, **accountName** and **containerName** are only required, but in case of operation being requested in blob level, for example, **getBlob**, **accountName**, **containerName** and **blobName** are required.

The blob will be created if it does not already exist. You can append query options to the URI in the following format,

```
?options=value&option2=value&...
```

8.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

8.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

8.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

8.3. COMPONENT OPTIONS

The Azure Storage Blob Service component supports 31 options, which are listed below.

Name	Description	Default	Type
blobName (common)	The blob name, to consume specific blob from a container. However on producer, is only required for the operations on the blob level.		String
blobOffset (common)	Set the blob offset for the upload or download operations, default is 0.	0	long
blobType (common)	The blob type in order to initiate the appropriate settings for each blob type. Enum values: <ul style="list-style-type: none"> • blockblob • appendblob • pageblob 	blockblob	BlobType
closeStreamAfterRead (common)	Close the stream after read or keep it open, default is true.	true	boolean
configuration (common)	The component configurations.		BlobConfiguration
credentials (common)	StorageSharedKeyCredential can be injected to create the azure client, this holds the important authentication information.		StorageSharedKeyCredential

Name	Description	Default	Type
dataCount (common)	How many bytes to include in the range. Must be greater than or equal to 0 if specified.		Long
fileDir (common)	The file directory where the downloaded blobs will be saved to, this can be used in both, producer and consumer.		String
maxResultsPerPage (common)	Specifies the maximum number of blobs to return, including all BlobPrefix elements. If the request does not specify maxResultsPerPage or specifies a value greater than 5,000, the server will return up to 5,000 items.		Integer
maxRetryRequests (common)	Specifies the maximum number of additional HTTP Get requests that will be made while reading the data from a response body.	0	int
prefix (common)	Filters the results to return only blobs whose names begin with the specified prefix. May be null to return all blobs.		String
regex (common)	Filters the results to return only blobs whose names match the specified regular expression. May be null to return all if both prefix and regex are set, regex takes the priority and prefix is ignored.		String
serviceClient (common)	Autowired Client to a storage account. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. It may also be used to construct URLs to blobs and containers. This client contains operations on a service account. Operations on a container are available on BlobContainerClient through <code>BlobServiceClient#getBlobContainerClient(String)</code> , and operations on a blob are available on BlobClient through <code>BlobContainerClient#getBlobClient(String)</code> .		BlobServiceClient
timeout (common)	An optional timeout value beyond which a RuntimeException will be raised.		Duration

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
blobSequenceNumber (producer)	A user-controlled value that you can use to track requests. The value of the sequence number must be between 0 and 263 - 1. The default value is 0.	0	Long
blockListType (producer)	Specifies which type of blocks to return. Enum values: <ul style="list-style-type: none"> ● committed ● uncommitted ● all 	COMMITTED	BlockListType
changeFeedContext (producer)	When using <code>getChangeFeed</code> producer operation, this gives additional context that is passed through the Http pipeline during the service call.		Context
changeFeedEndTime (producer)	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately before the end time. Note: A few events belonging to the next hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the end time up by an hour.		OffsetDateTime
changeFeedStartTime (producer)	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately after the start time. Note: A few events belonging to the previous hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the start time down by an hour.		OffsetDateTime
closeStreamAfterWrite (producer)	Close the stream after write or keep it open, default is true.	true	boolean

Name	Description	Default	Type
commitBlockListL ater (producer)	When is set to true, the staged blocks will not be committed directly.	true	boolean
createAppendBlo b (producer)	When is set to true, the append blocks will be created when committing append blocks.	true	boolean
createPageBlob (producer)	When is set to true, the page blob will be created when uploading page blob.	true	boolean
downloadLinkExp iration (producer)	Override the default expiration (millis) of URL download link.		Long
lazyStartProduce r (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>The blob operation that can be used with this component on the producer.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● listBlobContainers ● createBlobContainer ● deleteBlobContainer ● listBlobs ● getBlob ● deleteBlob ● downloadBlobToFile ● downloadLink ● uploadBlockBlob ● stageBlockBlobList ● commitBlobBlockList ● getBlobBlockList ● createAppendBlob ● commitAppendBlob ● createPageBlob ● uploadPageBlob ● resizePageBlob ● clearPageBlob ● getPageBlobRanges 	listBlobContainers	BlobOperationsDefinition
pageBlobSize (producer)	<p>Specifies the maximum size for the page blob, up to 8 TB. The page blob size must be aligned to a 512-byte boundary.</p>	512	Long
autowiredEnabled (advanced)	<p>Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.</p>	true	boolean

Name	Description	Default	Type
accessKey (security)	Access key for the associated azure account name to be used for authentication with azure blob services.		String
sourceBlobAccessKey (security)	Source Blob Access Key: for copyblob operation, sadly, we need to have an accessKey for the source blob we want to copy Passing an accessKey as header, it's unsafe so we could set as key.		String

8.4. ENDPOINT OPTIONS

The Azure Storage Blob Service endpoint is configured using URI syntax:

```
azure-storage-blob:accountName/containerName
```

with the following path and query parameters:

8.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
accountName (common)	Azure account name to be used for authentication with azure blob services.		String
containerName (common)	The blob container name.		String

8.4.2. Query Parameters (48 parameters)

Name	Description	Default	Type
blobName (common)	The blob name, to consume specific blob from a container. However on producer, is only required for the operations on the blob level.		String
blobOffset (common)	Set the blob offset for the upload or download operations, default is 0.	0	long

Name	Description	Default	Type
blobServiceClient (common)	Client to a storage account. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. It may also be used to construct URLs to blobs and containers. This client contains operations on a service account. Operations on a container are available on BlobContainerClient through <code>getBlobContainerClient(String)</code> , and operations on a blob are available on BlobClient through <code>getBlobContainerClient(String).getBlobClient(String)</code> .		BlobServiceClient
blobType (common)	The blob type in order to initiate the appropriate settings for each blob type. Enum values: <ul style="list-style-type: none"> • blockblob • appendblob • pageblob 	blockblob	BlobType
closeStreamAfterRead (common)	Close the stream after read or keep it open, default is true.	true	boolean
credentials (common)	StorageSharedKeyCredential can be injected to create the azure client, this holds the important authentication information.		StorageSharedKeyCredential
dataCount (common)	How many bytes to include in the range. Must be greater than or equal to 0 if specified.		Long
fileDir (common)	The file directory where the downloaded blobs will be saved to, this can be used in both, producer and consumer.		String
maxResultsPerPage (common)	Specifies the maximum number of blobs to return, including all BlobPrefix elements. If the request does not specify <code>maxResultsPerPage</code> or specifies a value greater than 5,000, the server will return up to 5,000 items.		Integer
maxRetryRequests (common)	Specifies the maximum number of additional HTTP Get requests that will be made while reading the data from a response body.	0	int

Name	Description	Default	Type
prefix (common)	Filters the results to return only blobs whose names begin with the specified prefix. May be null to return all blobs.		String
regex (common)	Filters the results to return only blobs whose names match the specified regular expression. May be null to return all if both prefix and regex are set, regex takes the priority and prefix is ignored.		String
serviceClient (common)	Autowired Client to a storage account. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. It may also be used to construct URLs to blobs and containers. This client contains operations on a service account. Operations on a container are available on BlobContainerClient through BlobServiceClient#getBlobContainerClient(String), and operations on a blob are available on BlobClient through BlobContainerClient#getBlobClient(String).		BlobServiceClient
timeout (common)	An optional timeout value beyond which a RuntimeException will be raised.		Duration
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
blobSequenceNumber (producer)	A user-controlled value that you can use to track requests. The value of the sequence number must be between 0 and 263 - 1. The default value is 0.	0	Long
blockListType (producer)	Specifies which type of blocks to return. Enum values: <ul style="list-style-type: none"> ● committed ● uncommitted ● all 	COMMITTED	BlockListType
changeFeedContext (producer)	When using <code>getChangeFeed</code> producer operation, this gives additional context that is passed through the Http pipeline during the service call.		Context
changeFeedEndTime (producer)	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately before the end time. Note: A few events belonging to the next hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the end time up by an hour.		OffsetDateTime

Name	Description	Default	Type
changeFeedStartTime (producer)	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately after the start time. Note: A few events belonging to the previous hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the start time down by an hour.		OffsetDateTime
closeStreamAfterWrite (producer)	Close the stream after write or keep it open, default is true.	true	boolean
commitBlockListLater (producer)	When is set to true, the staged blocks will not be committed directly.	true	boolean
createAppendBlob (producer)	When is set to true, the append blocks will be created when committing append blocks.	true	boolean
createPageBlob (producer)	When is set to true, the page blob will be created when uploading page blob.	true	boolean
downloadLinkExpiration (producer)	Override the default expiration (millis) of URL download link.		Long
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	<p>The blob operation that can be used with this component on the producer.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● listBlobContainers ● createBlobContainer ● deleteBlobContainer ● listBlobs ● getBlob ● deleteBlob ● downloadBlobToFile ● downloadLink ● uploadBlockBlob ● stageBlockBlobList ● commitBlobBlockList ● getBlobBlockList ● createAppendBlob ● commitAppendBlob ● createPageBlob ● uploadPageBlob ● resizePageBlob ● clearPageBlob ● getPageBlobRanges 	listBlobContainers	BlobOperationsDefinition
pageBlobSize (producer)	<p>Specifies the maximum size for the page blob, up to 8 TB. The page blob size must be aligned to a 512-byte boundary.</p>	512	Long
backoffErrorThreshold (scheduler)	<p>The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.</p>		int
backoffIdleThreshold (scheduler)	<p>The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.</p>		int

Name	Description	Default	Type
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If <code>greedy</code> is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value <code>spring</code> or <code>quartz</code> for built in scheduler.	none	Object

Name	Description	Default	Type
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> ● NANoseconds ● MICROseconds ● MILLIseconds ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
accessKey (security)	Access key for the associated azure account name to be used for authentication with azure blob services.		String
sourceBlobAccessKey (security)	Source Blob Access Key: for copyblob operation, sadly, we need to have an accessKey for the source blob we want to copy Passing an accessKey as header, it's unsafe so we could set as key.		String

Required information options

To use this component, you have 3 options in order to provide the required Azure authentication information:

- Provide **accountName** and **accessKey** for your Azure account, this is the simplest way to get started. The accessKey can be generated through your Azure portal.
- Provide a [StorageSharedKeyCredential](#) instance which can be provided into **credentials** option.
- Provide a [BlobServiceClient](#) instance which can be provided into **blobServiceClient**. Note: You don't need to create a specific client, e.g: BlockBlobClient, the BlobServiceClient represents the upper level which can be used to retrieve lower level clients.

8.5. USAGE

For example, in order to download a blob content from the block blob **hello.txt** located on the **container1** in the **camelazure** storage account, use the following snippet:

```
from("azure-storage-blob://camelazure/container1?
blobName=hello.txt&accessKey=yourAccessKey").
to("file://blobdirectory");
```

8.5.1. Message headers evaluated by the component producer

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobTimeout	BlobConstants.TIMEOUT	Duration	All	An optional timeout value beyond which a <code>{@link RuntimeException}</code> will be raised.
CamelAzureStorageBlobMetadata	BlobConstants.METADATA	Map<String,String>	Operations related to container and blob	Metadata to associate with the container or blob.
CamelAzureStorageBlobPublicAccessType	BlobConstants.PUBLIC_ACCESS_TYPE	PublicAccessType	createContainer	Specifies how the data in this container is available to the public. Pass null for no public access.
CamelAzureStorageBlobRequestCondition	BlobConstants.BLOB_REQUEST_CONDITION	BlobRequestConditions	Operations related to container and blob	This contains values which will restrict the successful operation of a variety of requests to the conditions present. These conditions are entirely optional.
CamelAzureStorageBlobListDetails	BlobConstants.BLOB_LIST_DETAILS	BlobListDetails	listBlobs	The details for listing specific blobs

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobPrefix	BlobConstants.PREFIX	String	listBlobs, getBlob	Filters the results to return only blobs whose names begin with the specified prefix. May be null to return all blobs.
CamelAzureStorageBlobMaxResultsPerPage	BlobConstants.MAX_RESULTS_PER_PAGE	Integer	listBlobs	Specifies the maximum number of blobs to return, including all BlobPrefix elements. If the request does not specify maxResultsPerPage or specifies a value greater than 5,000, the server will return up to 5,000 items.
CamelAzureStorageBlobListBlobOptions	BlobConstants.LIST_BLOB_OPTIONS	ListBlobsOptions	listBlobs	Defines options available to configure the behavior of a call to listBlobsFlatSegment on a {@link BlobContainerClient} object.
CamelAzureStorageBlobHttpHeaders	BlobConstants.BLOB_HTTP_HEADERS	BlobHttpHeaders	uploadBlockBlob, commitBlobBlockList, createAppendBlob, createPageBlob	Additional parameters for a set of operations.
CamelAzureStorageBlobAccessTier	BlobConstants.ACCESS_TIER	AccessTier	uploadBlockBlob, commitBlobBlockList	Defines values for AccessTier.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobContentMD5	BlobConstants.CONTENT_MD5	byte[]	Most operations related to upload blob	An MD5 hash of the block content. This hash is used to verify the integrity of the block during transport. When this header is specified, the storage service compares the hash of the content that has arrived with this header value. Note that this MD5 hash is not stored with the blob. If the two hashes do not match, the operation will fail.
CamelAzureStorageBlobPageBlobRange	BlobConstants.PAGE_BLOB_RANGE	PageRange	Operations related to page blob	A {@link PageRange} object. Given that pages must be aligned with 512-byte boundaries, the start offset must be a modulus of 512 and the end offset must be a modulus of 512 - 1. Examples of valid byte ranges are 0-511, 512-1023, etc.
CamelAzureStorageBlobCommitBlobBlockListLater	BlobConstants.COMMIT_BLOCK_LIST_LATER	boolean	stageBlockBlobList	When is set to true , the staged blocks will not be committed directly.
CamelAzureStorageBlobCreateAppendBlob	BlobConstants.CREATE_APPEND_BLOB	boolean	commitAppendBlob	When is set to true , the append blocks will be created when committing append blocks.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobCreatePageBlob	BlobConstants.CREATE_PAGE_BLOB	boolean	uploadPageBlob	When is set to true , the page blob will be created when uploading page blob.
CamelAzureStorageBlobBlockListType	BlobConstants.BLOCK_LIST_TYPE	BlockListType	getBlobBlockList	Specifies which type of blocks to return.
CamelAzureStorageBlobPageBlobSize	BlobConstants.PAGE_BLOB_SIZE	Long	createPageBlob , resizePageBlob	Specifies the maximum size for the page blob, up to 8 TB. The page blob size must be aligned to a 512-byte boundary.
CamelAzureStorageBlobSequenceNumber	BlobConstants.BLOB_SEQUENCE_NUMBER	Long	createPageBlob	A user-controlled value that you can use to track requests. The value of the sequence number must be between 0 and $2^{63} - 1$. The default value is 0.
CamelAzureStorageBlobDeleteSnapshotsOptionType	BlobConstants.DELETE_SNAPSHOT_OPTION_TYPE	DeleteSnapshotOptionType	deleteBlob	Specifies the behavior for deleting the snapshots on this blob. <code>Include</code> will delete the base blob and all snapshots. <code>Only</code> will delete only the snapshots. If a snapshot is being deleted, you must pass null.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobListBlobContainersOptions	BlobConstants.LIST_BLOB_CONTAINERS_OPTIONS	ListBlobContainersOptions	listBlobContainers	A {@link ListBlobContainersOptions} which specifies what data should be returned by the service.
CamelAzureStorageBlobParallelTransferOptions	BlobConstants.PARALLEL_TRANSFER_OPTIONS	ParallelTransferOptions	downloadBlobToFile	{@link ParallelTransferOptions} to use to download to file. Number of parallel transfers parameter is ignored.
CamelAzureStorageBlobFileDir	BlobConstants.FILE_DIR	String	downloadBlobToFile	The file directory where the downloaded blobs will be saved to.
CamelAzureStorageBlobDownloadLinkExpiration	BlobConstants.DOWNLOAD_LINK_EXPIRATION	Long	downloadLink	Override the default expiration (millis) of URL download link.
CamelAzureStorageBlobBlobName	BlobConstants.BLOB_NAME	String	Operations related to blob	Override/set the blob name on the exchange headers.
CamelAzureStorageBlobContainerName	BlobConstants.BLOB_CONTAINER_NAME	String	Operations related to container and blob	Override/set the container name on the exchange headers.
CamelAzureStorageBlobOperation	BlobConstants.BLOB_OPERATION	BlobOperationsDefinition	All	Specify the producer operation to execute, please see the doc on this page related to producer operation.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobRegex	BlobConstants. REGEX	String	listBlobs, getBlob	Filters the results to return only blobs whose names match the specified regular expression. May be null to return all. If both prefix and regex are set, regex takes the priority and prefix is ignored.
CamelAzureStorageBlobChangeFeedStartTime	BlobConstants. CHANGE_FEED _START_TIME	OffsetDateTime	getChangeFeed	It filters the results to return events approximately after the start time. Note: A few events belonging to the previous hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the start time down by an hour.
CamelAzureStorageBlobChangeFeedEndTime	BlobConstants. CHANGE_FEED _END_TIME	OffsetDateTime	getChangeFeed	It filters the results to return events approximately before the end time. Note: A few events belonging to the next hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the end time up by an hour.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageBlobChangeFeedContext	BlobConstants.CHANGE_FEED_CONTEXT	Context	getChangeFeed	This gives additional context that is passed through the Http pipeline during the service call.
CamelAzureStorageBlobSourceBlobAccountName	BlobConstants.SOURCE_BLOB_ACCOUNT_NAME	String	copyBlob	The source blob account name to be used as source account name in a copy blob operation
CamelAzureStorageBlobSourceBlobContainerName	BlobConstants.SOURCE_BLOB_CONTAINER_NAME	String	copyBlob	The source blob container name to be used as source container name in a copy blob operation

8.5.2. Message headers set by either component producer or consumer

Header	Variable Name	Type	Description
CamelAzureStorageBlobAccessTier	BlobConstants.ACCESS_TIER	AccessTier	Access tier of the blob.
CamelAzureStorageBlobAccessTierChangeTime	BlobConstants.ACCESS_TIER_CHANGE_TIME	OffsetDateTime	Datetime when the access tier of the blob last changed.
CamelAzureStorageBlobArchiveStatus	BlobConstants.ARCHIVE_STATUS	ArchiveStatus	Archive status of the blob.
CamelAzureStorageBlobCreationTime	BlobConstants.CREATION_TIME	OffsetDateTime	Creation time of the blob.
CamelAzureStorageBlobSequenceNumber	BlobConstants.BLOB_SEQUENCE_NUMBER	Long	The current sequence number for a page blob.

Header	Variable Name	Type	Description
CamelAzureStorageBlobBlobSize	BlobConstants.BLOB_SIZE	long	The size of the blob.
CamelAzureStorageBlobBlobType	BlobConstants.BLOB_TYPE	BlobType	The type of the blob.
CamelAzureStorageBlobCacheControl	BlobConstants.CACHE_CONTROL	String	Cache control specified for the blob.
CamelAzureStorageBlobCommittedBlockCount	BlobConstants.COMMITTED_BLOCK_COUNT	Integer	Number of blocks committed to an append blob
CamelAzureStorageBlobContentDisposition	BlobConstants.CONTENT_DISPOSITION	String	Content disposition specified for the blob.
CamelAzureStorageBlobContentEncoding	BlobConstants.CONTENT_ENCODING	String	Content encoding specified for the blob.
CamelAzureStorageBlobContentLanguage	BlobConstants.CONTENT_LANGUAGE	String	Content language specified for the blob.
CamelAzureStorageBlobContentMd5	BlobConstants.CONTENT_MD5	byte[]	Content MD5 specified for the blob.
CamelAzureStorageBlobContentType	BlobConstants.CONTENT_TYPE	String	Content type specified for the blob.
CamelAzureStorageBlobCopyCompletionTime	BlobConstants.COPY_COMPLETION_TIME	OffsetDateTime	Datetime when the last copy operation on the blob completed.
CamelAzureStorageBlobCopyDestinationSnapshot	BlobConstants.COPY_DESTINATION_SNAPSHOT	String	Snapshot identifier of the last incremental copy snapshot for the blob.
CamelAzureStorageBlobCopyId	BlobConstants.COPY_ID	String	Identifier of the last copy operation performed on the blob.

Header	Variable Name	Type	Description
CamelAzureStorageBlobCopyProgress	BlobConstants.COPY_PROGRESS	String	Progress of the last copy operation performed on the blob.
CamelAzureStorageBlobCopySource	BlobConstants.COPY_SOURCE	String	Source of the last copy operation performed on the blob.
CamelAzureStorageBlobCopyStatus	BlobConstants.COPY_STATUS	CopyStatusType	Status of the last copy operation performed on the blob.
CamelAzureStorageBlobCopyStatusDescription	BlobConstants.COPY_STATUS_DESCRIPTION	String	Description of the last copy operation on the blob.
CamelAzureStorageBlobETag	BlobConstants.E_TAG	String	The E Tag of the blob
CamelAzureStorageBlobsAccessTierInferred	BlobConstants.IS_ACCESS_TIER_INFERRRED	boolean	Flag indicating if the access tier of the blob was inferred from properties of the blob.
CamelAzureStorageBlobsIncrementalCopy	BlobConstants.IS_INCREMENTAL_COPY	boolean	Flag indicating if the blob was incrementally copied.
CamelAzureStorageBlobsServerEncrypted	BlobConstants.IS_SERVER_ENCRYPTED	boolean	Flag indicating if the blob's content is encrypted on the server.
CamelAzureStorageBlobLastModified	BlobConstants.LAST_MODIFIED	OffsetDateTime	Datetime when the blob was last modified.
CamelAzureStorageBlobLeaseDuration	BlobConstants.LEASE_DURATION	LeaseDurationType	Type of lease on the blob.
CamelAzureStorageBlobLeaseState	BlobConstants.LEASE_STATE	LeaseStateType	State of the lease on the blob.
CamelAzureStorageBlobLeaseStatus	BlobConstants.LEASE_STATUS	LeaseStatusType	Status of the lease on the blob.

Header	Variable Name	Type	Description
CamelAzureStorageBlobMetadata	BlobConstants.METADATA	Map<String, String>	Additional metadata associated with the blob.
CamelAzureStorageBlobAppendOffset	BlobConstants.APPEND_OFFSET	String	The offset at which the block was committed to the block blob.
CamelAzureStorageBlobFilename	BlobConstants.FILE_NAME	String	The downloaded filename from the operation downloadBlobToFile .
CamelAzureStorageBlobDownloadLink	BlobConstants.DOWNLOAD_LINK	String	The download link generated by downloadLink operation.
CamelAzureStorageBlobRawHttpHeaders	BlobConstants.RAW_HTTP_HEADERS	HttpHeaders	Returns non-parsed httpHeaders that can be used by the user.

8.5.3. Advanced Azure Storage Blob configuration

If your Camel Application is running behind a firewall or if you need to have more control over the **BlobServiceClient** instance configuration, you can create your own instance:

```
StorageSharedKeyCredential credential = new StorageSharedKeyCredential("yourAccountName",
"yourAccessKey");
String uri = String.format("https://%s.blob.core.windows.net", "yourAccountName");

BlobServiceClient client = new BlobServiceClientBuilder()
    .endpoint(uri)
    .credential(credential)
    .buildClient();
// This is camel context
context.getRegistry().bind("client", client);
```

Then refer to this instance in your Camel **azure-storage-blob** component configuration:

```
from("azure-storage-blob://cameldev/container1?blobName=myblob&serviceClient=#client")
.to("mock:result");
```

8.5.4. Automatic detection of BlobServiceClient client in registry

The component is capable of detecting the presence of an **BlobServiceClient** bean into the registry. If it's the only instance of that type it will be used as client and you won't have to define it as uri parameter, like the example above. This may be really useful for smarter configuration of the endpoint.

8.5.5. Azure Storage Blob Producer operations

Camel Azure Storage Blob component provides wide range of operations on the producer side:

Operations on the service level

For these operations, **accountName** is required.

Operation	Description
listBlobContainers	Get the content of the blob. You can restrict the output of this operation to a blob range.
getChangeFeed	Returns transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes.

Operations on the container level

For these operations, **accountName** and **containerName** are required.

Operation	Description
createBlobContainer	Creates a new container within a storage account. If a container with the same name already exists, the producer will ignore it.
deleteBlobContainer	Deletes the specified container in the storage account. If the container doesn't exist the operation fails.
listBlobs	Returns a list of blobs in this container, with folder structures flattened.

Operations on the blob level

For these operations, **accountName**, **containerName** and **blobName** are required.

Operation	Blob Type	Description
getBlob	Common	Get the content of the blob. You can restrict the output of this operation to a blob range.
deleteBlob	Common	Delete a blob.
downloadBlobToFile	Common	Downloads the entire blob into a file specified by the path. The file will be created and must not exist, if the file already exists a <code>{@link FileAlreadyExistsException}</code> will be thrown.
downloadLink	Common	Generates the download link for the specified blob using shared access signatures (SAS). This by default only limit to 1hour of allowed access. However, you can override the default expiration duration through the headers.

Operation	Blob Type	Description
uploadBlockBlob	BlockBlob	Creates a new block blob, or updates the content of an existing block blob. Updating an existing block blob overwrites any existing metadata on the blob. Partial updates are not supported with PutBlob; the content of the existing blob is overwritten with the new content.
stageBlockBlobList	BlockBlob	Uploads the specified block to the block blob's "staging area" to be later committed by a call to <code>commitBlobBlockList</code> . However in case header CamelAzureStorageBlobCommitBlockListLater or config commitBlockListLater is set to false, this will commit the blocks immediately after staging the blocks.
commitBlobBlockList	BlockBlob	Writes a blob by specifying the list of block IDs that are to make up the blob. In order to be written as part of a blob, a block must have been successfully written to the server in a prior stageBlockBlobList operation. You can call commitBlobBlockList to update a blob by uploading only those blocks that have changed, then committing the new and existing blocks together. Any blocks not specified in the block list and permanently deleted.
getBlobBlockList	BlockBlob	Returns the list of blocks that have been uploaded as part of a block blob using the specified block list filter.
createAppendBlob	AppendBlob	Creates a 0-length append blob. Call <code>commitAppendBlob</code> operation to append data to an append blob.
commitAppendBlob	AppendBlob	Commits a new block of data to the end of the existing append blob. In case of header CamelAzureStorageBlobCreateAppendBlob or config createAppendBlob is set to true, it will attempt to create the appendBlob through internal call to createAppendBlob operation first before committing.
createPageBlob	PageBlob	Creates a page blob of the specified length. Call uploadPageBlob operation to upload data data to a page blob.

Operation	Blob Type	Description
uploadPageBlob	PageBlob	Writes one or more pages to the page blob. The write size must be a multiple of 512. In case of header CamelAzureStorageBlobCreatePageBlob or config createPageBlob is set to true, it will attempt to create the appendBlob through internal call to createPageBlob operation first before uploading.
resizePageBlob	PageBlob	Resizes the page blob to the specified size (which must be a multiple of 512).
clearPageBlob	PageBlob	Frees the specified pages from the page blob. The size of the range must be a multiple of 512.
getPageBlobRanges	PageBlob	Returns the list of valid page ranges for a page blob or snapshot of a page blob.
copyBlob	Common	Copy a blob from one container to another one, even from different accounts.

Refer to the example section in this page to learn how to use these operations into your camel application.

8.5.6. Consumer Examples

To consume a blob into a file using file component, this can be done like this:

```
from("azure-storage-blob://camelazure/container1?
blobName=hello.txt&accountName=yourAccountName&accessKey=yourAccessKey").
to("file://blobdirectory");
```

However, you can also write to file directly without using the file component, you will need to specify **fileDir** folder path in order to save your blob in your machine.

```
from("azure-storage-blob://camelazure/container1?
blobName=hello.txt&accountName=yourAccountName&accessKey=yourAccessKey&fileDir=/var/to/awes
ome/dir").
to("mock:results");
```

Also, the component supports batch consumer, hence you can consume multiple blobs with only specifying the container name, the consumer will return multiple exchanges depending on the number of the blobs in the container.

Example

```
from("azure-storage-blob://camelazure/container1?
accountName=yourAccountName&accessKey=yourAccessKey&fileDir=/var/to/awesome/dir").
to("mock:results");
```

8.5.7. Producer Operations Examples

- **listBlobContainers**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.LIST_BLOB_CONTAINERS_OPTIONS, new
ListBlobContainersOptions().setMaxResultsPerPage(10));
  })
  .to("azure-storage-blob://camelazure?operation=listBlobContainers&client&serviceClient=#client")
  .to("mock:result");
```

- **createBlobContainer**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_CONTAINER_NAME, "newContainerName");
  })
  .to("azure-storage-blob://camelazure/container1?
operation=createBlobContainer&serviceClient=#client")
  .to("mock:result");
```

- **deleteBlobContainer:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_CONTAINER_NAME, "overridenName");
  })
  .to("azure-storage-blob://camelazure/container1?
operation=deleteBlobContainer&serviceClient=#client")
  .to("mock:result");
```

- **listBlobs:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_CONTAINER_NAME, "overridenName");
  })
  .to("azure-storage-blob://camelazure/container1?operation=listBlobs&serviceClient=#client")
  .to("mock:result");
```

- **getBlob:**

We can either set an **outputStream** in the exchange body and write the data to it. E.g:

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_CONTAINER_NAME, "overridenName");

    // set our body
    exchange.getIn().setBody(outputStream);
  })
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=getBlob&serviceClient=#client")
  .to("mock:result");
```

If we don't set a body, then this operation will give us an **InputStream** instance which can proceed further downstream:

```
from("direct:start")
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=getBlob&serviceClient=#client")
  .process(exchange -> {
    InputStream inputStream = exchange.getMessage().getBody(InputStream.class);
    // We use Apache common IO for simplicity, but you are free to do whatever dealing
    // with inputStream
    System.out.println(IOUtils.toString(inputStream, StandardCharsets.UTF_8.name()));
  })
  .to("mock:result");
```

- **deleteBlob:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_NAME, "overridenName");
  })
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=deleteBlob&serviceClient=#client")
  .to("mock:result");
```

- **downloadBlobToFile:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_NAME, "overridenName");
  })
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=downloadBlobToFile&fileDir=/var/mydir&serviceClient=#client")
  .to("mock:result");
```

- **downloadLink**

```
from("direct:start")
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=downloadLink&serviceClient=#client")
  .process(exchange -> {
    String link = exchange.getMessage().getHeader(BlobConstants.DOWNLOAD_LINK,
String.class);
    System.out.println("My link " + link);
  })
  .to("mock:result");
```

- **uploadBlockBlob**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(BlobConstants.BLOB_NAME, "overridenName");
    exchange.getIn().setBody("Block Blob");
  })
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=uploadBlockBlob&serviceClient=#client")
  .to("mock:result");
```

- **stageBlockBlobList**

```
from("direct:start")
  .process(exchange -> {
    final List<BlobBlock> blocks = new LinkedList<>();
    blocks.add(BlobBlock.createBlobBlock(new ByteArrayInputStream("Hello".getBytes())));
    blocks.add(BlobBlock.createBlobBlock(new ByteArrayInputStream("From".getBytes())));
    blocks.add(BlobBlock.createBlobBlock(new ByteArrayInputStream("Camel".getBytes())));

    exchange.getIn().setBody(blocks);
  })
  .to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=stageBlockBlobList&serviceClient=#client")
  .to("mock:result");
```

- **commitBlockBlobList**

```
from("direct:start")
  .process(exchange -> {
    // We assume here you have the knowledge of these blocks you want to commit
    final List<Block> blocksIds = new LinkedList<>();
    blocksIds.add(new Block().setName("id-1"));
    blocksIds.add(new Block().setName("id-2"));
    blocksIds.add(new Block().setName("id-3"));

    exchange.getIn().setBody(blocksIds);
  })
```



```
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=commitBlockBlobList&serviceClient=#client")
.to("mock:result");
```

- **getBlobBlockList**

```
from("direct:start")
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=getBlobBlockList&serviceClient=#client")
.log("${body}")
.to("mock:result");
```

- **createAppendBlob**

```
from("direct:start")
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=createAppendBlob&serviceClient=#client")
.to("mock:result");
```

- **commitAppendBlob**

```
from("direct:start")
.process(exchange -> {
    final String data = "Hello world from my awesome tests!";
    final InputStream dataStream = new
    ByteArrayInputStream(data.getBytes(StandardCharsets.UTF_8));

    exchange.getIn().setBody(dataStream);

    // of course you can set whatever headers you like, refer to the headers section to learn more
})
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=commitAppendBlob&serviceClient=#client")
.to("mock:result");
```

- **createPageBlob**

```
from("direct:start")
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=createPageBlob&serviceClient=#client")
.to("mock:result");
```

- **uploadPageBlob**

```
from("direct:start")
.process(exchange -> {
    byte[] dataBytes = new byte[512]; // we set range for the page from 0-511
    new Random().nextBytes(dataBytes);
    final InputStream dataStream = new ByteArrayInputStream(dataBytes);
    final PageRange pageRange = new PageRange().setStart(0).setEnd(511);

    exchange.getIn().setHeader(BlobConstants.PAGE_BLOB_RANGE, pageRange);
    exchange.getIn().setBody(dataStream);
})
```

```
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=uploadPageBlob&serviceClient=#client")
.to("mock:result");
```

- **resizePageBlob**

```
from("direct:start")
.process(exchange -> {
    final PageRange pageRange = new PageRange().setStart(0).setEnd(511);

    exchange.getIn().setHeader(BlobConstants.PAGE_BLOB_RANGE, pageRange);
})
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=resizePageBlob&serviceClient=#client")
.to("mock:result");
```

- **clearPageBlob**

```
from("direct:start")
.process(exchange -> {
    final PageRange pageRange = new PageRange().setStart(0).setEnd(511);

    exchange.getIn().setHeader(BlobConstants.PAGE_BLOB_RANGE, pageRange);
})
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=clearPageBlob&serviceClient=#client")
.to("mock:result");
```

- **getPageBlobRanges**

```
from("direct:start")
.process(exchange -> {
    final PageRange pageRange = new PageRange().setStart(0).setEnd(511);

    exchange.getIn().setHeader(BlobConstants.PAGE_BLOB_RANGE, pageRange);
})
.to("azure-storage-blob://camelazure/container1?
blobName=blob&operation=getPageBlobRanges&serviceClient=#client")
.log("${body}")
.to("mock:result");
```

- **copyBlob**

```
from("direct:copyBlob")
.process(exchange -> {
    exchange.getIn().setHeader(BlobConstants.BLOB_NAME, "file.txt");
    exchange.getMessage().setHeader(BlobConstants.SOURCE_BLOB_CONTAINER_NAME,
"containerblob1");
    exchange.getMessage().setHeader(BlobConstants.SOURCE_BLOB_ACCOUNT_NAME,
"account");
})
.to("azure-storage-blob://account/containerblob2?
operation=copyBlob&sourceBlobAccessKey=RAW(accessKey)")
.to("mock:result");
```

In this way the file.txt in the container containerblob1 of the account 'account', will be copied to the container containerblob2 of the same account.

8.5.8. Development Notes (Important)

All integration tests use [Testcontainers](#) and run by default. Obtaining of Azure accessKey and accountName is needed to be able to run all integration tests using Azure services. In addition to the mocked unit tests you **will need to run the integration tests with every change you make or even client upgrade as the Azure client can break things even on minor versions upgrade.** To run the integration tests, on this component directory, run the following maven command:

```
mvn verify -PfullTests -DaccountName=myacc -DaccessKey=mykey
```

Whereby **accountName** is your Azure account name and **accessKey** is the access key being generated from Azure portal.

8.6. SPRING BOOT AUTO-CONFIGURATION

When using azure-storage-blob with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-azure-storage-blob-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 32 options, which are listed below.

Name	Description	Default	Type
camel.component.azure-storage-blob.access-key	Access key for the associated azure account name to be used for authentication with azure blob services.		String
camel.component.azure-storage-blob.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.azure-storage-blob.blob-name	The blob name, to consume specific blob from a container. However on producer, is only required for the operations on the blob level.		String
camel.component.azure-storage-blob.blob-offset	Set the blob offset for the upload or download operations, default is 0.	0	Long

Name	Description	Default	Type
<code>camel.component.azure-storage-blob.blob-sequence-number</code>	A user-controlled value that you can use to track requests. The value of the sequence number must be between 0 and 263 - 1. The default value is 0.	0	Long
<code>camel.component.azure-storage-blob.blob-type</code>	The blob type in order to initiate the appropriate settings for each blob type.		BlobType
<code>camel.component.azure-storage-blob.block-list-type</code>	Specifies which type of blocks to return.		BlockListType
<code>camel.component.azure-storage-blob.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which means any exceptions occurred while the consumer is trying to pick up incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.azure-storage-blob.change-feed-context</code>	When using <code>getChangeFeed</code> producer operation, this gives additional context that is passed through the Http pipeline during the service call. The option is a <code>com.azure.core.util.Context</code> type.		Context
<code>camel.component.azure-storage-blob.change-feed-end-time</code>	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately before the end time. Note: A few events belonging to the next hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the end time up by an hour. The option is a <code>java.time.OffsetDateTime</code> type.		OffsetDateTime
<code>camel.component.azure-storage-blob.change-feed-start-time</code>	When using <code>getChangeFeed</code> producer operation, this filters the results to return events approximately after the start time. Note: A few events belonging to the previous hour can also be returned. A few events belonging to this hour can be missing; to ensure all events from the hour are returned, round the start time down by an hour. The option is a <code>java.time.OffsetDateTime</code> type.		OffsetDateTime

Name	Description	Default	Type
<code>camel.component.azure-storage-blob.close-stream-after-read</code>	Close the stream after read or keep it open, default is true.	true	Boolean
<code>camel.component.azure-storage-blob.close-stream-after-write</code>	Close the stream after write or keep it open, default is true.	true	Boolean
<code>camel.component.azure-storage-blob.commit-block-list-later</code>	When is set to true, the staged blocks will not be committed directly.	true	Boolean
<code>camel.component.azure-storage-blob.configuration</code>	The component configurations. The option is a <code>org.apache.camel.component.azure.storage.blob.BlobConfiguration</code> type.		BlobConfiguration
<code>camel.component.azure-storage-blob.create-append-blob</code>	When is set to true, the append blocks will be created when committing append blocks.	true	Boolean
<code>camel.component.azure-storage-blob.create-page-blob</code>	When is set to true, the page blob will be created when uploading page blob.	true	Boolean
<code>camel.component.azure-storage-blob.credentials</code>	<code>StorageSharedKeyCredential</code> can be injected to create the azure client, this holds the important authentication information. The option is a <code>com.azure.storage.common.StorageSharedKeyCredential</code> type.		StorageSharedKeyCredential
<code>camel.component.azure-storage-blob.data-count</code>	How many bytes to include in the range. Must be greater than or equal to 0 if specified.		Long
<code>camel.component.azure-storage-blob.download-link-expiration</code>	Override the default expiration (millis) of URL download link.		Long

Name	Description	Default	Type
<code>camel.component.azure-storage-blob.enabled</code>	Whether to enable auto configuration of the azure-storage-blob component. This is enabled by default.		Boolean
<code>camel.component.azure-storage-blob.file-dir</code>	The file directory where the downloaded blobs will be saved to, this can be used in both, producer and consumer.		String
<code>camel.component.azure-storage-blob.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.azure-storage-blob.max-results-per-page</code>	Specifies the maximum number of blobs to return, including all BlobPrefix elements. If the request does not specify maxResultsPerPage or specifies a value greater than 5,000, the server will return up to 5,000 items.		Integer
<code>camel.component.azure-storage-blob.max-retry-requests</code>	Specifies the maximum number of additional HTTP Get requests that will be made while reading the data from a response body.	0	Integer
<code>camel.component.azure-storage-blob.operation</code>	The blob operation that can be used with this component on the producer.		BlobOperationsDefinition
<code>camel.component.azure-storage-blob.page-blob-size</code>	Specifies the maximum size for the page blob, up to 8 TB. The page blob size must be aligned to a 512-byte boundary.	512	Long
<code>camel.component.azure-storage-blob.prefix</code>	Filters the results to return only blobs whose names begin with the specified prefix. May be null to return all blobs.		String
<code>camel.component.azure-storage-blob.regex</code>	Filters the results to return only blobs whose names match the specified regular expression. May be null to return all if both prefix and regex are set, regex takes the priority and prefix is ignored.		String

Name	Description	Default	Type
camel.component.azure-storage-blob.service-client	Client to a storage account. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. It may also be used to construct URLs to blobs and containers. This client contains operations on a service account. Operations on a container are available on BlobContainerClient through BlobServiceClient#getBlobContainerClient(String), and operations on a blob are available on BlobClient through BlobContainerClient#getBlobClient(String). The option is a com.azure.storage.blob.BlobServiceClient type.		BlobServiceClient
camel.component.azure-storage-blob.source-blob-access-key	Source Blob Access Key: for copyblob operation, sadly, we need to have an accessKey for the source blob we want to copy Passing an accessKey as header, it's unsafe so we could set as key.		String
camel.component.azure-storage-blob.timeout	An optional timeout value beyond which a RuntimeException will be raised. The option is a java.time.Duration type.		Duration

CHAPTER 9. AZURE STORAGE QUEUE SERVICE

Both producer and consumer are supported

The Azure Storage Queue component supports storing and retrieving the messages to/from [Azure Storage Queue](#) service using **Azure APIs v12**. However in case of versions above v12, we will see if this component can adopt these changes depending on how much breaking changes can result.

Prerequisites

You must have a valid Windows Azure Storage account. More information is available at [Azure Documentation Portal](#).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-azure-storage-queue</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

9.1. URI FORMAT

```
azure-storage-queue://accountName[/queueName][?options]
```

In case of consumer, `accountName` and `queueName` are required. In case of producer, it depends on the operation that being requested, for example if operation is on a service level, e.b: `listQueues`, only `accountName` is required, but in case of operation being requested on the queue level, for example, `createQueue`, `sendMessage`.. etc, both `accountName` and `queueName` are required.

The queue will be created if it does not already exist. You can append query options to the URI in the following format,

?options=value&option2=value&...

9.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

9.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (`application.properties|yaml`), or directly with Java code.

9.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

9.3. COMPONENT OPTIONS

The Azure Storage Queue Service component supports 15 options, which are listed below.

Name	Description	Default	Type
configuration (common)	The component configurations.		QueueConfigurati on
serviceClient (common)	Autowired Service client to a storage account to interact with the queue service. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. This client contains all the operations for interacting with a queue account in Azure Storage. Operations allowed by the client are creating, listing, and deleting queues, retrieving and updating properties of the account, and retrieving statistics of the account.		QueueServiceClie nt
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
createQueue (producer)	When is set to true, the queue will be automatically created when sending messages to the queue.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
operation (producer)	Queue service operation hint to the producer. Enum values: <ul style="list-style-type: none"> ● listQueues ● createQueue ● deleteQueue ● clearQueue ● sendMessage ● deleteMessage ● receiveMessages ● peekMessages ● updateMessage 		QueueOperationDefinition
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
maxMessages (queue)	Maximum number of messages to get, if there are less messages exist in the queue than requested all the messages will be returned. If left empty only 1 message will be retrieved, the allowed range is 1 to 32 messages.	1	Integer
messageId (queue)	The ID of the message to be deleted or updated.		String

Name	Description	Default	Type
popReceipt (queue)	Unique identifier that must match for the message to be deleted or updated.		String
timeout (queue)	An optional timeout applied to the operation. If a response is not returned before the timeout concludes a RuntimeException will be thrown.		Duration
timeToLive (queue)	How long the message will stay alive in the queue. If unset the value will default to 7 days, if -1 is passed the message will not expire. The time to live must be -1 or any positive number. The format should be in this form: PnDTnHnMn.nS., e.g: PT20.345S – parses as 20.345 seconds, P2D – parses as 2 days However, in case you are using EndpointDsl/ComponentDsl, you can do something like Duration.ofSeconds() since these Java APIs are typesafe.		Duration
visibilityTimeout (queue)	The timeout period for how long the message is invisible in the queue. The timeout must be between 1 seconds and 7 days. The format should be in this form: PnDTnHnMn.nS., e.g: PT20.345S – parses as 20.345 seconds, P2D – parses as 2 days However, in case you are using EndpointDsl/ComponentDsl, you can do something like Duration.ofSeconds() since these Java APIs are typesafe.		Duration
accessKey (security)	Access key for the associated azure account name to be used for authentication with azure queue services.		String
credentials (security)	StorageSharedKeyCredential can be injected to create the azure client, this holds the important authentication information.		StorageSharedKey Credential

9.4. ENDPOINT OPTIONS

The Azure Storage Queue Service endpoint is configured using URI syntax:

```
azure-storage-queue:accountName/queueName
```

with the following path and query parameters:

9.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
accountName (common)	Azure account name to be used for authentication with azure queue services.		String
queueName (common)	The queue resource name.		String

9.4.2. Query Parameters (31 parameters)

Name	Description	Default	Type
serviceClient (common)	Autowired Service client to a storage account to interact with the queue service. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. This client contains all the operations for interacting with a queue account in Azure Storage. Operations allowed by the client are creating, listing, and deleting queues, retrieving and updating properties of the account, and retrieving statistics of the account.		QueueServiceClient
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
createQueue (producer)	When is set to true, the queue will be automatically created when sending messages to the queue.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
operation (producer)	Queue service operation hint to the producer. Enum values: <ul style="list-style-type: none"> • listQueues • createQueue • deleteQueue • clearQueue • sendMessage • deleteMessage • receiveMessages • peekMessages • updateMessage 		QueueOperationDefinition
maxMessages (queue)	Maximum number of messages to get, if there are less messages exist in the queue than requested all the messages will be returned. If left empty only 1 message will be retrieved, the allowed range is 1 to 32 messages.	1	Integer
messageId (queue)	The ID of the message to be deleted or updated.		String
popReceipt (queue)	Unique identifier that must match for the message to be deleted or updated.		String
timeout (queue)	An optional timeout applied to the operation. If a response is not returned before the timeout concludes a RuntimeException will be thrown.		Duration
timeToLive (queue)	How long the message will stay alive in the queue. If unset the value will default to 7 days, if -1 is passed the message will not expire. The time to live must be -1 or any positive number. The format should be in this form: PnDTnHnMn.nS, e.g: PT20.345S – parses as 20.345 seconds, P2D – parses as 2 days However, in case you are using EndpointDsl/ComponentDsl, you can do something like Duration.ofSeconds() since these Java APIs are typesafe.		Duration

Name	Description	Default	Type
visibilityTimeout (queue)	The timeout period for how long the message is invisible in the queue. The timeout must be between 1 seconds and 7 days. The format should be in this form: PnDTnHnMn.nS., e.g: PT20.345S – parses as 20.345 seconds, P2D – parses as 2 days However, in case you are using EndpointDsl/ComponentDsl, you can do something like Duration.ofSeconds() since these Java APIs are typesafe.		Duration
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long

Name	Description	Default	Type
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	<p>Time unit for initialDelay and delay options.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit

Name	Description	Default	Type
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See <code>ScheduledExecutorService</code> in JDK for details.	true	boolean
accessKey (security)	Access key for the associated azure account name to be used for authentication with azure queue services.		String
credentials (security)	<code>StorageSharedKeyCredential</code> can be injected to create the azure client, this holds the important authentication information.		<code>StorageSharedKeyCredential</code>

Required information options

To use this component, you have 3 options in order to provide the required Azure authentication information:

- Provide **accountName** and **accessKey** for your Azure account, this is the simplest way to get started. The accessKey can be generated through your Azure portal.
- Provide a `StorageSharedKeyCredential` instance which can be provided into **credentials** option.
- Provide a `QueueServiceClient` instance which can be provided into **serviceClient**. Note: You don't need to create a specific client, e.g: `QueueClient`, the `QueueServiceClient` represents the upper level which can be used to retrieve lower level clients.

9.5. USAGE

For example in order to get a message content from the queue **messageQueue** in the **storageAccount** storage account and, use the following snippet:

```
from("azure-storage-queue://storageAccount/messageQueue?accessKey=yourAccessKey").
to("file://queuedirectory");
```

9.5.1. Message headers evaluated by the component producer

Header	Variable Name	Type	Operations	Description
CamelAzureStorageQueueSegmentOptions	QueueConstants.QUEUES_SEGMENT_OPTIONS	QueuesSegmentOptions	listQueues	Options for listing queues
CamelAzureStorageQueueTimeout	QueueConstants.TIMEOUT	Duration	All	An optional timeout value beyond which a <code>\{@link RuntimeException}</code> will be raised.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageQueueMetadata	QueueConstants.METADATA	Map<String,String>	createQueue	Metadata to associate with the queue
CamelAzureStorageQueueTimeToLive	QueueConstants.TIME_TO_LIVE	Duration	sendMessage	How long the message will stay alive in the queue. If unset the value will default to 7 days, if -1 is passed the message will not expire. The time to live must be -1 or any positive number.
CamelAzureStorageQueueVisibilityTimeout	QueueConstants.VISIBILITY_TIMEOUT	Duration	sendMessage, receiveMessages, updateMessage	The timeout period for how long the message is invisible in the queue. If unset the value will default to 0 and the message will be instantly visible. The timeout must be between 0 seconds and 7 days.
CamelAzureStorageQueueCreateQueue	QueueConstants.CREATE_QUEUE	boolean	sendMessage	When is set to true , the queue will be automatically created when sending messages to the queue.
CamelAzureStorageQueuePopReceipt	QueueConstants.POP_RECEIPT	String	deleteMessage, updateMessage	Unique identifier that must match for the message to be deleted or updated.
CamelAzureStorageQueueMessageId	QueueConstants.MESSAGE_ID	String	deleteMessage, updateMessage	The ID of the message to be deleted or updated.

Header	Variable Name	Type	Operations	Description
CamelAzureStorageQueueMaxMessages	QueueConstants.MAX_MESSAGES	Integer	receiveMessages, peekMessages	Maximum number of messages to get, if there are less messages exist in the queue than requested all the messages will be returned. If left empty only 1 message will be retrieved, the allowed range is 1 to 32 messages.
CamelAzureStorageQueueOperation	QueueConstants.QUEUE_OPERATION	QueueOperationDefinition	All	Specify the producer operation to execute, please see the doc on this page related to producer operation.
CamelAzureStorageQueueName	QueueConstants.QUEUE_NAME	String	All	Override the queue name.

9.5.2. Message headers set by either component producer or consumer

Header	Variable Name	Type	Description
CamelAzureStorageQueueMessageId	QueueConstants.MESSAGE_ID	String	The ID of message that being sent to the queue.
CamelAzureStorageQueueInsertionTime	QueueConstants.INSERTION_TIME	OffsetDateTime	The time the Message was inserted into the Queue.
CamelAzureStorageQueueExpirationTime	QueueConstants.EXPIRATION_TIME	OffsetDateTime	The time that the Message will expire and be automatically deleted.
CamelAzureStorageQueuePopReceipt	QueueConstants.POP_RECEIPT	String	This value is required to delete/update the Message. If deletion fails using this popreceipt then the message has been dequeued by another client.

Header	Variable Name	Type	Description
CamelAzureStorageQueueTimeNextVisible	QueueConstants.TIME_NEXT_VISIBLE	OffsetDateTime	The time that the message will again become visible in the Queue.
CamelAzureStorageQueueDequeueCount	QueueConstants.DEQUEUE_COUNT	long	The number of times the message has been dequeued.
CamelAzureStorageQueueRawHttpHeaders	QueueConstants.RAW_HTTP_HEADERS	HttpHeaders	Returns non-parsed httpHeaders that can be used by the user.

9.5.3. Advanced Azure Storage Queue configuration

If your Camel Application is running behind a firewall or if you need to have more control over the **QueueServiceClient** instance configuration, you can create your own instance:

```
StorageSharedKeyCredential credential = new StorageSharedKeyCredential("yourAccountName",
"yourAccessKey");
String uri = String.format("https://%s.queue.core.windows.net", "yourAccountName");

QueueServiceClient client = new QueueServiceClientBuilder()
    .endpoint(uri)
    .credential(credential)
    .buildClient();
// This is camel context
context.getRegistry().bind("client", client);
```

Then refer to this instance in your Camel **azure-storage-queue** component configuration:

```
from("azure-storage-queue://cameldev/queue1?serviceClient=#client")
.to("file://outputFolder?fileName=output.txt&fileExist=Append");
```

9.5.4. Automatic detection of QueueServiceClient client in registry

The component is capable of detecting the presence of an QueueServiceClient bean into the registry. If it's the only instance of that type it will be used as client and you won't have to define it as uri parameter, like the example above. This may be really useful for smarter configuration of the endpoint.

9.5.5. Azure Storage Queue Producer operations

Camel Azure Storage Queue component provides wide range of operations on the producer side:

Operations on the service level

For these operations, **accountName** is required.

Operation	Description
listQueues	Lists the queues in the storage account that pass the filter starting at the specified marker.

Operations on the queue level

For these operations, **accountName** and **queueName** are required.

Operation	Description
createQueue	Creates a new queue.
deleteQueue	Permanently deletes the queue.
clearQueue	Deletes all messages in the queue..
sendMessage	Default Producer Operation Sends a message with a given time-to-live and a timeout period where the message is invisible in the queue. The message text is evaluated from the exchange message body. By default, if the queue doesn't exist, it will create an empty queue first. If you want to disable this, set the config createQueue or header CamelAzureStorageQueueCreateQueue to false .
deleteMessage	Deletes the specified message in the queue.
receiveMessages	Retrieves up to the maximum number of messages from the queue and hides them from other operations for the timeout period. However it will not dequeue the message from the queue due to reliability reasons.
peekMessages	Peek messages from the front of the queue up to the maximum number of messages.
updateMessage	Updates the specific message in the queue with a new message and resets the visibility timeout. The message text is evaluated from the exchange message body.

Refer to the example section in this page to learn how to use these operations into your camel application.

9.5.6. Consumer Examples

To consume a queue into a file component with maximum 5 messages in one batch, this can be done like this:

```
from("azure-storage-queue://cameldev/queue1?serviceClient=#client&maxMessages=5")
.to("file://outputFolder?fileName=output.txt&fileExist=Append");
```

9.5.7. Producer Operations Examples

- **listQueues:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g, to only returns list of queues with 'awesome' prefix:
    exchange.getIn().setHeader(QueueConstants.QUEUES_SEGMENT_OPTIONS, new
QueuesSegmentOptions().setPrefix("awesome"));
  })
  .to("azure-storage-queue://cameldev?serviceClient=#client&operation=listQueues")
  .log("${body}")
  .to("mock:result");
```

- **createQueue:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(QueueConstants.QUEUE_NAME, "overrideName");
  })
  .to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=createQueue");
```

- **deleteQueue:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(QueueConstants.QUEUE_NAME, "overrideName");
  })
  .to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=deleteQueue");
```

- **clearQueue:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
    exchange.getIn().setHeader(QueueConstants.QUEUE_NAME, "overrideName");
  })
  .to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=clearQueue");
```

- **sendMessage:**

```
from("direct:start")
  .process(exchange -> {
    // set the header you want the producer to evaluate, refer to the previous
    // section to learn about the headers that can be set
    // e.g:
```

```

exchange.getIn().setBody("message to send");
// we set a visibility of 1 min
exchange.getIn().setHeader(QueueConstants.VISIBILITY_TIMEOUT, Duration.ofMinutes(1));
})
.to("azure-storage-queue://cameldev/test?serviceClient=#client");

```

- **deleteMessage:**

```

from("direct:start")
.process(exchange -> {
// set the header you want the producer to evaluate, refer to the previous
// section to learn about the headers that can be set
// e.g:
// Mandatory header:
exchange.getIn().setHeader(QueueConstants.MESSAGE_ID, "1");
// Mandatory header:
exchange.getIn().setHeader(QueueConstants.POP_RECEIPT, "PAAAAHEEERXXX-1");
})
.to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=deleteMessage");

```

- **receiveMessages:**

```

from("direct:start")
.to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=receiveMessages")
.process(exchange -> {
    final List<QueueMessageItem> messageItems = exchange.getMessage().getBody(List.class);
    messageItems.forEach(messageItem -> System.out.println(messageItem.getMessageText()));
})
.to("mock:result");

```

- **peekMessages:**

```

from("direct:start")
.to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=peekMessages")
.process(exchange -> {
    final List<PeekedMessageItem> messageItems = exchange.getMessage().getBody(List.class);
    messageItems.forEach(messageItem -> System.out.println(messageItem.getMessageText()));
})
.to("mock:result");

```

- **updateMessage:**

```

from("direct:start")
.process(exchange -> {
// set the header you want the producer to evaluate, refer to the previous
// section to learn about the headers that can be set
// e.g:
exchange.getIn().setBody("new message text");
// Mandatory header:
exchange.getIn().setHeader(QueueConstants.MESSAGE_ID, "1");
// Mandatory header:
exchange.getIn().setHeader(QueueConstants.POP_RECEIPT, "PAAAAHEEERXXX-1");
// Mandatory header:

```

```
exchange.getIn().setHeader(QueueConstants.VISIBILITY_TIMEOUT, Duration.ofMinutes(1));
})
.to("azure-storage-queue://cameldev/test?serviceClient=#client&operation=updateMessage");
```

9.5.8. Development Notes (Important)

When developing on this component, you will need to obtain your Azure accessKey in order to run the integration tests. In addition to the mocked unit tests you **will need to run the integration tests with every change you make or even client upgrade as the Azure client can break things even on minor versions upgrade**. To run the integration tests, on this component directory, run the following maven command:

```
mvn verify -PfullTests -DaccountName=myacc -DaccessKey=mykey
```

Whereby **accountName** is your Azure account name and **accessKey** is the access key being generated from Azure portal.

9.6. SPRING BOOT AUTO-CONFIGURATION

When using azure-storage-queue with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-azure-storage-queue-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 16 options, which are listed below.

Name	Description	Default	Type
camel.component.azure-storage-queue.access-key	Access key for the associated azure account name to be used for authentication with azure queue services.		String
camel.component.azure-storage-queue.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.azure-storage-queue.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.azure-storage-queue.configuration</code>	The component configurations. The option is a <code>org.apache.camel.component.azure.storage.queue.QueueConfiguration</code> type.		QueueConfiguration
<code>camel.component.azure-storage-queue.create-queue</code>	When is set to true, the queue will be automatically created when sending messages to the queue.	false	Boolean
<code>camel.component.azure-storage-queue.credentials</code>	<code>StorageSharedKeyCredential</code> can be injected to create the azure client, this holds the important authentication information. The option is a <code>com.azure.storage.common.StorageSharedKeyCredential</code> type.		StorageSharedKeyCredential
<code>camel.component.azure-storage-queue.enabled</code>	Whether to enable auto configuration of the <code>azure-storage-queue</code> component. This is enabled by default.		Boolean
<code>camel.component.azure-storage-queue.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.azure-storage-queue.max-messages</code>	Maximum number of messages to get, if there are less messages exist in the queue than requested all the messages will be returned. If left empty only 1 message will be retrieved, the allowed range is 1 to 32 messages.	1	Integer

Name	Description	Default	Type
<code>camel.component.azure-storage-queue.message-id</code>	The ID of the message to be deleted or updated.		String
<code>camel.component.azure-storage-queue.operation</code>	Queue service operation hint to the producer.		QueueOperationDefinition
<code>camel.component.azure-storage-queue.pop-receipt</code>	Unique identifier that must match for the message to be deleted or updated.		String
<code>camel.component.azure-storage-queue.service-client</code>	Service client to a storage account to interact with the queue service. This client does not hold any state about a particular storage account but is instead a convenient way of sending off appropriate requests to the resource on the service. This client contains all the operations for interacting with a queue account in Azure Storage. Operations allowed by the client are creating, listing, and deleting queues, retrieving and updating properties of the account, and retrieving statistics of the account. The option is a <code>com.azure.storage.queue.QueueServiceClient</code> type.		QueueServiceClient
<code>camel.component.azure-storage-queue.time-to-live</code>	How long the message will stay alive in the queue. If unset the value will default to 7 days, if -1 is passed the message will not expire. The time to live must be -1 or any positive number. The format should be in this form: <code>PnDTnHnMn.nS</code> , e.g: <code>PT20.345S</code> – parses as 20.345 seconds, <code>P2D</code> – parses as 2 days However, in case you are using <code>EndpointDsl/ComponentDsl</code> , you can do something like <code>Duration.ofSeconds()</code> since these Java APIs are typesafe. The option is a <code>java.time.Duration</code> type.		Duration
<code>camel.component.azure-storage-queue.timeout</code>	An optional timeout applied to the operation. If a response is not returned before the timeout concludes a <code>RuntimeException</code> will be thrown. The option is a <code>java.time.Duration</code> type.		Duration

Name	Description	Default	Type
<code>camel.component.azure-storage-queue.visibility-timeout</code>	The timeout period for how long the message is invisible in the queue. The timeout must be between 1 seconds and 7 days. The format should be in this form: PnDTnHnMn.nS, e.g: PT20.345S – parses as 20.345 seconds, P2D – parses as 2 days However, in case you are using EndpointDsl/ComponentDsl, you can do something like Duration.ofSeconds() since these Java APIs are typesafe. The option is a java.time.Duration type.		Duration

CHAPTER 10. BEAN

Only producer is supported

The Bean component binds beans to Camel message exchanges.

10.1. URI FORMAT

```
bean:beanName[?options]
```

Where **beanID** can be any string which is used to look up the bean in the Registry

10.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

10.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

10.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

10.3. COMPONENT OPTIONS

The Bean component supports 4 options, which are listed below.

Name	Description	Default	Type
cache (producer)	Deprecated Use singleton option instead.	true	Boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
scope (producer)	<p>Scope of bean. When using singleton scope (default) the bean is created or looked up only once and reused for the lifetime of the endpoint. The bean should be thread-safe in case concurrent threads is calling the bean at the same time. When using request scope the bean is created or looked up once per request (exchange). This can be used if you want to store state on a bean while processing a request and you want to call the same bean instance multiple times while processing the request. The bean does not have to be thread-safe as the instance is only called from the same request. When using delegate scope, then the bean will be looked up or created per call. However in case of lookup then this is delegated to the bean registry such as Spring or CDI (if in use), which depends on their configuration can act as either singleton or prototype scope. so when using prototype then this depends on the delegated registry.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● Singleton ● Request ● Prototype 	Singleton	BeanScope
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

10.4. ENDPOINT OPTIONS

The Bean endpoint is configured using URI syntax:

```
bean:beanName
```

with the following path and query parameters:

10.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
beanName (common)	Required Sets the name of the bean to invoke.		String

10.4.2. Query Parameters (5 parameters)

Name	Description	Default	Type
cache (common)	Deprecated Use scope option instead.		Boolean
method (common)	Sets the name of the method to invoke on the bean.		String

Name	Description	Default	Type
scope (common)	<p>Scope of bean. When using singleton scope (default) the bean is created or looked up only once and reused for the lifetime of the endpoint. The bean should be thread-safe in case concurrent threads is calling the bean at the same time. When using request scope the bean is created or looked up once per request (exchange). This can be used if you want to store state on a bean while processing a request and you want to call the same bean instance multiple times while processing the request. The bean does not have to be thread-safe as the instance is only called from the same request. When using prototype scope, then the bean will be looked up or created per call. However in case of lookup then this is delegated to the bean registry such as Spring or CDI (if in use), which depends on their configuration can act as either singleton or prototype scope. so when using prototype then this depends on the delegated registry.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • Singleton • Request • Prototype 	Singleton	BeanScope
lazyStartProducer (producer)	<p>Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.</p>	false	boolean
parameters (advanced)	<p>Used for configuring additional properties on the bean.</p>		Map

10.5. USING

The object instance that is used to consume messages must be explicitly registered with the Registry. For example, if you are using Spring you must define the bean in the Spring configuration XML file.

You can also register beans manually via Camel's **Registry** with the **bind** method.

Once an endpoint has been registered, you can build Camel routes that use it to process exchanges.

A **bean**: endpoint cannot be defined as the input to the route; i.e. you cannot consume from it, you can only route from some inbound message Endpoint to the bean endpoint as output. So consider using a **direct**: or **queue**: endpoint as the input.

You can use the **createProxy()** methods on [ProxyHelper](#) to create a proxy that will generate exchanges and send them to any endpoint:

And the same route using XML DSL:

```
<route>
  <from uri="direct:hello"/>
  <to uri="bean:bye"/>
</route>
```

10.6. BEAN AS ENDPOINT

Camel also supports invoking [Bean](#) as an Endpoint. What happens is that when the exchange is routed to the **myBean** Camel will use the Bean Binding to invoke the bean. The source for the bean is just a plain POJO.

Camel will use Bean Binding to invoke the **sayHello** method, by converting the Exchange's In body to the **String** type and storing the output of the method on the Exchange Out body.

10.7. JAVA DSL BEAN SYNTAX

Java DSL comes with syntactic sugar for the component. Instead of specifying the bean explicitly as the endpoint (i.e. **to("bean:beanName")**) you can use the following syntax:

```
// Send message to the bean endpoint
// and invoke method resolved using Bean Binding.
from("direct:start").bean("beanName");

// Send message to the bean endpoint
// and invoke given method.
from("direct:start").bean("beanName", "methodName");
```

Instead of passing name of the reference to the bean (so that Camel will lookup for it in the registry), you can specify the bean itself:

```
// Send message to the given bean instance.
from("direct:start").bean(new ExampleBean());

// Explicit selection of bean method to be invoked.
from("direct:start").bean(new ExampleBean(), "methodName");

// Camel will create the instance of bean and cache it for you.
from("direct:start").bean(ExampleBean.class);
```

10.8. BEAN BINDING

How bean methods to be invoked are chosen (if they are not specified explicitly through the **method** parameter) and how parameter values are constructed from the Message are all defined by the Bean Binding mechanism which is used throughout all of the various Bean Integration mechanisms in Camel.

10.9. SPRING BOOT AUTO-CONFIGURATION

When using bean with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-bean-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

The component supports 13 options, which are listed below.

Name	Description	Default	Type
camel.component.bean.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.bean.enabled	Whether to enable auto configuration of the bean component. This is enabled by default.		Boolean
camel.component.bean.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

Name	Description	Default	Type
camel.component.bean.scope	Scope of bean. When using singleton scope (default) the bean is created or looked up only once and reused for the lifetime of the endpoint. The bean should be thread-safe in case concurrent threads is calling the bean at the same time. When using request scope the bean is created or looked up once per request (exchange). This can be used if you want to store state on a bean while processing a request and you want to call the same bean instance multiple times while processing the request. The bean does not have to be thread-safe as the instance is only called from the same request. When using delegate scope, then the bean will be looked up or created per call. However in case of lookup then this is delegated to the bean registry such as Spring or CDI (if in use), which depends on their configuration can act as either singleton or prototype scope. so when using prototype then this depends on the delegated registry.		BeanScope
camel.component.class.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.class.enabled	Whether to enable auto configuration of the class component. This is enabled by default.		Boolean
camel.component.class.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

Name	Description	Default	Type
<code>camel.component.class.scope</code>	Scope of bean. When using singleton scope (default) the bean is created or looked up only once and reused for the lifetime of the endpoint. The bean should be thread-safe in case concurrent threads is calling the bean at the same time. When using request scope the bean is created or looked up once per request (exchange). This can be used if you want to store state on a bean while processing a request and you want to call the same bean instance multiple times while processing the request. The bean does not have to be thread-safe as the instance is only called from the same request. When using delegate scope, then the bean will be looked up or created per call. However in case of lookup then this is delegated to the bean registry such as Spring or CDI (if in use), which depends on their configuration can act as either singleton or prototype scope. so when using prototype then this depends on the delegated registry.		BeanScope
<code>camel.language.bean.enabled</code>	Whether to enable auto configuration of the bean language. This is enabled by default.		Boolean
<code>camel.language.bean.scope</code>	Scope of bean. When using singleton scope (default) the bean is created or looked up only once and reused for the lifetime of the endpoint. The bean should be thread-safe in case concurrent threads is calling the bean at the same time. When using request scope the bean is created or looked up once per request (exchange). This can be used if you want to store state on a bean while processing a request and you want to call the same bean instance multiple times while processing the request. The bean does not have to be thread-safe as the instance is only called from the same request. When using prototype scope, then the bean will be looked up or created per call. However in case of lookup then this is delegated to the bean registry such as Spring or CDI (if in use), which depends on their configuration can act as either singleton or prototype scope. So when using prototype scope then this depends on the bean registry implementation.	Singleton	String
<code>camel.language.bean.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.component.bean.cache</code>	Deprecated Use singleton option instead.	true	Boolean

Name	Description	Default	Type
<code>camel.component.class.cache</code>	Deprecated Use singleton option instead.	true	Boolean

CHAPTER 11. BEAN VALIDATOR

Only producer is supported

The Validator component performs bean validation of the message body using the Java Bean Validation API (). Camel uses the reference implementation, which is [Hibernate Validator](#).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-bean-validator</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

11.1. URI FORMAT

```
bean-validator:label[?options]
```

Where **label** is an arbitrary text value describing the endpoint. You can append query options to the URI in the following format,

?option=value&option=value&...

11.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

11.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

11.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

11.3. COMPONENT OPTIONS

The Bean Validator component supports 8 options, which are listed below.

Name	Description	Default	Type
ignoreXmlConfiguration (producer)	Whether to ignore data from the META-INF/validation.xml file.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
constraintValidatorFactory (advanced)	To use a custom ConstraintValidatorFactory.		ConstraintValidatorFactory
messageInterpolator (advanced)	To use a custom MessageInterpolator.		MessageInterpolator
traversableResolver (advanced)	To use a custom TraversableResolver.		TraversableResolver
validationProviderResolver (advanced)	To use a a custom ValidationProviderResolver.		ValidationProviderResolver

Name	Description	Default	Type
validatorFactory (advanced)	Autowired To use a custom ValidatorFactory.		ValidatorFactory

11.4. ENDPOINT OPTIONS

The Bean Validator endpoint is configured using URI syntax:

```
bean-validator:label
```

with the following path and query parameters:

11.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
label (producer)	Required Where label is an arbitrary text value describing the endpoint.		String

11.4.2. Query Parameters (8 parameters)

Name	Description	Default	Type
group (producer)	To use a custom validation group.	javax.validation.groups.Default	String
ignoreXmlConfiguration (producer)	Whether to ignore data from the META-INF/validation.xml file.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
constraintValidatorFactory (advanced)	To use a custom ConstraintValidatorFactory.		ConstraintValidatorFactory
messageInterpolator (advanced)	To use a custom MessageInterpolator.		MessageInterpolator
traversableResolver (advanced)	To use a custom TraversableResolver.		TraversableResolver
validationProviderResolver (advanced)	To use a a custom ValidationProviderResolver.		ValidationProviderResolver
validatorFactory (advanced)	To use a custom ValidatorFactory.		ValidatorFactory

11.5. OSGI DEPLOYMENT

To use Hibernate Validator in the OSGi environment use dedicated **ValidationProviderResolver** implementation, just as

org.apache.camel.component.bean.validator.HibernateValidationProviderResolver. The snippet below demonstrates this approach. You can also use **HibernateValidationProviderResolver**.

Using HibernateValidationProviderResolver

```
from("direct:test").
  to("bean-validator://ValidationProviderResolverTest?
validationProviderResolver=#myValidationProviderResolver");
```

```
<bean id="myValidationProviderResolver"
class="org.apache.camel.component.bean.validator.HibernateValidationProviderResolver"/>
```

If no custom **ValidationProviderResolver** is defined and the validator component has been deployed into the OSGi environment, the **HibernateValidationProviderResolver** will be automatically used.

11.6. EXAMPLE

Assumed we have a java bean with the following annotations

Car.java

```
public class Car {

  @NotNull
  private String manufacturer;
```



```

@NotNull
@Size(min = 5, max = 14, groups = OptionalChecks.class)
private String licensePlate;

// getter and setter
}

```

and an interface definition for our custom validation group

OptionalChecks.java

```

public interface OptionalChecks {
}

```

with the following Camel route, only the **@NotNull** constraints on the attributes manufacturer and licensePlate will be validated (Camel uses the default group **javax.validation.groups.Default**).

```

from("direct:start")
.to("bean-validator://x")
.to("mock:end")

```

If you want to check the constraints from the group **OptionalChecks**, you have to define the route like this

```

from("direct:start")
.to("bean-validator://x?group=OptionalChecks")
.to("mock:end")

```

If you want to check the constraints from both groups, you have to define a new interface first

AllChecks.java

```

@GroupSequence({Default.class, OptionalChecks.class})
public interface AllChecks {
}

```

and then your route definition should look like this

```

from("direct:start")
.to("bean-validator://x?group=AllChecks")
.to("mock:end")

```

And if you have to provide your own message interpolator, traversable resolver and constraint validator factory, you have to write a route like this

```

<bean id="myMessageInterpolator" class="my.ConstraintValidatorFactory" />
<bean id="myTraversableResolver" class="my.TraversableResolver" />
<bean id="myConstraintValidatorFactory" class="my.ConstraintValidatorFactory" />

```

```

from("direct:start")
.to("bean-validator://x?group=AllChecks&messageInterpolator=#myMessageInterpolator
&traversableResolver=#myTraversableResolver&constraintValidatorFactory=#myConstraintValidatorFac

```

```
tory")
.to("mock:end")
```

It's also possible to describe your constraints as XML and not as Java annotations. In this case, you have to provide the file **META-INF/validation.xml** which could look like this

validation.xml

```
<validation-config
  xmlns="http://jboss.org/xml/ns/javax/validation/configuration"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://jboss.org/xml/ns/javax/validation/configuration">

  <default-provider>org.hibernate.validator.HibernateValidator</default-provider>
  <message-
interpolator>org.hibernate.validator.engine.ResourceBundleMessageInterpolator</message-
interpolator>
  <traversable-
resolver>org.hibernate.validator.engine.resolver.DefaultTraversableResolver</traversable-resolver>
  <constraint-validator-
factory>org.hibernate.validator.engine.ConstraintValidatorFactoryImpl</constraint-validator-factory>
  <constraint-mapping>/constraints-car.xml</constraint-mapping>

</validation-config>
```

and the **constraints-car.xml** file

constraints-car.xml

```
<constraint-mappings xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://jboss.org/xml/ns/javax/validation/mapping validation-mapping-1.0.xsd"
  xmlns="http://jboss.org/xml/ns/javax/validation/mapping">

  <default-package>org.apache.camel.component.bean.validator</default-package>

  <bean class="CarWithoutAnnotations" ignore-annotations="true">
    <field name="manufacturer">
      <constraint annotation="javax.validation.constraints.NotNull" />
    </field>

    <field name="licensePlate">
      <constraint annotation="javax.validation.constraints.NotNull" />

      <constraint annotation="javax.validation.constraints.Size">
        <groups>
          <value>org.apache.camel.component.bean.validator.OptionalChecks</value>
        </groups>
        <element name="min">5</element>
        <element name="max">14</element>
      </constraint>
    </field>
  </bean>
</constraint-mappings>
```

Here is the XML syntax for the example route definition for [OrderedChecks](#).

Note that the body should include an instance of a class to validate.

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="direct:start"/>
      <to uri="bean-validator://x?
group=org.apache.camel.component.bean.validator.OrderedChecks"/>
    </route>
  </camelContext>
</beans>
```

11.7. SPRING BOOT AUTO-CONFIGURATION

When using bean-validator with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-bean-validator-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 9 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.bean-validator.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.bean-validator.constraint-validator-factory</code>	To use a custom ConstraintValidatorFactory. The option is a <code>javax.validation.ConstraintValidatorFactory</code> type.		ConstraintValidatorFactory
<code>camel.component.bean-validator.enabled</code>	Whether to enable auto configuration of the bean-validator component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.bean-validator.ignore-xml-configuration</code>	Whether to ignore data from the META-INF/validation.xml file.	false	Boolean
<code>camel.component.bean-validator.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.bean-validator.message-interpolator</code>	To use a custom MessageInterpolator. The option is a javax.validation.MessageInterpolator type.		MessageInterpolator
<code>camel.component.bean-validator.traversable-resolver</code>	To use a custom TraversableResolver. The option is a javax.validation.TraversableResolver type.		TraversableResolver
<code>camel.component.bean-validator.validation-provider-resolver</code>	To use a a custom ValidationProviderResolver. The option is a javax.validation.ValidationProviderResolver type.		ValidationProviderResolver
<code>camel.component.bean-validator.validator-factory</code>	To use a custom ValidatorFactory. The option is a javax.validation.ValidatorFactory type.		ValidatorFactory

CHAPTER 12. BROWSE

Both producer and consumer are supported

The Browse component provides a simple `BrowsableEndpoint` which can be useful for testing, visualisation tools or debugging. The exchanges sent to the endpoint are all available to be browsed.

12.1. URI FORMAT

```
browse:someName[?options]
```

Where `someName` can be any string to uniquely identify the endpoint.

12.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

12.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (`application.properties|yaml`), or directly with Java code.

12.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

12.3. COMPONENT OPTIONS

The Browse component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

12.4. ENDPOINT OPTIONS

The Browse endpoint is configured using URI syntax:

```
browse:name
```

with the following path and query parameters:

12.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required A name which can be any string to uniquely identify the endpoint.		String

12.4.2. Query Parameters (4 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● <code>InOnly</code> ● <code>InOut</code> ● <code>InOptionalOut</code> 		<code>ExchangePattern</code>
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

12.5. SAMPLE

In the route below, we insert a **browse:** component to be able to browse the Exchanges that are passing through:

```
from("activemq:order.in").to("browse:orderReceived").to("bean:processOrder");
```

We can now inspect the received exchanges from within the Java code:

```
private CamelContext context;
```

```

public void inspectReceivedOrders() {
    BrowsableEndpoint browse = context.getEndpoint("browse:orderReceived",
    BrowsableEndpoint.class);
    List<Exchange> exchanges = browse.getExchanges();

    // then we can inspect the list of received exchanges from Java
    for (Exchange exchange : exchanges) {
        String payload = exchange.getIn().getBody();
        // do something with payload
    }
}

```

12.6. SPRING BOOT AUTO-CONFIGURATION

When using browse with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-browse-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
camel.component.browse.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.browse.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.browse.enabled	Whether to enable auto configuration of the browse component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component .browse.lazy- start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 13. CASSANDRA CQL

Both producer and consumer are supported

[Apache Cassandra](#) is an open source NoSQL database designed to handle large amounts on commodity hardware. Like Amazon's DynamoDB, Cassandra has a peer-to-peer and master-less architecture to avoid single point of failure and garanty high availability. Like Google's BigTable, Cassandra data is structured using column families which can be accessed through the Thrift RPC API or a SQL-like API called CQL.



NOTE

This component aims at integrating Cassandra 2.0+ using the CQL3 API (not the Thrift API). It's based on [Cassandra Java Driver](#) provided by DataStax.

13.1. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

13.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

13.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

13.2. COMPONENT OPTIONS

The Cassandra CQL component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

13.3. ENDPOINT OPTIONS

The Cassandra CQL endpoint is configured using URI syntax:

```
cql:beanRef:hosts:port/keyspace
```

with the following path and query parameters:

13.3.1. Path Parameters (4 parameters)

Name	Description	Default	Type
beanRef (common)	beanRef is defined using bean:id.		String

Name	Description	Default	Type
hosts (common)	Hostname(s) Cassandra server(s). Multiple hosts can be separated by comma.		String
port (common)	Port number of Cassandra server(s).		Integer
keyspace (common)	Keyspace to use.		String

13.3.2. Query Parameters (30 parameters)

Name	Description	Default	Type
clusterName (common)	Cluster name.		String
consistencyLevel (common)	Consistency level to use. Enum values: <ul style="list-style-type: none"> ● ANY ● ONE ● TWO ● THREE ● QUORUM ● ALL ● LOCAL_ONE ● LOCAL_QUORUM ● EACH_QUORUM ● SERIAL ● LOCAL_SERIAL 		DefaultConsistencyLevel
cql (common)	CQL query to perform. Can be overridden with the message header with key CamelCqlQuery.		String
datacenter (common)	Datacenter to use.	datacenter1	String
loadBalancingPolicyClass (common)	To use a specific LoadBalancingPolicyClass.		String

Name	Description	Default	Type
password (common)	Password for session authentication.		String
prepareStatements (common)	Whether to use PreparedStatements or regular Statements.	true	boolean
resultSetConversionStrategy (common)	To use a custom class that implements logic for converting ResultSet into message body ALL, ONE, LIMIT_10, LIMIT_100...		ResultSetConversionStrategy
session (common)	To use the Session instance (you would normally not use this option).		CqlSession
username (common)	Username for session authentication.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long

Name	Description	Default	Type
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	<p>Time unit for initialDelay and delay options.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit

Name	Description	Default	Type
<code>useFixedDelay</code> (scheduler)	Controls if fixed delay or fixed rate is used. See <code>ScheduledExecutorService</code> in JDK for details.	true	boolean

13.4. ENDPOINT CONNECTION SYNTAX

The endpoint can initiate the Cassandra connection or use an existing one.

URI	Description
<code>cql:localhost/keyspace</code>	Single host, default port, usual for testing
<code>cql:host1,host2/keyspace</code>	Multi host, default port
<code>cql:host1,host2:9042/keyspace</code>	Multi host, custom port
<code>cql:host1,host2</code>	Default port and keyspace
<code>cql:bean:sessionRef</code>	Provided Session reference
<code>cql:bean:clusterRef/keyspace</code>	Provided Cluster reference

To fine tune the Cassandra connection (SSL options, pooling options, load balancing policy, retry policy, reconnection policy...), create your own Cluster instance and give it to the Camel endpoint.

13.5. MESSAGES

13.5.1. Incoming Message

The Camel Cassandra endpoint expects a bunch of simple objects (**Object** or **Object[]** or **Collection<Object>**) which will be bound to the CQL statement as query parameters. If message body is null or empty, then CQL query will be executed without binding parameters.

Headers

- **CamelCqlQuery** (optional, **String** or **RegularStatement**)
CQL query either as a plain String or built using the **QueryBuilder**.

13.5.2. Outgoing Message

The Camel Cassandra endpoint produces one or many a Cassandra Row objects depending on the **resultSetConversionStrategy**:

- **List<Row>** if **resultSetConversionStrategy** is **ALL** or **LIMIT_[0-9]+**
- **Single `Row`** if **resultSetConversionStrategy** is **ONE**

- Anything else, if **resultSetConversionStrategy** is a custom implementation of the **ResultSetConversionStrategy**

13.6. REPOSITORIES

Cassandra can be used to store message keys or messages for the idempotent and aggregation EIP.

Cassandra might not be the best tool for queuing use cases yet, read [Cassandra anti-patterns queues and queue like datasets](#). It's advised to use `LeveledCompaction` and a small GC grace setting for these tables to allow tombstoned rows to be removed quickly.

13.7. IDEMPOTENT REPOSITORY

The **NamedCassandraIdempotentRepository** stores messages keys in a Cassandra table like this:

CAMEL_IDEMPOTENT.cql

```
CREATE TABLE CAMEL_IDEMPOTENT (
  NAME varchar, -- Repository name
  KEY varchar, -- Message key
  PRIMARY KEY (NAME, KEY)
) WITH compaction = {'class': 'LeveledCompactionStrategy'}
AND gc_grace_seconds = 86400;
```

This repository implementation uses lightweight transactions (also known as Compare and Set) and requires Cassandra 2.0.7+.

Alternatively, the **CassandraIdempotentRepository** does not have a **NAME** column and can be extended to use a different data model.

Option	Default	Description
table	CAMEL_IDEMPOTENT	Table name
pkColumns	NAME, `KEY`	Primary key columns
name		Repository name, value used for NAME column
ttl		Key time to live
writeConsistencyLevel		Consistency level used to insert/delete key: ANY, ONE, TWO, QUORUM, LOCAL_QUORUM...
readConsistencyLevel		Consistency level used to read/check key: ONE, TWO, QUORUM, LOCAL_QUORUM...

13.8. AGGREGATION REPOSITORY

The **NamedCassandraAggregationRepository** stores exchanges by correlation key in a Cassandra table like this:

CAMEL_AGGREGATION.cql

```
CREATE TABLE CAMEL_AGGREGATION (
  NAME varchar,      -- Repository name
  KEY varchar,      -- Correlation id
  EXCHANGE_ID varchar, -- Exchange id
  EXCHANGE blob,    -- Serialized exchange
  PRIMARY KEY (NAME, KEY)
) WITH compaction = {'class':'LeveledCompactionStrategy'}
AND gc_grace_seconds = 86400;
```

Alternatively, the **CassandraAggregationRepository** does not have a **NAME** column and can be extended to use a different data model.

Option	Default	Description
table	CAMEL_AGGREGATION	Table name
pkColumns	NAME,KEY	Primary key columns
exchangeIdColumn	EXCHANGE_ID	Exchange Id column
exchangeColumn	EXCHANGE	Exchange content column
name		Repository name, value used for NAME column
ttl		Exchange time to live
writeConsistencyLevel		Consistency level used to insert/delete exchange: ANY, ONE, TWO, QUORUM, LOCAL_QUORUM...
readConsistencyLevel		Consistency level used to read/check exchange: ONE, TWO, QUORUM, LOCAL_QUORUM...

13.9. EXAMPLES

To insert something on a table you can use the following code:

```
String CQL = "insert into camel_user(login, first_name, last_name) values (?, ?, ?)";
from("direct:input")
.to("cql://localhost/camel_ks?cql=" + CQL);
```

At this point you should be able to insert data by using a list as body

```
Arrays.asList("davsclaus", "Claus", "Ibsen")
```

The same approach can be used for updating or querying the table.

13.10. SPRING BOOT AUTO-CONFIGURATION

When using cql with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-cassandraql-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.cql.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.cql.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.cql.enabled</code>	Whether to enable auto configuration of the cql component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.cql.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 14. CONTROL BUS

Only producer is supported

The [Control Bus](#) from the EIP patterns allows for the integration system to be monitored and managed from within the framework.

Use a Control Bus to manage an enterprise integration system. The Control Bus uses the same messaging mechanism used by the application data, but uses separate channels to transmit data that is relevant to the management of components involved in the message flow.

In Camel you can manage and monitor using JMX, or by using a Java API from the **CamelContext**, or from the **org.apache.camel.api.management** package, or use the event notifier which has an example [here](#).

The ControlBus component provides easy management of Camel applications based on the [Control Bus](#) EIP pattern. For example, by sending a message to an Endpoint you can control the lifecycle of routes, or gather performance statistics.

```
controlbus:command[?options]
```

Where **command** can be any string to identify which type of command to use.

14.1. COMMANDS

Command	Description
route	To control routes using the routeld and action parameter.
language	Allows you to specify a to use for evaluating the message body. If there is any result from the evaluation, then the result is put in the message body.

14.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

14.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

14.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

14.3. COMPONENT OPTIONS

The Control Bus component supports 2 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

14.4. ENDPOINT OPTIONS

The Control Bus endpoint is configured using URI syntax:

```
controlbus:command:language
```

with the following path and query parameters:

14.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
command (producer)	<p>Required Command can be either route or language.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● route ● language 		String
language (producer)	<p>Allows you to specify the name of a Language to use for evaluating the message body. If there is any result from the evaluation, then the result is put in the message body.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● bean ● constant ● el ● exchangeProperty ● file ● groovy ● header ● jsonpath ● mvel ● ognl ● ref ● simple ● spel ● sql ● terser ● tokenize ● xpath ● xquery ● xtokenize 		Language

14.4.1.1. Query Parameters (6 parameters)

Name	Description	Default	Type
action (producer)	<p>To denote an action that can be either: start, stop, or status. To either start or stop a route, or to get the status of the route as output in the message body. You can use suspend and resume from Camel 2.11.1 onwards to either suspend or resume a route. And from Camel 2.11.1 onwards you can use stats to get performance statics returned in XML format; the routeld option can be used to define which route to get the performance stats for, if routeld is not defined, then you get statistics for the entire CamelContext. The restart action will restart the route.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● start ● stop ● suspend ● resume ● restart ● status ● stats 		String
async (producer)	Whether to execute the control bus task asynchronously. Important: If this option is enabled, then any result from the task is not set on the Exchange. This is only possible if executing tasks synchronously.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
loggingLevel (producer)	Logging level used for logging when task is done, or if any exceptions occurred during processing the task. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	INFO	LogLevel
restartDelay (producer)	The delay in millis to use when restarting a route.	1000	int
routeId (producer)	To specify a route by its id. The special keyword <code>current</code> indicates the current route.		String

14.5. USING ROUTE COMMAND

The route command allows you to do common tasks on a given route very easily, for example to start a route, you can send an empty message to this endpoint:

```
template.sendBody("controlbus:route?routeId=foo&action=start", null);
```

To get the status of the route, you can do:

```
String status = template.requestBody("controlbus:route?routeId=foo&action=status", null, String.class);
```

14.6. GETTING PERFORMANCE STATISTICS

This requires JMX to be enabled (is by default) then you can get the performance statistics per route, or for the CamelContext. For example to get the statistics for a route named `foo`, we can do:

```
String xml = template.requestBody("controlbus:route?routeId=foo&action=stats", null, String.class);
```

The returned statistics is in XML format. Its the same data you can get from JMX with the **dumpRouteStatsAsXml** operation on the **ManagedRouteMBean**.

To get statistics for the entire CamelContext you just omit the `routeId` parameter as shown below:

```
String xml = template.requestBody("controlbus:route?action=stats", null, String.class);
```

14.7. USING SIMPLE LANGUAGE

You can use the [Simple](#) language with the control bus, for example to stop a specific route, you can send a message to the **"controlbus:language:simple"** endpoint containing the following message:

```
template.sendBody("controlbus:language:simple",
"${camelContext.getRouteController().stopRoute('myRoute')}");
```

As this is a void operation, no result is returned. However, if you want the route status you can do:


```
String status = template.requestBody("controlbus:language:simple",
"${camelContext.getRouteStatus('myRoute')}", String.class);
```

It's easier to use the **route** command to control lifecycle of routes. The **language** command allows you to execute a language script that has stronger powers such as [Groovy](#) or to some extend the [Simple](#) language.

For example to shutdown Camel itself you can do:

```
template.sendBody("controlbus:language:simple?async=true", "${camelContext.stop()}");
```

We use **async=true** to stop Camel asynchronously as otherwise we would be trying to stop Camel while it was in-flight processing the message we sent to the control bus component.



NOTE

You can also use other languages such as [Groovy](#), etc.

14.8. SPRING BOOT AUTO-CONFIGURATION

When using controlbus with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-controlbus-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
camel.component.controlbus.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.controlbus.enabled</code>	Whether to enable auto configuration of the controlbus component. This is enabled by default.		Boolean
<code>camel.component.controlbus.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 15. CRON

Only consumer is supported

The Cron component is a generic interface component that allows triggering events at specific time interval specified using the Unix cron syntax (e.g. `0/2 * * * * ?` to trigger an event every two seconds).

Being an interface component, the Cron component does not contain a default implementation, instead it requires that the users plug the implementation of their choice.

The following standard Camel components support the Cron endpoints:

- Camel-quartz
- Camel-spring

The Cron component is also supported in **Camel K**, which can use the Kubernetes scheduler to trigger the routes when required by the cron expression. Camel K does not require additional libraries to be plugged when using cron expressions compatible with Kubernetes cron syntax.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-cron</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

Additional libraries may be needed in order to plug a specific implementation.

15.1. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

15.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

15.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

15.2. COMPONENT OPTIONS

The Cron component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
cronService (advanced)	The id of the CamelCronService to use when multiple implementations are provided.		String

15.3. ENDPOINT OPTIONS

The Cron endpoint is configured using URI syntax:

```
cron:name
```

with the following path and query parameters:

15.3.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (consumer)	Required The name of the cron trigger.		String

15.3.2. Query Parameters (4 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
schedule (consumer)	Required A cron expression that will be used to generate events.		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 		ExchangePattern

15.4. USAGE

The component can be used to trigger events at specified times, as in the following example:

```
from("cron:tab?schedule=0/1+*+*+*+*+*")
  .setBody().constant("event")
  .log("${body}");
```

The schedule expression `0/3+10+*+?` can be also written as `0/3 10 * * * ?` and triggers an event every three seconds only in the tenth minute of each hour.

Parts in the schedule expression means (in order):

- Seconds (optional)
- Minutes
- Hours
- Day of month
- Month
- Day of week
- Year (optional)

Schedule expressions can be made of 5 to 7 parts. When expressions are composed of 6 parts, the first items is the "seconds" part (and year is considered missing).

Other valid examples of schedule expressions are:

- **0/2 * * * ?** (5 parts, an event every two minutes)
- **0 0/2 * * * MON-FRI 2030** (7 parts, an event every two minutes only in year 2030)

Routes can also be written using the XML DSL.

```
<route>
  <from uri="cron:tab?schedule=0/1+*+*+*+*+*" />
  <setBody>
    <constant>event</constant>
  </setBody>
  <to uri="log:info" />
</route>
```

15.5. SPRING BOOT AUTO-CONFIGURATION

When using cron with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-cron-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.cron.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.cron.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.cron.cron-service</code>	The id of the CamelCronService to use when multiple implementations are provided.		String
<code>camel.component.cron.enabled</code>	Whether to enable auto configuration of the cron component. This is enabled by default.		Boolean

CHAPTER 16. CXF

Both producer and consumer are supported

The CXF component provides integration with [Apache CXF](#) for connecting to [JAX-WS](#) services hosted in CXF.

TIP

When using CXF in streaming modes (see `DataFormat` option), then also read about Stream caching.

Maven users must add the following dependency to their `pom.xml` for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-cxf-soap</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

16.1. URI FORMAT

There are two URI formats for this endpoint: `cxfEndpoint` and `someAddress`.

```
cxf:bean:cxfEndpoint[?options]
```

Where `cxfEndpoint` represents a bean ID that references a bean in the Spring bean registry. With this URI format, most of the endpoint details are specified in the bean definition.

```
cxf://someAddress[?options]
```

Where `someAddress` specifies the CXF endpoint's address. With this URI format, most of the endpoint details are specified using options.

For either style above, you can append options to the URI as follows:

```
cxf:bean:cxfEndpoint?wsdlURL=wsdl/hello_world.wsdl&dataFormat=PAYLOAD
```

16.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

16.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

16.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a *type safe* way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

16.3. COMPONENT OPTIONS

The CXF component supports 6 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
allowStreaming (advanced)	This option controls whether the CXF component, when running in PAYLOAD mode, will DOM parse the incoming messages into DOM Elements or keep the payload as a javax.xml.transform.Source object that would allow streaming in some cases.		Boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
headerFilterStrategy (filter)	To use a custom org.apache.camel.spi.HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean

16.4. ENDPOINT OPTIONS

The CXF endpoint is configured using URI syntax:

```
cxf:beanId:address
```

with the following path and query parameters:

16.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
beanId (common)	To lookup an existing configured CxfEndpoint. Must used bean: as prefix.		String
address (service)	The service publish address.		String

16.4.2. Query Parameters (35 parameters)

Name	Description	Default	Type
dataFormat (common)	<p>The data type messages supported by the CXF endpoint.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● PAYLOAD ● RAW ● MESSAGE ● CXF_MESSAGE ● POJO 	POJO	DataFormat
wrappedStyle (common)	<p>The WSDL style that describes how parameters are represented in the SOAP body. If the value is false, CXF will chose the document-literal unwrapped style, If the value is true, CXF will chose the document-literal wrapped style.</p>		Boolean
bridgeErrorHandler (consumer)	<p>Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.</p>	false	boolean
exceptionHandler (consumer (advanced))	<p>To let the consumer use a custom <code>ExceptionHandler</code>. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.</p>		ExceptionHandler
exchangePattern (consumer (advanced))	<p>Sets the exchange pattern when the consumer creates an exchange.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
cookieHandler (producer)	<p>Configure a cookie handler to maintain a HTTP session.</p>		CookieHandler

Name	Description	Default	Type
defaultOperationName (producer)	This option will set the default operationName that will be used by the CxfProducer which invokes the remote service.		String
defaultOperationNamespace (producer)	This option will set the default operationNamespace that will be used by the CxfProducer which invokes the remote service.		String
hostnameVerifier (producer)	The hostname verifier to be used. Use the # notation to reference a HostnameVerifier from the registry.		HostnameVerifier
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
sslContextParameters (producer)	The Camel SSL setting reference. Use the # notation to reference the SSL Context.		SSLContextParameters
wrapped (producer)	Which kind of operation that CXF endpoint producer will invoke.	false	boolean
synchronous (producer (advanced))	Sets whether synchronous processing should be strictly used.	false	boolean
allowStreaming (advanced)	This option controls whether the CXF component, when running in PAYLOAD mode, will DOM parse the incoming messages into DOM Elements or keep the payload as a javax.xml.transform.Source object that would allow streaming in some cases.		Boolean
bus (advanced)	To use a custom configured CXF Bus.		Bus
continuationTimeout (advanced)	This option is used to set the CXF continuation timeout which could be used in CxfConsumer by default when the CXF server is using Jetty or Servlet transport.	30000	long

Name	Description	Default	Type
cxfBinding (advanced)	To use a custom CxfBinding to control the binding between Camel Message and CXF Message.		CxfBinding
cxfConfigurer (advanced)	This option could apply the implementation of org.apache.camel.component.cxf.CxfEndpointConfigurer which supports to configure the CXF endpoint in programmatic way. User can configure the CXF server and client by implementing configure{ServerClient} method of CxfEndpointConfigurer.		CxfConfigurer
defaultBus (advanced)	Will set the default bus when CXF endpoint create a bus by itself.	false	boolean
headerFilterStrategy (advanced)	To use a custom HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
mergeProtocolHeaders (advanced)	Whether to merge protocol headers. If enabled then propagating headers between Camel and CXF becomes more consistent and similar. For more details see CAMEL-6393.	false	boolean
mtomEnabled (advanced)	To enable MTOM (attachments). This requires to use POJO or PAYLOAD data format mode.	false	boolean
properties (advanced)	To set additional CXF options using the key/value pairs from the Map. For example to turn on stacktraces in SOAP faults, properties.faultStackTraceEnabled=true.		Map
skipPayloadMessagePartCheck (advanced)	Sets whether SOAP message validation should be disabled.	false	boolean
loggingFeatureEnabled (logging)	This option enables CXF Logging Feature which writes inbound and outbound SOAP messages to log.	false	boolean
loggingSizeLimit (logging)	To limit the total size of number of bytes the logger will output when logging feature has been enabled and -1 for no limit.	49152	int
skipFaultLogging (logging)	This option controls whether the PhaseInterceptorChain skips logging the Fault that it catches.	false	boolean
password (security)	This option is used to set the basic authentication information of password for the CXF client.		String

Name	Description	Default	Type
username (security)	This option is used to set the basic authentication information of username for the CXF client.		String
bindingId (service)	The bindingId for the service model to use.		String
portName (service)	The endpoint name this service is implementing, it maps to the wsdl:portname. In the format of ns:PORT_NAME where ns is a namespace prefix valid at this scope.		String
publishedEndpointUrl (service)	This option can override the endpointUrl that published from the WSDL which can be accessed with service address url plus wsdl.		String
serviceClass (service)	The class name of the SEI (Service Endpoint Interface) class which could have JSR181 annotation or not.		Class
serviceName (service)	The service name this service is implementing, it maps to the wsdl:servicename.		String
wsdlURL (service)	The location of the WSDL. Can be on the classpath, file system, or be hosted remotely.		String

The **serviceName** and **portName** are [QNames](#), so if you provide them be sure to prefix them with their {namespace} as shown in the examples above.

16.4.3. Descriptions of the dataformats

In Apache Camel, the Camel CXF component is the key to integrating routes with Web services. You can use the Camel CXF component to create a CXF endpoint, which can be used in either of the following ways:

- **Consumer** – (at the start of a route) represents a Web service instance, which integrates with the route. The type of payload injected into the route depends on the value of the endpoint's dataFormat option.
- **Producer** – (at other points in the route) represents a WS client proxy, which converts the current exchange object into an operation invocation on a remote Web service. The format of the current exchange must match the endpoint's dataFormat setting.

DataFormat	Description
POJO	POJOs (Plain old Java objects) are the Java parameters to the method being invoked on the target server. Both Protocol and Logical JAX-WS handlers are supported.

DataFormat	Description
PAYLOAD	PAYLOAD is the message payload (the contents of the soap:body) after message configuration in the CXF endpoint is applied. Only Protocol JAX-WS handler is supported. Logical JAX-WS handler is not supported.
RAW	RAW mode provides the raw message stream that is received from the transport layer. It is not possible to touch or change the stream, some of the CXF interceptors will be removed if you are using this kind of DataFormat, so you can't see any soap headers after the camel-cxf consumer. JAX-WS handler is not supported.
CXF_MESSAGE	CXF_MESSAGE allows for invoking the full capabilities of CXF interceptors by converting the message from the transport layer into a raw SOAP message

You can determine the data format mode of an exchange by retrieving the exchange property, **CamelCXFDataFormat**. The exchange key constant is defined in **org.apache.camel.component.cxf.common.message.CxfConstants.DATA_FORMAT_PROPERTY**.

16.4.4. How to enable CXF's LoggingOutInterceptor in RAW mode

CXF's **LoggingOutInterceptor** outputs outbound message that goes on the wire to logging system (Java Util Logging). Since the **LoggingOutInterceptor** is in **PRE_STREAM** phase (but **PRE_STREAM** phase is removed in **RAW** mode), you have to configure **LoggingOutInterceptor** to be run during the **WRITE** phase. The following is an example.

```
@Bean
public CxfEndpoint serviceEndpoint(LoggingOutInterceptor loggingOutInterceptor) {
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setAddress("http://localhost:" + port
        + "/services" + SERVICE_ADDRESS);
    cxfEndpoint.setServiceClass(org.apache.camel.component.cxf.HelloService.class);
    Map<String, Object> properties = new HashMap<String, Object>();
    properties.put("dataFormat", "RAW");
    cxfEndpoint.setProperties(properties);
    cxfEndpoint.getOutInterceptors().add(loggingOutInterceptor);
    return cxfEndpoint;
}

@Bean
public LoggingOutInterceptor loggingOutInterceptor() {
    LoggingOutInterceptor logger = new LoggingOutInterceptor("write");
    return logger;
}
```

16.4.5. Description of relayHeaders option

There are *in-band* and *out-of-band* on-the-wire headers from the perspective of a JAXWS WSDL-first developer.

The *in-band* headers are headers that are explicitly defined as part of the WSDL binding contract for an endpoint such as SOAP headers.

The *out-of-band* headers are headers that are serialized over the wire, but are not explicitly part of the WSDL binding contract.

Headers relaying/filtering is bi-directional.

When a route has a CXF endpoint and the developer needs to have on-the-wire headers, such as SOAP headers, be relayed along the route to be consumed say by another JAXWS endpoint, then **relayHeaders** should be set to **true**, which is the default value.

16.4.6. Available only in POJO mode

The **relayHeaders=true** expresses an intent to relay the headers. The actual decision on whether a given header is relayed is delegated to a pluggable instance that implements the **MessageHeadersRelay** interface. A concrete implementation of **MessageHeadersRelay** will be consulted to decide if a header needs to be relayed or not. There is already an implementation of **SoapMessageHeadersRelay** which binds itself to well-known SOAP name spaces. Currently only out-of-band headers are filtered, and in-band headers will always be relayed when **relayHeaders=true**. If there is a header on the wire whose name space is unknown to the runtime, then a fall back **DefaultMessageHeadersRelay** will be used, which simply allows all headers to be relayed.

The **relayHeaders=false** setting specifies that all headers in-band and out-of-band should be dropped.

You can plugin your own **MessageHeadersRelay** implementations overriding or adding additional ones to the list of relays. In order to override a preloaded relay instance just make sure that your **MessageHeadersRelay** implementation services the same name spaces as the one you looking to override. Also note, that the overriding relay has to service all of the name spaces as the one you looking to override, or else a runtime exception on route start up will be thrown as this would introduce an ambiguity in name spaces to relay instance mappings.

```
<cxf:cxfEndpoint ...>
  <cxf:properties>
    <entry key="org.apache.camel.cxf.message.headers.relays">
      <list>
        <ref bean="customHeadersRelay"/>
      </list>
    </entry>
  </cxf:properties>
</cxf:cxfEndpoint>
<bean id="customHeadersRelay"
class="org.apache.camel.component.cxf.soap.headers.CustomHeadersRelay"/>
```

Take a look at the tests that show how you'd be able to relay/drop headers here:

<https://github.com/apache/camel/blob/main/components/camel-cxf/src/test/java/org/apache/camel/component/cxf/soap/headers/CxfMessageHeadersRelayTest.java>

- **POJO** and **PAYLOAD** modes are supported. In **POJO** mode, only out-of-band message headers are available for filtering as the in-band headers have been processed and removed from header list by CXF. The in-band headers are incorporated into the **MessageContentList** in POJO mode. The **camel-cxf** component does not make any attempt to remove the in-band headers from the **MessageContentList**. If filtering of in-band headers is required, please use **PAYLOAD** mode or plug in a (pretty straightforward) CXF interceptor/JAXWS Handler to the CXF endpoint.

- The Message Header Relay mechanism has been merged into **CxfHeaderFilterStrategy**. The **relayHeaders** option, its semantics, and default value remain the same, but it is a property of **CxfHeaderFilterStrategy**. Here is an example of configuring it.

```
@Bean
public HeaderFilterStrategy dropAllMessageHeadersStrategy() {
    CxfHeaderFilterStrategy headerFilterStrategy = new CxfHeaderFilterStrategy();
    headerFilterStrategy.setRelayHeaders(false);
    return headerFilterStrategy;
}
```

Then, your endpoint can reference the **CxfHeaderFilterStrategy**.

```
@Bean
public CxfEndpoint routerNoRelayEndpoint(HeaderFilterStrategy dropAllMessageHeadersStrategy) {
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setServiceClass(org.apache.camel.component.cxf.soap.headers.HeaderTester.class);
    cxfEndpoint.setAddress("/CxfMessageHeadersRelayTest/HeaderService/routerNoRelayEndpoint");
    cxfEndpoint.setWsdIURL("soap_header.wsdl");
    cxfEndpoint.setEndpointNameAsQName(
        QName.valueOf("{http://apache.org/camel/component/cxf/soap/headers}SoapPortNoRelay"));
    cxfEndpoint.setServiceNameAsQName(SERVICENAME);
    Map<String, Object> properties = new HashMap<String, Object>();
    properties.put("dataFormat", "PAYLOAD");
    cxfEndpoint.setProperties(properties);
    cxfEndpoint.setHeaderFilterStrategy(dropAllMessageHeadersStrategy);
    return cxfEndpoint;
}
```

```
@Bean
public CxfEndpoint serviceNoRelayEndpoint(HeaderFilterStrategy dropAllMessageHeadersStrategy)
{
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setServiceClass(org.apache.camel.component.cxf.soap.headers.HeaderTester.class);
    cxfEndpoint.setAddress("http://localhost:" + port +
"/services/CxfMessageHeadersRelayTest/HeaderService/routerNoRelayEndpointBackend");
    cxfEndpoint.setWsdIURL("soap_header.wsdl");
    cxfEndpoint.setEndpointNameAsQName(
        QName.valueOf("{http://apache.org/camel/component/cxf/soap/headers}SoapPortNoRelay"));
    cxfEndpoint.setServiceNameAsQName(SERVICENAME);
    Map<String, Object> properties = new HashMap<String, Object>();
    properties.put("dataFormat", "PAYLOAD");
    cxfEndpoint.setProperties(properties);
    cxfEndpoint.setHeaderFilterStrategy(dropAllMessageHeadersStrategy);
    return cxfEndpoint;
}
```

Then configure the route as follows:

```
from("cxf:bean:routerNoRelayEndpoint")
.to("cxf:bean:serviceNoRelayEndpoint");
```

- The **MessageHeadersRelay** interface has changed slightly and has been renamed to **MessageHeaderFilter**. It is a property of **CxfHeaderFilterStrategy**. Here is an example of configuring user defined Message Header Filters:

```

@Bean
public HeaderFilterStrategy customMessageFilterStrategy() {
    CxfHeaderFilterStrategy headerFilterStrategy = new CxfHeaderFilterStrategy();
    List<MessageHeaderFilter> headerFilterList = new ArrayList<MessageHeaderFilter>();
    headerFilterList.add(new SoapMessageHeaderFilter());
    headerFilterList.add(new CustomHeaderFilter());
    headerFilterStrategy.setMessageHeaderFilters(headerFilterList);
    return headerFilterStrategy;
}

```

- In addition to **relayHeaders**, the following properties can be configured in **CxfHeaderFilterStrategy**.

Name	Required	Description
relayHeaders	No	All message headers will be processed by Message Header Filters <i>Type: boolean Default: true</i>
relayAllMessage Headers	No	All message headers will be propagated (without processing by Message Header Filters) <i>Type: boolean Default: false</i>
allowFilterName spaceClash	No	If two filters overlap in activation namespace, the property control how it should be handled. If the value is true , last one wins. If the value is false , it will throw an exception <i>Type: boolean Default: false</i>

16.5. CONFIGURE THE CXF ENDPOINTS WITH SPRING

You can configure the CXF endpoint with the Spring configuration file shown below, and you can also embed the endpoint into the **camelContext** tags. When you are invoking the service endpoint, you can set the **operationName** and **operationNamespace** headers to explicitly state which operation you are calling.

```

<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:cxf="http://camel.apache.org/schema/cxf"
    xsi:schemaLocation="
        http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans.xsd
        http://camel.apache.org/schema/cxf http://camel.apache.org/schema/cxf/camel-cxf.xsd
        http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-
        spring.xsd">
    <cxf:cxfEndpoint id="routerEndpoint" address="http://localhost:9003/CamelContext/RouterPort"
        serviceClass="org.apache.hello_world_soap_http.GreeterImpl"/>
    <cxf:cxfEndpoint id="serviceEndpoint" address="http://localhost:9000/SoapContext/SoapPort"
        wsdlURL="testutils/hello_world.wsdl"
        serviceClass="org.apache.hello_world_soap_http.Greeter"
        endpointName="s:SoapPort"
        serviceName="s:SOAPService"
        xmlns:s="http://apache.org/hello_world_soap_http" />
    <camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
        <route>

```

```

    <from uri="cxf:bean:routerEndpoint" />
    <to uri="cxf:bean:serviceEndpoint" />
  </route>
</camelContext>
</beans>

```

Be sure to include the JAX-WS **schemaLocation** attribute specified on the root beans element. This allows CXF to validate the file and is required. Also note the namespace declarations at the end of the **<cxf:cxfEndpoint/>** tag. These declarations are required because the combined **{namespace}localName** syntax is presently not supported for this tag's attribute values.

The **cxf:cxfEndpoint** element supports many additional attributes:


Name	Value
PortName	The endpoint name this service is implementing, it maps to the wsdl:port@name . In the format of ns:PORT_NAME where ns is a namespace prefix valid at this scope.
serviceName	The service name this service is implementing, it maps to the wsdl:service@name . In the format of ns:SERVICE_NAME where ns is a namespace prefix valid at this scope.
wsdlURL	The location of the WSDL. Can be on the classpath, file system, or be hosted remotely.
bindingId	The bindingId for the service model to use.
address	The service publish address.
bus	The bus name that will be used in the JAX-WS endpoint.
serviceClass	The class name of the SEI (Service Endpoint Interface) class which could have JSR181 annotation or not.

It also supports many child elements:

Name	Value
cxf:inInterceptors	The incoming interceptors for this endpoint. A list of <bean> or <ref> .
cxf:inFaultInterceptors	The incoming fault interceptors for this endpoint. A list of <bean> or <ref> .
cxf:outInterceptors	The outgoing interceptors for this endpoint. A list of <bean> or <ref> .
cxf:outFaultInterceptors	The outgoing fault interceptors for this endpoint. A list of <bean> or <ref> .

Name	Value
cxf:properties	A properties map which should be supplied to the JAX-WS endpoint. See below.
cxf:handlers	A JAX-WS handler list which should be supplied to the JAX-WS endpoint. See below.
cxf:dataBinding	You can specify the which DataBinding will be use in the endpoint. This can be supplied using the Spring <code><bean class="MyDataBinding"/></code> syntax.
cxf:binding	You can specify the BindingFactory for this endpoint to use. This can be supplied using the Spring <code><bean class="MyBindingFactory"/></code> syntax.
cxf:features	The features that hold the interceptors for this endpoint. A list of beans or refs
cxf:schemaLocations	The schema locations for endpoint to use. A list of schemaLocations
cxf:serviceFactory	The service factory for this endpoint to use. This can be supplied using the Spring <code><bean class="MyServiceFactory"/></code> syntax

You can find more advanced examples that show how to provide interceptors, properties and handlers on the CXF [JAX-WS Configuration page](#).



NOTE

You can use `cxf:properties` to set the camel-cxf endpoint's `dataFormat` and `setDefaultBus` properties from spring configuration file.

```
<cxf:cxfEndpoint id="testEndpoint" address="http://localhost:9000/router"
  serviceClass="org.apache.camel.component.cxf.HelloService"
  endpointName="s:PortName"
  serviceName="s:ServiceName"
  xmlns:s="http://www.example.com/test">
  <cxf:properties>
    <entry key="dataFormat" value="RAW"/>
    <entry key="setDefaultBus" value="true"/>
  </cxf:properties>
</cxf:cxfEndpoint>
```



NOTE

In SpringBoot, you can use Spring XML files to configure **camel-cxf** and use code similar to the following example to create XML configured beans:

```
@ImportResource({
    "classpath:spring-configuration.xml"
})
```

However, the use of Java code configured beans (as shown in other examples) is best practice in SpringBoot.

16.6. HOW TO MAKE THE CAMEL-CXF COMPONENT USE LOG4J INSTEAD OF JAVA.UTIL.LOGGING

CXF's default logger is **java.util.logging**. If you want to change it to log4j, proceed as follows. Create a file, in the classpath, named **META-INF/cxf/org.apache.cxf.logger**. This file should contain the fully-qualified name of the class, **org.apache.cxf.common.logging.Log4jLogger**, with no comments, on a single line.

16.7. HOW TO LET CAMEL-CXF RESPONSE START WITH XML PROCESSING INSTRUCTION

If you are using some SOAP client such as PHP, you will get this kind of error, because CXF doesn't add the XML processing instruction `<?xml version="1.0" encoding="utf-8"?>`:

```
Error:sendSms: SoapFault exception: [Client] looks like we got no XML document in [...]
```

To resolve this issue, you just need to tell `StaxOutInterceptor` to write the XML start document for you, as in the [WriteXmlDeclarationInterceptor](#) below:

```
public class WriteXmlDeclarationInterceptor extends AbstractPhaseInterceptor<SoapMessage> {
    public WriteXmlDeclarationInterceptor() {
        super(Phase.PRE_STREAM);
        addBefore(StaxOutInterceptor.class.getName());
    }

    public void handleMessage(SoapMessage message) throws Fault {
        message.put("org.apache.cxf.stax.force-start-document", Boolean.TRUE);
    }
}
```

As an alternative you can add a message header for it as demonstrated in [CxfConsumerTest](#):

```
// set up the response context which force start document
Map<String, Object> map = new HashMap<String, Object>();
map.put("org.apache.cxf.stax.force-start-document", Boolean.TRUE);
exchange.getOut().setHeader(Client.RESPONSE_CONTEXT, map);
```

16.8. HOW TO OVERRIDE THE CXF PRODUCER ADDRESS FROM MESSAGE HEADER

The **camel-cxf** producer supports to override the target service address by setting a message header **CamelDestinationOverrideUrl**.

```
// set up the service address from the message header to override the setting of CXF endpoint
exchange.getIn().setHeader(Exchange.DESTINATION_OVERRIDE_URL,
constant(getServiceAddress()));
```

16.9. HOW TO CONSUME A MESSAGE FROM A CAMEL-CXF ENDPOINT IN POJO DATA FORMAT

The **camel-cxf** endpoint consumer POJO data format is based on the [CXF invoker](#), so the message header has a property with the name of **CxfConstants.OPERATION_NAME** and the message body is a list of the SEI method parameters.

Consider the [PersonProcessor](#) example code:

```
public class PersonProcessor implements Processor {

    private static final Logger LOG = LoggerFactory.getLogger(PersonProcessor.class);

    @Override
    @SuppressWarnings("unchecked")
    public void process(Exchange exchange) throws Exception {
        LOG.info("processing exchange in camel");

        BindingOperationInfo boi = (BindingOperationInfo)
exchange.getProperty(BindingOperationInfo.class.getName());
        if (boi != null) {
            LOG.info("boi.isUnwrapped" + boi.isUnwrapped());
        }
        // Get the parameters list which element is the holder.
        MessageContentsList msgList = (MessageContentsList) exchange.getIn().getBody();
        Holder<String> personId = (Holder<String>) msgList.get(0);
        Holder<String> ssn = (Holder<String>) msgList.get(1);
        Holder<String> name = (Holder<String>) msgList.get(2);

        if (personId.value == null || personId.value.length() == 0) {
            LOG.info("person id 123, so throwing exception");
            // Try to throw out the soap fault message
            org.apache.camel.wsd1_first.types.UnknownPersonFault personFault
                = new org.apache.camel.wsd1_first.types.UnknownPersonFault();
            personFault.setPersonId("");
            org.apache.camel.wsd1_first.UnknownPersonFault fault
                = new org.apache.camel.wsd1_first.UnknownPersonFault("Get the null value of person
name", personFault);
            exchange.getMessage().setBody(fault);
            return;
        }

        name.value = "Bonjour";
        ssn.value = "123";
```

```

LOG.info("setting Bonjour as the response");
// Set the response message, first element is the return value of the operation,
// the others are the holders of method parameters
exchange.getMessage().setBody(new Object[] { null, personId, ssn, name });
}
}

```

16.10. HOW TO PREPARE THE MESSAGE FOR THE CAMEL-CXF ENDPOINT IN POJO DATA FORMAT

The **camel-cxf** endpoint producer is based on the [CXF client API](#). First you need to specify the operation name in the message header, then add the method parameters to a list, and initialize the message with this parameter list. The response message's body is a `messageContentsList`, you can get the result from that list.

If you don't specify the operation name in the message header, **CxfProducer** will try to use the **defaultOperationName** from **CxfEndpoint**, if there is no **defaultOperationName** set on **CxfEndpoint**, it will pick up the first `operationName` from the `Operation` list.

If you want to get the object array from the message body, you can get the body using **`message.getBody(Object[].class)`**, as shown in [CxfProducerRouterTest.testInvokingSimpleServerWithParams](#):

```

Exchange senderExchange = new DefaultExchange(context, ExchangePattern.InOut);
final List<String> params = new ArrayList<>();
// Prepare the request message for the camel-cxf procedure
params.add(TEST_MESSAGE);
senderExchange.getIn().setBody(params);
senderExchange.getIn().setHeader(CxfConstants.OPERATION_NAME, ECHO_OPERATION);

Exchange exchange = template.send("direct:EndpointA", senderExchange);

org.apache.camel.Message out = exchange.getMessage();
// The response message's body is an MessageContentsList which first element is the return value of
the operation,
// If there are some holder parameters, the holder parameter will be filled in the reset of List.
// The result will be extract from the MessageContentsList with the String class type
MessageContentsList result = (MessageContentsList) out.getBody();
LOG.info("Received output text: " + result.get(0));
Map<String, Object> responseContext = CastUtils.cast((Map<?, ?>)
out.getHeader(Client.RESPONSE_CONTEXT));
assertNotNull(responseContext);
assertEquals("UTF-8", responseContext.get(org.apache.cxf.message.Message.ENCODING),
    "We should get the response context here");
assertEquals("echo " + TEST_MESSAGE, result.get(0), "Reply body on Camel is wrong");

```

16.11. HOW TO DEAL WITH THE MESSAGE FOR A CAMEL-CXF ENDPOINT IN PAYLOAD DATA FORMAT

PAYLOAD means that you process the payload from the SOAP envelope as a native `CxfPayload`. **`Message.getBody()`** will return a **`org.apache.camel.component.cxf.CxfPayload`** object, with getters for SOAP message headers and the SOAP body.

See [CxfConsumerPayloadTest](#):

```
protected RouteBuilder createRouteBuilder() {
    return new RouteBuilder() {
        public void configure() {
            from(simpleEndpointURI + "&dataFormat=PAYLOAD").to("log:info").process(new Processor()
            {
                @SuppressWarnings("unchecked")
                public void process(final Exchange exchange) throws Exception {
                    CxfPayload<SoapHeader> requestPayload =
exchange.getIn().getBody(CxfPayload.class);
                    List<Source> inElements = requestPayload.getBodySources();
                    List<Source> outElements = new ArrayList<>();
                    // You can use a customer toStringConverter to turn a CxfPayLoad message into String
as you want
                    String request = exchange.getIn().getBody(String.class);
                    XmlConverter converter = new XmlConverter();
                    String documentString = ECHO_RESPONSE;

                    Element in = new XmlConverter().toDOMElement(inElements.get(0));
                    // Just check the element namespace
                    if (!in.getNamespaceURI().equals(ELEMENT_NAMESPACE)) {
                        throw new IllegalArgumentException("Wrong element namespace");
                    }
                    if (in.getLocalName().equals("echoBoolean")) {
                        documentString = ECHO_BOOLEAN_RESPONSE;
                        checkRequest("ECHO_BOOLEAN_REQUEST", request);
                    } else {
                        documentString = ECHO_RESPONSE;
                        checkRequest("ECHO_REQUEST", request);
                    }
                    Document outDocument = converter.toDOMDocument(documentString, exchange);
                    outElements.add(new DOMSource(outDocument.getDocumentElement()));
                    // set the payload header with null
                    CxfPayload<SoapHeader> responsePayload = new CxfPayload<>(null, outElements,
null);
                    exchange.getMessage().setBody(responsePayload);
                }
            });
        }
    };
}
```

16.12. HOW TO GET AND SET SOAP HEADERS IN POJO MODE

POJO means that the data format is a "list of Java objects" when the camel-cxf endpoint produces or consumes Camel exchanges. Even though Camel exposes the message body as POJOs in this mode, camel-cxf still provides access to read and write SOAP headers. However, since CXF interceptors remove in-band SOAP headers from the header list after they have been processed, only out-of-band SOAP headers are available to camel-cxf in POJO mode.

The following example illustrates how to get/set SOAP headers. Suppose we have a route that forwards from one Camel-cxf endpoint to another. That is, SOAP Client → Camel → CXF service. We can attach two processors to obtain/insert SOAP headers at (1) before a request goes out to the CXF service and

(2) before the response comes back to the SOAP Client. Processor (1) and (2) in this example are `InsertRequestOutHeaderProcessor` and `InsertResponseOutHeaderProcessor`. Our route looks like this:

```
from("cxf:bean:routerRelayEndpointWithInsertion")
    .process(new InsertRequestOutHeaderProcessor())
    .to("cxf:bean:serviceRelayEndpointWithInsertion")
    .process(new InsertResponseOutHeaderProcessor());
```

The Bean `routerRelayEndpointWithInsertion` and `serviceRelayEndpointWithInsertion` are defined as follows:

```
@Bean
public CxfEndpoint routerRelayEndpointWithInsertion() {
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setServiceClass(org.apache.camel.component.cxf.soap.headers.HeaderTester.class);

    cxfEndpoint.setAddress("/CxfMessageHeadersRelayTest/HeaderService/routerRelayEndpointWithInsertion");
    cxfEndpoint.setWsdIURL("soap_header.wsdl");
    cxfEndpoint.setEndpointNameAsQName(
        QName.valueOf("
{http://apache.org/camel/component/cxf/soap/headers}SoapPortRelayWithInsertion"));
    cxfEndpoint.setServiceNameAsQName(SERVICENAME);
    cxfEndpoint.getFeatures().add(new LoggingFeature());
    return cxfEndpoint;
}

@Bean
public CxfEndpoint serviceRelayEndpointWithInsertion() {
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setServiceClass(org.apache.camel.component.cxf.soap.headers.HeaderTester.class);
    cxfEndpoint.setAddress("http://localhost:" + port +
"/services/CxfMessageHeadersRelayTest/HeaderService/routerRelayEndpointWithInsertionBackend");

    cxfEndpoint.setWsdIURL("soap_header.wsdl");
    cxfEndpoint.setEndpointNameAsQName(
        QName.valueOf("
{http://apache.org/camel/component/cxf/soap/headers}SoapPortRelayWithInsertion"));
    cxfEndpoint.setServiceNameAsQName(SERVICENAME);
    cxfEndpoint.getFeatures().add(new LoggingFeature());
    return cxfEndpoint;
}
```

SOAP headers are propagated to and from Camel Message headers. The Camel message header name is `org.apache.cxf.headers.Header.list` which is a constant defined in CXF (`org.apache.cxf.headers.Header.HEADER_LIST`). The header value is a List of CXF `SoapHeader` objects (`org.apache.cxf.binding.soap.SoapHeader`). The following snippet is the `InsertResponseOutHeaderProcessor` (that insert a new SOAP header in the response message). The way to access SOAP headers in both `InsertResponseOutHeaderProcessor` and `InsertRequestOutHeaderProcessor` are actually the same. The only difference between the two processors is setting the direction of the inserted SOAP header.

You can find the `InsertResponseOutHeaderProcessor` example in [CxfMessageHeadersRelayTest](#):

```
public static class InsertResponseOutHeaderProcessor implements Processor {
```

```

public void process(Exchange exchange) throws Exception {
    List<SoapHeader> soapHeaders = CastUtils.cast((List<?
>)exchange.getIn().getHeader(Header.HEADER_LIST));

    // Insert a new header
    String xml = "<?xml version=\"1.0\" encoding=\"utf-8\"?><outofbandHeader "
        + "xmlns=\"http://cxf.apache.org/outofband/Header\" hdrAttribute=\"testHdrAttribute\" "
        + "xmlns:soap=\"http://schemas.xmlsoap.org/soap/envelope/\" soap:mustUnderstand=\"1\">"
        + "<name>New_testOobHeader</name><value>New_testOobHeaderValue</value>"
    </outofbandHeader>";
    SoapHeader newHeader = new SoapHeader(soapHeaders.get(0).getName(),
        DOMUtils.readXml(new StringReader(xml)).getDocumentElement());
    // make sure direction is OUT since it is a response message.
    newHeader.setDirection(Direction.DIRECTION_OUT);
    //newHeader.setMustUnderstand(false);
    soapHeaders.add(newHeader);
}
}

```

16.13. HOW TO GET AND SET SOAP HEADERS IN PAYLOAD MODE

We've already shown how to access the SOAP message as `CxfPayload` object in PAYLOAD mode in the section [How to deal with the message for a camel-cxf endpoint in PAYLOAD data format](#) .

Once you obtain a `CxfPayload` object, you can invoke the `CxfPayload.getHeaders()` method that returns a List of DOM Elements (SOAP headers).

For an example see [CxfPayloadSoapHeaderTest](#):

```

from(getRouterEndpointURI()).process(new Processor() {
    @SuppressWarnings("unchecked")
    public void process(Exchange exchange) throws Exception {
        CxfPayload<SoapHeader> payload = exchange.getIn().getBody(CxfPayload.class);
        List<Source> elements = payload.getBodySources();
        assertNotNull(elements, "We should get the elements here");
        assertEquals(1, elements.size(), "Get the wrong elements size");

        Element el = new XmlConverter().toDOMElement(elements.get(0));
        elements.set(0, new DOMSource(el));
        assertEquals("http://camel.apache.org/pizza/types",
            el.getNamespaceURI(), "Get the wrong namespace URI");

        List<SoapHeader> headers = payload.getHeaders();
        assertNotNull(headers, "We should get the headers here");
        assertEquals(1, headers.size(), "Get the wrong headers size");
        assertEquals("http://camel.apache.org/pizza/types",
            ((Element) (headers.get(0).getObject())).getNamespaceURI(), "Get the wrong namespace
URI");
        // alternatively you can also get the SOAP header via the camel header:
        headers = exchange.getIn().getHeader(Header.HEADER_LIST, List.class);
        assertNotNull(headers, "We should get the headers here");
        assertEquals(1, headers.size(), "Get the wrong headers size");
    }
}

```

```

    assertEquals("http://camel.apache.org/pizza/types",
        ((Element) (headers.get(0).getObject())).getNamespaceURI(), "Get the wrong namespace
URI");
    }
})
.to(getServiceEndpointURI());

```

You can also use the same way as described in sub-chapter "How to get and set SOAP headers in POJO mode" to set or get the SOAP headers. So, you can use the header "org.apache.cxf.headers.Header.list" to get and set a list of SOAP headers. This does also mean that if you have a route that forwards from one Camel-cxf endpoint to another (SOAP Client → Camel → CXF service), now also the SOAP headers sent by the SOAP client are forwarded to the CXF service. If you do not want that these headers are forwarded you have to remove them in the Camel header "org.apache.cxf.headers.Header.list".

16.14. SOAP HEADERS ARE NOT AVAILABLE IN RAW MODE

SOAP headers are not available in RAW mode as SOAP processing is skipped.

16.15. HOW TO THROW A SOAP FAULT FROM CAMEL

If you are using a **camel-cxf** endpoint to consume the SOAP request, you may need to throw the SOAP Fault from the camel context.

Basically, you can use the **throwFault** DSL to do that; it works for **POJO**, **PAYLOAD** and **MESSAGE** data format.

You can define the soap fault as shown in [CxfCustomizedExceptionTest](#):

```

SOAP_FAULT = new SoapFault(EXCEPTION_MESSAGE, SoapFault.FAULT_CODE_CLIENT);
Element detail = SOAP_FAULT.getOrCreateDetail();
Document doc = detail.getOwnerDocument();
Text tn = doc.createTextNode(DETAIL_TEXT);
detail.appendChild(tn);

```

Then throw it as you like

```

from(routerEndpointURI).setFaultBody(constant(SOAP_FAULT));

```

If your CXF endpoint is working in the **MESSAGE** data format, you could set the SOAP Fault message in the message body and set the response code in the message header as demonstrated by [CxfMessageStreamExceptionTest](#)

```

from(routerEndpointURI).process(new Processor() {
    public void process(Exchange exchange) throws Exception {
        Message out = exchange.getOut();
        // Set the message body with the
        out.setBody(this.getClass().getResourceAsStream("SoapFaultMessage.xml"));
        // Set the response code here
        out.setHeader(org.apache.cxf.message.Message.RESPONSE_CODE, new Integer(500));
    }
});

```

Same for using POJO data format. You can set the SOAPFault on the out body.

16.16. HOW TO PROPAGATE A CAMEL-CXF ENDPOINT'S REQUEST AND RESPONSE CONTEXT

[CXF client API](#) provides a way to invoke the operation with request and response context. If you are using a **camel-cxf** endpoint producer to invoke the outside web service, you can set the request context and get response context with the following code:

```

    CxfExchange exchange = (CxfExchange)template.send(getJaxwsEndpointUri(), new
Processor() {
    public void process(final Exchange exchange) {
        final List<String> params = new ArrayList<String>();
        params.add(TEST_MESSAGE);
        // Set the request context to the inMessage
        Map<String, Object> requestContext = new HashMap<String, Object>();
        requestContext.put(BindingProvider.ENDPOINT_ADDRESS_PROPERTY,
JAXWS_SERVER_ADDRESS);
        exchange.getIn().setBody(params);
        exchange.getIn().setHeader(Client.REQUEST_CONTEXT , requestContext);
        exchange.getIn().setHeader(CxfConstants.OPERATION_NAME,
GREET_ME_OPERATION);
    }
});
org.apache.camel.Message out = exchange.getOut();
// The output is an object array, the first element of the array is the return value
Object[] output = out.getBody(Object[].class);
LOG.info("Received output text: " + output[0]);
// Get the response context form outMessage
Map<String, Object> responseContext =
CastUtils.cast((Map)out.getHeader(Client.RESPONSE_CONTEXT));
assertNotNull(responseContext);
assertEquals("Get the wrong wsdl operation name", "
{http://apache.org/hello_world_soap_http}greetMe",
responseContext.get("javax.xml.ws.wsdl.operation").toString());

```

16.17. ATTACHMENT SUPPORT

POJO Mode: Both SOAP with Attachment and MTOM are supported (see example in Payload Mode for enabling MTOM). However, SOAP with Attachment is not tested. Since attachments are marshalled and unmarshalled into POJOs, users typically do not need to deal with the attachment themselves. Attachments are propagated to Camel message's attachments if the MTOM is not enabled. So, it is possible to retrieve attachments by Camel Message API

```
DataHandler Message.getAttachment(String id)
```

Payload Mode: MTOM is supported by the component. Attachments can be retrieved by Camel Message APIs mentioned above. SOAP with Attachment (SwA) is supported and attachments can be retrieved. SwA is the default (same as setting the CXF endpoint property "mtom-enabled" to false).

To enable MTOM, set the CXF endpoint property "mtom-enabled" to *true*.

```
@Bean
```

```

public CxfEndpoint routerEndpoint() {
    CxfSpringEndpoint cxfEndpoint = new CxfSpringEndpoint();
    cxfEndpoint.setServiceNameAsQName(SERVICE_QNAME);
    cxfEndpoint.setEndpointNameAsQName(PORT_QNAME);
    cxfEndpoint.setAddress("/") + getClass().getSimpleName() + "/jaxws-mtom/hello");
    cxfEndpoint.setWsdIURL("mtom.wsdl");
    Map<String, Object> properties = new HashMap<String, Object>();
    properties.put("dataFormat", "PAYLOAD");
    properties.put("mtom-enabled", true);
    cxfEndpoint.setProperties(properties);
    return cxfEndpoint;
}

```

You can produce a Camel message with attachment to send to a CXF endpoint in Payload mode.

```

Exchange exchange = context.createProducerTemplate().send("direct:testEndpoint", new
Processor() {

    public void process(Exchange exchange) throws Exception {
        exchange.setPattern(ExchangePattern.InOut);
        List<Source> elements = new ArrayList<Source>();
        elements.add(new DOMSource(DOMUtils.readXml(new
StringReader(MtomTestHelper.REQ_MESSAGE)).getDocumentElement()));
        CxfPayload<SoapHeader> body = new CxfPayload<SoapHeader>(new ArrayList<SoapHeader>
(),
        elements, null);
        exchange.getIn().setBody(body);
        exchange.getIn().addAttachment(MtomTestHelper.REQ_PHOTO_CID,
        new DataHandler(new ByteArrayDataSource(MtomTestHelper.REQ_PHOTO_DATA,
"application/octet-stream")));

        exchange.getIn().addAttachment(MtomTestHelper.REQ_IMAGE_CID,
        new DataHandler(new ByteArrayDataSource(MtomTestHelper.requestJpeg, "image/jpeg")));

    }

});

// process response

CxfPayload<SoapHeader> out = exchange.getOut().getBody(CxfPayload.class);
Assert.assertEquals(1, out.getBody().size());

Map<String, String> ns = new HashMap<String, String>();
ns.put("ns", MtomTestHelper.SERVICE_TYPES_NS);
ns.put("xop", MtomTestHelper.XOP_NS);

XPathUtils xu = new XPathUtils(ns);
Element oute = new XmlConverter().toDOMElement(out.getBody().get(0));
Element ele = (Element)xu.getValue("//ns:DetailResponse/ns:photo/xop:Include", oute,
        XPathConstants.NODE);
String photold = ele.getAttribute("href").substring(4); // skip "cid:"

ele = (Element)xu.getValue("//ns:DetailResponse/ns:image/xop:Include", oute,
        XPathConstants.NODE);
String imageld = ele.getAttribute("href").substring(4); // skip "cid:"

```

```

DataHandler dr = exchange.getOut().getAttachment(photoid);
Assert.assertEquals("application/octet-stream", dr.getContentType());
MtomTestHelper.assertEquals(MtomTestHelper.RESP_PHOTO_DATA,
IOUtils.readBytesFromStream(dr.getInputStream()));

dr = exchange.getOut().getAttachment(imageld);
Assert.assertEquals("image/jpeg", dr.getContentType());

BufferedImage image = ImageIO.read(dr.getInputStream());
Assert.assertEquals(560, image.getWidth());
Assert.assertEquals(300, image.getHeight());

```

You can also consume a Camel message received from a CXF endpoint in Payload mode. The [CxfMtomConsumerPayloadModeTest](#) illustrates how this works:

```

public static class MyProcessor implements Processor {

    @SuppressWarnings("unchecked")
    public void process(Exchange exchange) throws Exception {
        CxfPayload<SoapHeader> in = exchange.getIn().getBody(CxfPayload.class);

        // verify request
        Assert.assertEquals(1, in.getBody().size());

        Map<String, String> ns = new HashMap<String, String>();
        ns.put("ns", MtomTestHelper.SERVICE_TYPES_NS);
        ns.put("xop", MtomTestHelper.XOP_NS);

        XPathUtils xu = new XPathUtils(ns);
        Element body = new XmlConverter().toDOMElement(in.getBody().get(0));
        Element ele = (Element)xu.getValue("//ns:Detail/ns:photo/xop:Include", body,
            XPathConstants.NODE);
        String photoid = ele.getAttribute("href").substring(4); // skip "cid:"
        Assert.assertEquals(MtomTestHelper.REQ_PHOTO_CID, photoid);

        ele = (Element)xu.getValue("//ns:Detail/ns:image/xop:Include", body,
            XPathConstants.NODE);
        String imageld = ele.getAttribute("href").substring(4); // skip "cid:"
        Assert.assertEquals(MtomTestHelper.REQ_IMAGE_CID, imageld);

        DataHandler dr = exchange.getIn().getAttachment(photoid);
        Assert.assertEquals("application/octet-stream", dr.getContentType());
        MtomTestHelper.assertEquals(MtomTestHelper.REQ_PHOTO_DATA,
IOUtils.readBytesFromStream(dr.getInputStream()));

        dr = exchange.getIn().getAttachment(imageld);
        Assert.assertEquals("image/jpeg", dr.getContentType());
        MtomTestHelper.assertEquals(MtomTestHelper.requestJpeg,
IOUtils.readBytesFromStream(dr.getInputStream()));

        // create response
        List<Source> elements = new ArrayList<Source>();
        elements.add(new DOMSource(DOMUtils.readXml(new
StringReader(MtomTestHelper.RESP_MESSAGE)).getDocumentElement()));
        CxfPayload<SoapHeader> sbody = new CxfPayload<SoapHeader>(new

```

```

ArrayList<SoapHeader>(),
    elements, null);
exchange.getOut().setBody(sbody);
exchange.getOut().addAttachment(MtomTestHelper.RESP_PHOTO_CID,
    new DataHandler(new ByteArrayDataSource(MtomTestHelper.RESP_PHOTO_DATA,
"application/octet-stream")));

exchange.getOut().addAttachment(MtomTestHelper.RESP_IMAGE_CID,
    new DataHandler(new ByteArrayDataSource(MtomTestHelper.responseJpeg, "image/jpeg")));
    }
}

```

Raw Mode: Attachments are not supported as it does not process the message at all.

CXF_RAW Mode: MTOM is supported, and Attachments can be retrieved by Camel Message APIs mentioned above. Note that when receiving a multipart (i.e. MTOM) message the default SOAPMessage to String converter will provide the complete multipart payload on the body. If you require just the SOAP XML as a String, you can set the message body with `message.getSOAPPart()`, and Camel convert can do the rest of work for you.

16.18. STREAMING SUPPORT IN PAYLOAD MODE

The camel-cxf component now supports streaming of incoming messages when using PAYLOAD mode. Previously, the incoming messages would have been completely DOM parsed. For large messages, this is time consuming and uses a significant amount of memory. The incoming messages can remain as a `javax.xml.transform.Source` while being routed and, if nothing modifies the payload, can then be directly streamed out to the target destination. For common "simple proxy" use cases (example: `from("cxf:...").to("cxf:...")`), this can provide very significant performance increases as well as significantly lowered memory requirements.

However, there are cases where streaming may not be appropriate or desired. Due to the streaming nature, invalid incoming XML may not be caught until later in the processing chain. Also, certain actions may require the message to be DOM parsed anyway (like WS-Security or message tracing and such) in which case the advantages of the streaming is limited. At this point, there are two ways to control the streaming:

- Endpoint property: you can add "allowStreaming=false" as an endpoint property to turn the streaming on/off.
- Component property: the `CxfComponent` object also has an `allowStreaming` property that can set the default for endpoints created from that component.

Global system property: you can add a system property of "org.apache.camel.component.cxf.streaming" to "false" to turn it off. That sets the global default, but setting the endpoint property above will override this value for that endpoint.

16.19. USING THE GENERIC CXF DISPATCH MODE

The camel-cxf component supports the generic [CXF dispatch mode](#) that can transport messages of arbitrary structures (i.e., not bound to a specific XML schema). To use this mode, you simply omit specifying the `wSDLURL` and `serviceClass` attributes of the CXF endpoint.

```

<cxf:cxfEndpoint id="testEndpoint" address="http://localhost:9000/SoapContext/SoapAnyPort">
  <cxf:properties>

```



```

<entry key="dataFormat" value="PAYLOAD"/>
</cxf:properties>
</cxf:cxfEndpoint>

```

It is noted that the default CXF dispatch client does not send a specific SOAPAction header. Therefore, when the target service requires a specific SOAPAction value, it is supplied in the Camel header using the key SOAPAction (case-insensitive).

16.20. SPRING BOOT AUTO-CONFIGURATION

When using cxf with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-cxf-soap-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 13 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.cxf.allow-streaming</code>	This option controls whether the CXF component, when running in PAYLOAD mode, will DOM parse the incoming messages into DOM Elements or keep the payload as a <code>javax.xml.transform.Source</code> object that would allow streaming in some cases.		Boolean
<code>camel.component.cxf.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.cxf.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.cxf.enabled</code>	Whether to enable auto configuration of the cxf component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.cxf.header-filter-strategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		<code>HeaderFilterStrategy</code>
<code>camel.component.cxf.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	<code>false</code>	<code>Boolean</code>
<code>camel.component.cxf.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	<code>false</code>	<code>Boolean</code>
<code>camel.component.cxf.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	<code>true</code>	<code>Boolean</code>
<code>camel.component.cxf.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at <code>WARN</code> or <code>ERROR</code> level and ignored.	<code>false</code>	<code>Boolean</code>
<code>camel.component.cxf.enabled</code>	Whether to enable auto configuration of the <code>cxf</code> component. This is enabled by default.		<code>Boolean</code>
<code>camel.component.cxf.header-filter-strategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		<code>HeaderFilterStrategy</code>

Name	Description	Default	Type
<code>camel.component.cxfrcs.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.cxfrcs.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	false	Boolean

CHAPTER 17. DATA FORMAT

Only producer is supported

The Dataformat component allows to use the [Data Format](#) as a Camel Component.

17.1. URI FORMAT

```
dataformat:name:(marshal|unmarshal)[?options]
```

Where **name** is the name of the Data Format. And then followed by the operation which must either be **marshal** or **unmarshal**. The options is used for configuring the [Data Format](#) in use. See the Data Format documentation for which options it support.

17.2. DATAFORMAT OPTIONS

17.2.1. Configuring Options

Camel components are configured on two separate levels:

- component level
- endpoint level

17.2.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

17.2.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

17.3. COMPONENT OPTIONS

The Data Format component supports 2 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

17.4. ENDPOINT OPTIONS

The Data Format endpoint is configured using URI syntax:

```
dataformat:name:operation
```

with the following path and query parameters:

17.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
name (producer)	Required Name of data format.		String
operation (producer)	Required Operation to use either marshal or unmarshal. Enum values: <ul style="list-style-type: none"> • marshal • unmarshal 		String

17.4.2. Query Parameters (1 parameters)

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

17.5. SAMPLES

For example to use the [JAXB Data Format](#) we can do as follows:

```
from("activemq:My.Queue").
  to("dataformat:jaxb:unmarshal?contextPath=com.acme.model").
  to("mqseries:Another.Queue");
```

And in XML DSL you do:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="activemq:My.Queue"/>
    <to uri="dataformat:jaxb:unmarshal?contextPath=com.acme.model"/>
    <to uri="mqseries:Another.Queue"/>
  </route>
</camelContext>
```

17.6. SPRING BOOT AUTO-CONFIGURATION

When using dataformat with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-dataformat-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
camel.component.dataformat.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.dataformat.enabled	Whether to enable auto configuration of the dataformat component. This is enabled by default.		Boolean
camel.component.dataformat.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 18. DATASET

Both producer and consumer are supported

Testing of distributed and asynchronous processing is notoriously difficult. The [Mock](#), [Test](#) and [DataSet](#) endpoints work great with the Camel Testing Framework to simplify your unit and integration testing using [Enterprise Integration Patterns](#) and Camel's large range of Components together with the powerful Bean Integration.

The DataSet component provides a mechanism to easily perform load & soak testing of your system. It works by allowing you to create [DataSet instances](#) both as a source of messages and as a way to assert that the data set is received.

Camel will use the [throughput logger](#) when sending datasets.

18.1. URI FORMAT

```
dataset:name[?options]
```

Where **name** is used to find the [DataSet instance](#) in the Registry

Camel ships with a support implementation of **org.apache.camel.component.dataset.DataSet**, the **org.apache.camel.component.dataset.DataSetSupport** class, that can be used as a base for implementing your own DataSet. Camel also ships with some implementations that can be used for testing: **org.apache.camel.component.dataset.SimpleDataSet**, **org.apache.camel.component.dataset.ListDataSet** and **org.apache.camel.component.dataset.FileDataSet**, all of which extend **DataSetSupport**.

18.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

18.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

18.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

18.3. COMPONENT OPTIONS

The Dataset component supports 5 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
log (producer)	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
exchangeFormatter (advanced)	Autowired Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to DefaultExchangeFormatter.		ExchangeFormatter

18.4. ENDPOINT OPTIONS

The Dataset endpoint is configured using URI syntax:

```
dataset:name
```

with the following path and query parameters:

18.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required Name of DataSet to lookup in the registry.		DataSet

18.4.2. Query Parameters (21 parameters)

Name	Description	Default	Type
dataSetIndex (common)	<p>Controls the behaviour of the CamelDataSetIndex header. For Consumers: - off = the header will not be set - strict/lenient = the header will be set For Producers: - off = the header value will not be verified, and will not be set if it is not present = strict = the header value must be present and will be verified = lenient = the header value will be verified if it is present, and will be set if it is not present.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • strict • lenient • off 	lenient	String

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
initialDelay (consumer)	Time period in millis to wait before starting sending messages.	1000	long
minRate (consumer)	Wait until the DataSet contains at least this number of messages.	0	int
preloadSize (consumer)	Sets how many messages should be preloaded (sent) before the route completes its initialization.	0	long
produceDelay (consumer)	Allows a delay to be specified which causes a delay when a message is sent by the consumer (to simulate slow processing).	3	long
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
assertPeriod (producer)	Sets a grace period after which the mock endpoint will re-assert to ensure the preliminary assertion is still valid. This is used for example to assert that exactly a number of messages arrives. For example if <code>expectedMessageCount(int)</code> was set to 5, then the assertion is satisfied when 5 or more message arrives. To ensure that exactly 5 messages arrives, then you would need to wait a little period to ensure no further message arrives. This is what you can use this method for. By default this period is disabled.		long
consumeDelay (producer)	Allows a delay to be specified which causes a delay when a message is consumed by the producer (to simulate slow processing).	0	long
expectedCount (producer)	Specifies the expected number of message exchanges that should be received by this endpoint. Beware: If you want to expect that 0 messages, then take extra care, as 0 matches when the tests starts, so you need to set a assert period time to let the test run for a while to make sure there are still no messages arrived; for that use <code>setAssertPeriod(long)</code> . An alternative is to use <code>NotifyBuilder</code> , and use the notifier to know when Camel is done routing some messages, before you call the <code>assertIsSatisfied()</code> method on the mocks. This allows you to not use a fixed assert period, to speedup testing times. If you want to assert that exactly n'th message arrives to this mock endpoint, then see also the <code>setAssertPeriod(long)</code> method for further details.	-1	int
failFast (producer)	Sets whether <code>assertIsSatisfied()</code> should fail fast at the first detected failed expectation while it may otherwise wait for all expected messages to arrive before performing expectations verifications. Is by default true. Set to false to use behavior as in Camel 2.x.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
log (producer)	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the org.apache.camel.component.mock.MockEndpoint class.	false	boolean
reportGroup (producer)	A number that is used to turn on throughput logging based on groups of the size.		int
resultMinimumWaitTime (producer)	Sets the minimum expected amount of time (in millis) the assertIsSatisfied() will wait on a latch until it is satisfied.		long
resultWaitTime (producer)	Sets the maximum amount of time (in millis) the assertIsSatisfied() will wait on a latch until it is satisfied.		long
retainFirst (producer)	Specifies to only retain the first n'th number of received Exchanges. This is used when testing with big data, to reduce memory consumption by not storing copies of every Exchange this mock endpoint receives. Important: When using this limitation, then the getReceivedCounter() will still return the actual number of received Exchanges. For example if we have received 5000 Exchanges, and have configured to only retain the first 10 Exchanges, then the getReceivedCounter() will still return 5000 but there is only the first 10 Exchanges in the getExchanges() and getReceivedExchanges() methods. When using this method, then some of the other expectation methods is not supported, for example the expectedBodiesReceived(Object...) sets a expectation on the first number of bodies received. You can configure both setRetainFirst(int) and setRetainLast(int) methods, to limit both the first and last received.	-1	int

Name	Description	Default	Type
retainLast (producer)	Specifies to only retain the last n'th number of received Exchanges. This is used when testing with big data, to reduce memory consumption by not storing copies of every Exchange this mock endpoint receives. Important: When using this limitation, then the <code>getReceivedCounter()</code> will still return the actual number of received Exchanges. For example if we have received 5000 Exchanges, and have configured to only retain the last 20 Exchanges, then the <code>getReceivedCounter()</code> will still return 5000 but there is only the last 20 Exchanges in the <code>getExchanges()</code> and <code>getReceivedExchanges()</code> methods. When using this method, then some of the other expectation methods is not supported, for example the <code>expectedBodiesReceived(Object...)</code> sets a expectation on the first number of bodies received. You can configure both <code>setRetainFirst(int)</code> and <code>setRetainLast(int)</code> methods, to limit both the first and last received.	-1	int
sleepForEmptyTest (producer)	Allows a sleep to be specified to wait to check that this endpoint really is empty when <code>expectedMessageCount(int)</code> is called with zero.		long
copyOnExchange (producer (advanced))	Sets whether to make a deep copy of the incoming Exchange when received at this mock endpoint. Is by default true.	true	boolean

18.5. CONFIGURING DATASET

Camel will lookup in the Registry for a bean implementing the `DataSet` interface. So you can register your own `DataSet` as:

```
<bean id="myDataSet" class="com.mycompany.MyDataSet">
  <property name="size" value="100"/>
</bean>
```

18.6. EXAMPLE

For example, to test that a set of messages are sent to a queue and then consumed from the queue without losing any messages:

```
// send the dataset to a queue
from("dataset:foo").to("activemq:SomeQueue");

// now lets test that the messages are consumed correctly
from("activemq:SomeQueue").to("dataset:foo");
```

The above would look in the Registry to find the **foo** DataSet instance which is used to create the messages.

Then you create a DataSet implementation, such as using the **SimpleDataSet** as described below, configuring things like how big the data set is and what the messages look like etc.

18.7. DATASETSUPPORT (ABSTRACT CLASS)

The DataSetSupport abstract class is a nice starting point for new DataSets, and provides some useful features to derived classes.

18.7.1. Properties on DataSetSupport

Property	Type	Default	Description
defaultHeaders	Map<String, Object>	null	Specifies the default message body. For SimpleDataSet it is a constant payload; though if you want to create custom payloads per message, create your own derivation of DataSetSupport .
outputTransformer	org.apache.camel.Processor	null	
size	long	10	Specifies how many messages to send/consume.
reportCount	long	-1	Specifies the number of messages to be received before reporting progress. Useful for showing progress of a large load test. If < 0, then size / 5, if is 0 then size , else set to reportCount value.

18.8. SIMPLEDATASET

The **SimpleDataSet** extends **DataSetSupport**, and adds a default body.

18.8.1. Additional Properties on SimpleDataSet

Property	Type	Default	Description
defaultBody	Object	<hello>world! </hello>	Specifies the default message body. By default, the SimpleDataSet produces the same constant payload for each exchange. If you want to customize the payload for each exchange, create a Camel Processor and configure the SimpleDataSet to use it by setting the outputTransformer property.

18.9. LISTDATASET

The `ListDataSet` extends **DataSetSupport**, and adds a list of default bodies.

18.9.1. Additional Properties on ListDataSet

Property	Type	Default	Description
defaultBodies	List<Object>	empty LinkedList<Object>	Specifies the default message body. By default, the ListDataSet selects a constant payload from the list of defaultBodies using the CamelDataSetIndex . If you want to customize the payload, create a Camel Processor and configure the ListDataSet to use it by setting the outputTransformer property.
size	long	the size of the defaultBodies list	Specifies how many messages to send/consume. This value can be different from the size of the defaultBodies list. If the value is less than the size of the defaultBodies list, some of the list elements will not be used. If the value is greater than the size of the defaultBodies list, the payload for the exchange will be selected using the modulus of the CamelDataSetIndex and the size of the defaultBodies list (i.e. CamelDataSetIndex % defaultBodies.size())

18.10. FILEDATASET

The **FileDataSet** extends **ListDataSet**, and adds support for loading the bodies from a file.

18.10.1. Additional Properties on FileDataSet

Property	Type	Default	Description
sourceFile	File	null	Specifies the source file for payloads
delimiter	String	<code>\z</code>	Specifies the delimiter pattern used by a java.util.Scanner to split the file into multiple payloads.

18.11. SPRING BOOT AUTO-CONFIGURATION

When using dataset with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-dataset-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 11 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.dataset-test.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.dataset-test.enabled</code>	Whether to enable auto configuration of the dataset-test component. This is enabled by default.		Boolean
<code>camel.component.dataset-test.exchange-formatter</code>	Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to DefaultExchangeFormatter. The option is a <code>org.apache.camel.spi.ExchangeFormatter</code> type.		ExchangeFormatter
<code>camel.component.dataset-test.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.dataset-test.log</code>	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	Boolean


Name	Description	Default	Type
<code>camel.component.dataset.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.dataset.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.dataset.enabled</code>	Whether to enable auto configuration of the dataset component. This is enabled by default.		Boolean
<code>camel.component.dataset.exchange-formatter</code>	Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to <code>DefaultExchangeFormatter</code> . The option is a <code>org.apache.camel.spi.ExchangeFormatter</code> type.		ExchangeFormatter
<code>camel.component.dataset.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.dataset.log</code>	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	Boolean

CHAPTER 19. DIRECT

Both producer and consumer are supported

The Direct component provides direct, synchronous invocation of any consumers when a producer sends a message exchange.

This endpoint can be used to connect existing routes in the **same** camel context.



NOTE

Asynchronous

The [SEDA](#) component provides asynchronous invocation of any consumers when a producer sends a message exchange.

19.1. URI FORMAT

```
direct:someName[?options]
```

Where **someName** can be any string to uniquely identify the endpoint

19.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

19.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

19.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

19.3. COMPONENT OPTIONS

The Direct component supports 5 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
block (producer)	If sending a message to a direct endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
timeout (producer)	The timeout value to use if block is enabled.	30000	long
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

19.4. ENDPOINT OPTIONS

The Direct endpoint is configured using URI syntax:

```
direct:name
```

with the following path and query parameters:

19.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required Name of direct endpoint.		String

19.4.2. Query Parameters (8 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • <code>InOnly</code> • <code>InOut</code> • <code>InOptionalOut</code> 		<code>ExchangePattern</code>
block (producer)	If sending a message to a direct endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	boolean
failIfNoConsumers (producer)	Whether the producer should fail by throwing an exception, when sending to a DIRECT endpoint with no active consumers.	true	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
timeout (producer)	The timeout value to use if block is enabled.	30000	long
synchronous (advanced)	Whether synchronous processing is forced. If enabled then the producer thread, will be forced to wait until the message has been completed before the same thread will continue processing. If disabled (default) then the producer thread may be freed and can do other work while the message is continued processed by other threads (reactive).	false	boolean

19.5. SAMPLES

In the route below we use the direct component to link the two routes together:

```
from("activemq:queue:order.in")
  .to("bean:orderServer?method=validate")
  .to("direct:processOrder");

from("direct:processOrder")
  .to("bean:orderService?method=process")
  .to("activemq:queue:order.out");
```

And the sample using spring DSL:

```
<route>
  <from uri="activemq:queue:order.in"/>
  <to uri="bean:orderService?method=validate"/>
  <to uri="direct:processOrder"/>
</route>

<route>
  <from uri="direct:processOrder"/>
```

```

<to uri="bean:orderService?method=process"/>
<to uri="activemq:queue:order.out"/>
</route>

```

See also samples from the [SEDA](#) component, how they can be used together.

19.6. SPRING BOOT AUTO-CONFIGURATION

When using direct with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-direct-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 6 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.direct.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.direct.block</code>	If sending a message to a direct endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	Boolean
<code>camel.component.direct.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.direct.enabled</code>	Whether to enable auto configuration of the direct component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component .direct.lazy-start- producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component .direct.timeout	The timeout value to use if block is enabled.	30000	Long

CHAPTER 20. FHIR

Both producer and consumer are supported

The FHIR component integrates with the [HAPI-FHIR](#) library which is an open-source implementation of the [FHIR](#) (Fast Healthcare Interoperability Resources) specification in Java.

Maven users will need to add the following dependency to their pom.xml for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-fhir</artifactId>
  <version>${camel-version}</version>
</dependency>
```

20.1. URI FORMAT

The FHIR Component uses the following URI format:

```
fhir://endpoint-prefix/endpoint?[options]
```

Endpoint prefix can be one of:

- capabilities
- create
- delete
- history
- load-page
- meta
- operation
- patch
- read
- search
- transaction
- update
- validate

20.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level

- endpoint level

20.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

20.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

20.3. COMPONENT OPTIONS

The FHIR component supports 27 options, which are listed below.

Name	Description	Default	Type
encoding (common)	Encoding to use for all request. Enum values: <ul style="list-style-type: none"> • JSON • XML 		String

Name	Description	Default	Type
fhirVersion (common)	The FHIR Version to use. Enum values: <ul style="list-style-type: none"> • DSTU2 • DSTU2_HL7ORG • DSTU2_1 • DSTU3 • R4 • R5 	R4	String
log (common)	Will log every requests and responses.	false	boolean
prettyPrint (common)	Pretty print all request.	false	boolean
serverUrl (common)	The FHIR server base URL.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
client (advanced)	To use the custom client.		IGenericClient
clientFactory (advanced)	To use the custom client factory.		IRestfulClientFactory
compress (advanced)	Compresses outgoing (POST/PUT) contents to the GZIP format.	false	boolean
configuration (advanced)	To use the shared configuration.		FhirConfiguration
connectionTimeout (advanced)	How long to try and establish the initial TCP connection (in ms).	10000	Integer
deferModelScanning (advanced)	When this option is set, model classes will not be scanned for children until the child list for the given type is actually accessed.	false	boolean
fhirContext (advanced)	FhirContext is an expensive object to create. To avoid creating multiple instances, it can be set directly.		FhirContext
forceConformanceCheck (advanced)	Force conformance check.	false	boolean
sessionCookie (advanced)	HTTP session cookie to add to every request.		String
socketTimeout (advanced)	How long to block for individual read/write operations (in ms).	10000	Integer

Name	Description	Default	Type
summary (advanced)	Request that the server modify the response using the <code>_summary</code> param. Enum values: <ul style="list-style-type: none"> ● COUNT ● TEXT ● DATA ● TRUE ● FALSE 		String
validationMode (advanced)	When should Camel validate the FHIR Server's conformance statement. Enum values: <ul style="list-style-type: none"> ● NEVER ● ONCE 	ONCE	String
proxyHost (proxy)	The proxy host.		String
proxyPassword (proxy)	The proxy password.		String
proxyPort (proxy)	The proxy port.		Integer
proxyUser (proxy)	The proxy username.		String
accessToken (security)	OAuth access token.		String
password (security)	Username to use for basic authentication.		String
username (security)	Username to use for basic authentication.		String

20.4. ENDPOINT OPTIONS

The FHIR endpoint is configured using URI syntax:

```
fhir:apiName/methodName
```

with the following path and query parameters:

20.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
apiName (common)	<p>Required What kind of operation to perform.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● CAPABILITIES ● CREATE ● DELETE ● HISTORY ● LOAD_PAGE ● META ● OPERATION ● PATCH ● READ ● SEARCH ● TRANSACTION ● UPDATE ● VALIDATE 		FhirApiName
methodName (common)	<p>Required What sub operation to use for the selected operation.</p>		String

20.4.2. Query Parameters (44 parameters)

Name	Description	Default	Type
encoding (common)	<p>Encoding to use for all request.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● JSON ● XML 		String

Name	Description	Default	Type
fhirVersion (common)	The FHIR Version to use. Enum values: <ul style="list-style-type: none"> • DSTU2 • DSTU2_HL7ORG • DSTU2_1 • DSTU3 • R4 • R5 	R4	String
inBody (common)	Sets the name of a parameter to be passed in the exchange In Body.		String
log (common)	Will log every requests and responses.	false	boolean
prettyPrint (common)	Pretty print all request.	false	boolean
serverUrl (common)	The FHIR server base URL.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
client (advanced)	To use the custom client.		IGenericClient
clientFactory (advanced)	To use the custom client factory.		IRestfulClientFactory
compress (advanced)	Compresses outgoing (POST/PUT) contents to the GZIP format.	false	boolean
connectionTimeout (advanced)	How long to try and establish the initial TCP connection (in ms).	10000	Integer
deferModelScanning (advanced)	When this option is set, model classes will not be scanned for children until the child list for the given type is actually accessed.	false	boolean
fhirContext (advanced)	FhirContext is an expensive object to create. To avoid creating multiple instances, it can be set directly.		FhirContext

Name	Description	Default	Type
forceConformanceCheck (advanced)	Force conformance check.	false	boolean
sessionCookie (advanced)	HTTP session cookie to add to every request.		String
socketTimeout (advanced)	How long to block for individual read/write operations (in ms).	10000	Integer
summary (advanced)	Request that the server modify the response using the <code>_summary</code> param. Enum values: <ul style="list-style-type: none"> ● COUNT ● TEXT ● DATA ● TRUE ● FALSE 		String
validationMode (advanced)	When should Camel validate the FHIR Server's conformance statement. Enum values: <ul style="list-style-type: none"> ● NEVER ● ONCE 	ONCE	String
proxyHost (proxy)	The proxy host.		String
proxyPassword (proxy)	The proxy password.		String
proxyPort (proxy)	The proxy port.		Integer
proxyUser (proxy)	The proxy username.		String
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.		int

Name	Description	Default	Type
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If <code>greedy</code> is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value <code>spring</code> or <code>quartz</code> for built in scheduler.	none	Object

Name	Description	Default	Type
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
accessToken (security)	OAuth access token.		String
password (security)	Username to use for basic authentication.		String
username (security)	Username to use for basic authentication.		String

20.5. API PARAMETERS (13 APIS)

The @FHIR endpoint is an API based component and has additional parameters based on which API name and API method is used. The API name and API method is located in the endpoint URI as the **apiName/methodName** path parameters:

```
fhir:apiName/methodName
```

There are 13 API names as listed in the table below:

API Name	Type	Description
capabilities	Both	API to Fetch the capability statement for the server
create	Both	API for the create operation, which creates a new resource instance on the server
delete	Both	API for the delete operation, which performs a logical delete on a server resource
history	Both	API for the history method
load-page	Both	API that Loads the previous/next bundle of resources from a paged set, using the link specified in the link type=next tag within the atom bundle
meta	Both	API for the meta operations, which can be used to get, add and remove tags and other Meta elements from a resource or across the server
operation	Both	API for extended FHIR operations
patch	Both	API for the patch operation, which performs a logical patch on a server resource
read	Both	API method for read operations
search	Both	API to search for resources matching a given set of criteria
transaction	Both	API for sending a transaction (collection of resources) to the server to be executed as a single unit
update	Both	API for the update operation, which performs a logical delete on a server resource
validate	Both	API for validating resources

Each API is documented in the following sections to come.

20.5.1. API: capabilities

Both producer and consumer are supported

The capabilities API is defined in the syntax as follows:

```
fhir:capabilities/methodName?[parameters]
```

The method is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
ofType	Retrieve the conformance statement using the given model type

20.5.1.1. Method ofType

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseConformance`
`ofType(Class<org.hl7.fhir.instance.model.api.IBaseConformance> type,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/ofType` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>type</code>	The model type	Class

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.2. API: create

Both producer and consumer are supported

The `create` API is defined in the syntax as follows:

```
fhir:create/methodName?[parameters]
```

The 1 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
resource	Creates a <code>IBaseResource</code> on the server

20.5.2.1. Method resource

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome resource(String resourceAsString, String url, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `ca.uhn.fhir.rest.api.MethodOutcome resource(org.hl7.fhir.instance.model.api.IBaseResource resource, String url, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resource` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>preferReturn</code>	Add a <code>Prefer</code> header to the request, which requests that the server include or suppress the resource body as a part of the result. If a resource is returned by the server it will be parsed an accessible to the client via <code>MethodOutcome#getResource()</code> , may be null	<code>PreferReturnEnum</code>
<code>resource</code>	The resource to create	<code>IBaseResource</code>
<code>resourceAsString</code>	The resource to create	String
<code>url</code>	The search URL to use. The format of this URL should be of the form <code>ResourceTypeParameters</code> , for example: <code>Patientname=Smith&identifier=13.2.4.11.4%7C847366</code> , may be null	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.3. API: delete

Both producer and consumer are supported

The delete API is defined in the syntax as follows:

```
fhir:delete/methodName?[parameters]
```

The 3 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
resource	Deletes the given resource
resourceById	Deletes the resource by resource type e
resourceConditionalByUrl	Specifies that the delete should be performed as a conditional delete against a given search URL

20.5.3.1. Method resource

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseOperationOutcome resource(org.hl7.fhir.instance.model.api.IBaseResource resource, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resource` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
resource	The <code>IBaseResource</code> to delete	<code>IBaseResource</code>

20.5.3.2. Method resourceById

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseOperationOutcome resourceById(String type, String stringId, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseOperationOutcome resourceById(org.hl7.fhir.instance.model.api.IIdType id, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resourceById` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
id	The <code>IIdType</code> referencing the resource	<code>IIdType</code>
stringId	It's id	String

Parameter	Description	Type
<code>type</code>	The resource type e.g Patient	String

20.5.3.3. Method `resourceConditionalByUrl`

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseOperationOutcome resourceConditionalByUrl(String url, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resourceConditionalByUrl` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>url</code>	The search URL to use. The format of this URL should be of the form <code>ResourceTypeParameters</code> , for example: <code>Patientname=Smith&identifier=13.2.4.11.4%7C847366</code>	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.4. API: history

Both producer and consumer are supported

The history API is defined in the syntax as follows:

```
fhir:history/methodName?[parameters]
```

The 3 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
<code>onInstance</code>	Perform the operation across all versions of a specific resource (by ID and type) on the server
<code>onServer</code>	Perform the operation across all versions of all resources of all types on the server

Method	Description
onType	Perform the operation across all versions of all resources of the given type on the server

20.5.4.1. Method onInstance

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseBundle onInstance(org.hl7.fhir.instance.model.api.IIdType id, Class<org.hl7.fhir.instance.model.api.IBaseBundle> returnType, Integer count, java.util.Date cutoff, org.hl7.fhir.instance.model.api.IPrimitiveType<java.util.Date> iCutoff, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The `fhir/onInstance` API method has the parameters listed in the table below:

| Parameter              | Description                                                                                                                                                            | Type           |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| <b>count</b>           | Request that the server return only up to theCount number of resources, may be NULL                                                                                    | Integer        |
| <b>cutoff</b>          | Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL                                            | Date           |
| <b>extraParameters</b> | See ExtraParameters for a full list of parameters that can be passed, may be NULL                                                                                      | Map            |
| <b>iCutoff</b>         | Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL                                            | IPrimitiveType |
| <b>id</b>              | The IIdType which must be populated with both a resource type and a resource ID at                                                                                     | IIdType        |
| <b>returnType</b>      | Request that the method return a Bundle resource (such as <code>ca.uhn.fhir.model.dstu2.resource.Bundle</code> ). Use this method if you are accessing a DSTU2 server. | Class          |

#### 20.5.4.2. Method onServer

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseBundle onServer(Class<org.hl7.fhir.instance.model.api.IBaseBundle> returnType, Integer count, java.util.Date cutoff, org.hl7.fhir.instance.model.api.IPrimitiveType<java.util.Date> iCutoff, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The `fhir/onServer` API method has the parameters listed in the table below:

Parameter	Description	Type
count	Request that the server return only up to theCount number of resources, may be NULL	Integer
cutoff	Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL	Date
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
iCutoff	Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL	IPrimitiveType
returnType	Request that the method return a Bundle resource (such as ca.uhn.fhir.model.dstu2.resource.Bundle). Use this method if you are accessing a DSTU2 server.	Class

20.5.4.3. Method onType

Signatures:

- org.hl7.fhir.instance.model.api.IBaseBundle
 onType(Class<org.hl7.fhir.instance.model.api.IBaseResource> resourceType,
 Class<org.hl7.fhir.instance.model.api.IBaseBundle> returnType, Integer count, java.util.Date
 cutoff, org.hl7.fhir.instance.model.api.IPrimitiveType<java.util.Date> iCutoff,
 java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);

The fhir/onType API method has the parameters listed in the table below:

Parameter	Description	Type
count	Request that the server return only up to theCount number of resources, may be NULL	Integer
cutoff	Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL	Date
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
iCutoff	Request that the server return only resource versions that were created at or after the given time (inclusive), may be NULL	IPrimitiveType
resourceType	The resource type to search for	Class

Parameter	Description	Type
returnType	Request that the method return a Bundle resource (such as <code>ca.uhn.fhir.model.dstu2.resource.Bundle</code>). Use this method if you are accessing a DSTU2 server.	Class

In addition to the parameters above, the fhir API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.5. API: load-page

Both producer and consumer are supported

The load-page API is defined in the syntax as follows:

```
fhir:load-page/methodName?[parameters]
```

The 3 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
byUrl	Load a page of results using the given URL and bundle type and return a DSTU1 Atom bundle
next	Load the next page of results using the link with relation next in the bundle
previous	Load the previous page of results using the link with relation prev in the bundle

20.5.5.1. Method byUrl

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseBundle byUrl(String url, Class<org.hl7.fhir.instance.model.api.IBaseBundle> returnType, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The fhir/byUrl API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map

Parameter	Description	Type
<code>returnType</code>	The return type	Class
<code>url</code>	The search url	String

20.5.5.2. Method next

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseBundle next(org.hl7.fhir.instance.model.api.IBaseBundle bundle, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/next` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>bundle</code>	The <code>IBaseBundle</code>	<code>IBaseBundle</code>
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be <code>NULL</code>	Map

20.5.5.3. Method previous

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseBundle previous(org.hl7.fhir.instance.model.api.IBaseBundle bundle, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/previous` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>bundle</code>	The <code>IBaseBundle</code>	<code>IBaseBundle</code>
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be <code>NULL</code>	Map

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.6. API: meta

Both producer and consumer are supported

The meta API is defined in the syntax as follows:

```
fhir:meta/methodName?[parameters]
```

The 5 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
add	Add the elements in the given metadata to the already existing set (do not remove any)
delete	Delete the elements in the given metadata from the given id
getFromResource	Fetch the current metadata from a specific resource
getFromServer	Fetch the current metadata from the whole Server
getFromType	Fetch the current metadata from a specific type

20.5.6.1. Method add

Signatures:

- org.hl7.fhir.instance.model.api.IBaseMetaType
 add(org.hl7.fhir.instance.model.api.IBaseMetaType meta, org.hl7.fhir.instance.model.api.IIdType id, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);

The fhir/add API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
id	The id	IIdType
meta	The IBaseMetaType class	IBaseMetaType

20.5.6.2. Method delete

Signatures:

- org.hl7.fhir.instance.model.api.IBaseMetaType
 delete(org.hl7.fhir.instance.model.api.IBaseMetaType meta, org.hl7.fhir.instance.model.api.IIdType id, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);

The fhir/delete API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
id	The id	IIdType
meta	The IBaseMetaType class	IBaseMetaType

20.5.6.3. Method getFromResource

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseMetaType
getFromResource(Class<org.hl7.fhir.instance.model.api.IBaseMetaType> metaType,
org.hl7.fhir.instance.model.api.IIdType id,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The fhir/getFromResource API method has the parameters listed in the table below:

| Parameter              | Description                                                                       | Type    |
|------------------------|-----------------------------------------------------------------------------------|---------|
| <b>extraParameters</b> | See ExtraParameters for a full list of parameters that can be passed, may be NULL | Map     |
| <b>id</b>              | The id                                                                            | IIdType |
| <b>metaType</b>        | The IBaseMetaType class                                                           | Class   |

### 20.5.6.4. Method getFromServer

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseMetaType
getFromServer(Class<org.hl7.fhir.instance.model.api.IBaseMetaType> metaType,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The fhir/getFromServer API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
metaType	The type of the meta datatype for the given FHIR model version (should be MetaDt.class or MetaType.class)	Class

20.5.6.5. Method getFromType

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseMetaType`
`getFromType(Class<org.hl7.fhir.instance.model.api.IBaseMetaType> metaType, String resourceType, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/getFromType` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>metaType</code>	The <code>IBaseMetaType</code> class	Class
<code>resourceType</code>	The resource type e.g Patient	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.7. API: operation

Both producer and consumer are supported

The operation API is defined in the syntax as follows:

```
fhir:operation/methodName?[parameters]
```

The 5 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
<code>onInstance</code>	Perform the operation across all versions of a specific resource (by ID and type) on the server
<code>onInstanceVersion</code>	This operation operates on a specific version of a resource
<code>onServer</code>	Perform the operation across all versions of all resources of all types on the server
<code>onType</code>	Perform the operation across all versions of all resources of the given type on the server

Method	Description
processMessage	This operation is called \$process-message as defined by the FHIR specification

20.5.7.1. Method onInstance

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseResource onInstance(org.hl7.fhir.instance.model.api.IIdType id, String name, org.hl7.fhir.instance.model.api.IBaseParameters parameters, Class<org.hl7.fhir.instance.model.api.IBaseParameters> outputParameterType, boolean useHttpGet, Class<org.hl7.fhir.instance.model.api.IBaseResource> returnType, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The `fhir/onInstance` API method has the parameters listed in the table below:

| Parameter                  | Description                                                                                                                                                                                                                                                                                                                            | Type            |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>extraParameters</b>     | See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL                                                                                                                                                                                                                                         | Map             |
| <b>id</b>                  | Resource (version will be stripped)                                                                                                                                                                                                                                                                                                    | IIdType         |
| <b>name</b>                | Operation name                                                                                                                                                                                                                                                                                                                         | String          |
| <b>outputParameterType</b> | The type to use for the output parameters (this should be set to <code>Parameters.class</code> drawn from the version of the FHIR structures you are using), may be NULL                                                                                                                                                               | Class           |
| <b>parameters</b>          | The parameters to use as input. May also be null if the operation does not require any input parameters.                                                                                                                                                                                                                               | IBaseParameters |
| <b>returnType</b>          | If this operation returns a single resource body as its return type instead of a <code>Parameters</code> resource, use this method to specify that resource type. This is useful for certain operations (e.g. <code>Patient/NNN/\$everything</code> ) which return a bundle instead of a <code>Parameters</code> resource, may be NULL | Class           |
| <b>useHttpGet</b>          | Use HTTP GET verb                                                                                                                                                                                                                                                                                                                      | Boolean         |

### 20.5.7.2. Method onInstanceVersion

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseResource onInstanceVersion(org.hl7.fhir.instance.model.api.IIdType id, String name,
```



```
org.hl7.fhir.instance.model.api.IBaseParameters parameters,
Class<org.hl7.fhir.instance.model.api.IBaseParameters> outputParameterType, boolean
useHttpGet, Class<org.hl7.fhir.instance.model.api.IBaseResource> returnType,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The `fhir/onInstanceVersion` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>id</code>	Resource version	IIdType
<code>name</code>	Operation name	String
<code>outputParameterType</code>	The type to use for the output parameters (this should be set to <code>Parameters.class</code> drawn from the version of the FHIR structures you are using), may be NULL	Class
<code>parameters</code>	The parameters to use as input. May also be null if the operation does not require any input parameters.	IBaseParameters
<code>returnType</code>	If this operation returns a single resource body as its return type instead of a <code>Parameters</code> resource, use this method to specify that resource type. This is useful for certain operations (e.g. <code>Patient/NNN/\$everything</code>) which return a bundle instead of a <code>Parameters</code> resource, may be NULL	Class
<code>useHttpGet</code>	Use HTTP GET verb	Boolean

20.5.7.3. Method onServer

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseResource onServer(String name, org.hl7.fhir.instance.model.api.IBaseParameters parameters, Class<org.hl7.fhir.instance.model.api.IBaseParameters> outputParameterType, boolean useHttpGet, Class<org.hl7.fhir.instance.model.api.IBaseResource> returnType, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/onServer` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map

Parameter	Description	Type
name	Operation name	String
outputParameterType	The type to use for the output parameters (this should be set to Parameters.class drawn from the version of the FHIR structures you are using), may be NULL	Class
parameters	The parameters to use as input. May also be null if the operation does not require any input parameters.	IBaseParameters
returnType	If this operation returns a single resource body as its return type instead of a Parameters resource, use this method to specify that resource type. This is useful for certain operations (e.g. Patient/NNN/\$everything) which return a bundle instead of a Parameters resource, may be NULL	Class
useHttpGet	Use HTTP GET verb	Boolean

20.5.7.4. Method onType

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseResource
onType(Class<org.hl7.fhir.instance.model.api.IBaseResource> resourceType, String name,
org.hl7.fhir.instance.model.api.IBaseParameters parameters,
Class<org.hl7.fhir.instance.model.api.IBaseParameters> outputParameterType, boolean
useHttpGet, Class<org.hl7.fhir.instance.model.api.IBaseResource> returnType,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The fhir/onType API method has the parameters listed in the table below:

| Parameter                  | Description                                                                                                                                                 | Type            |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>extraParameters</b>     | See ExtraParameters for a full list of parameters that can be passed, may be NULL                                                                           | Map             |
| <b>name</b>                | Operation name                                                                                                                                              | String          |
| <b>outputParameterType</b> | The type to use for the output parameters (this should be set to Parameters.class drawn from the version of the FHIR structures you are using), may be NULL | Class           |
| <b>parameters</b>          | The parameters to use as input. May also be null if the operation does not require any input parameters.                                                    | IBaseParameters |
| <b>resourceType</b>        | The resource type to operate on                                                                                                                             | Class           |

| Parameter               | Description                                                                                                                                                                                                                                                                                    | Type    |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
| <code>returnType</code> | If this operation returns a single resource body as its return type instead of a Parameters resource, use this method to specify that resource type. This is useful for certain operations (e.g. Patient/NNN/\$everything) which return a bundle instead of a Parameters resource, may be NULL | Class   |
| <code>useHttpGet</code> | Use HTTP GET verb                                                                                                                                                                                                                                                                              | Boolean |

### 20.5.7.5. Method processMessage

Signatures:

- ```
org.hl7.fhir.instance.model.api.IBaseBundle processMessage(String respondToUri,
org.hl7.fhir.instance.model.api.IBaseBundle msgBundle, boolean asynchronous,
Class<org.hl7.fhir.instance.model.api.IBaseBundle> responseClass,
java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
```

The `fhir/processMessage` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>asynchronous</code>	Whether to process the message asynchronously or synchronously, defaults to synchronous.	Boolean
<code>extraParameters</code>	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
<code>msgBundle</code>	Set the Message Bundle to POST to the messaging server	IBaseBundle
<code>respondToUri</code>	An optional query parameter indicating that responses from the receiving server should be sent to this URI, may be NULL	String
<code>responseClass</code>	The response class	Class

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.8. API: patch

Both producer and consumer are supported

The patch API is defined in the syntax as follows:

```
fhir:patch/methodName?[parameters]
```

The 2 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
<code>patchById</code>	Applies the patch to the given resource ID
<code>patchByUrl</code>	Specifies that the update should be performed as a conditional create against a given search URL

20.5.8.1. Method `patchById`

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome patchById(String patchBody, String stringId, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `ca.uhn.fhir.rest.api.MethodOutcome patchById(String patchBody, org.hl7.fhir.instance.model.api.IIdType id, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/patchById` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>id</code>	The resource ID to patch	IIdType
<code>patchBody</code>	The body of the patch document serialized in either XML or JSON which conforms to	String
<code>preferReturn</code>	Add a <code>Prefer</code> header to the request, which requests that the server include or suppress the resource body as a part of the result. If a resource is returned by the server it will be parsed and accessible to the client via <code>MethodOutcome#getResource()</code>	PreferReturnEnum
<code>stringId</code>	The resource ID to patch	String

20.5.8.2. Method `patchByUrl`

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome patchByUrl(String patchBody, String url, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/patchByUrl` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>patchBody</code>	The body of the patch document serialized in either XML or JSON which conforms to	String
<code>preferReturn</code>	Add a <code>Prefer</code> header to the request, which requests that the server include or suppress the resource body as a part of the result. If a resource is returned by the server it will be parsed and accessible to the client via <code>MethodOutcome#getResource()</code>	<code>PreferReturnEnum</code>
<code>url</code>	The search URL to use. The format of this URL should be of the form <code>ResourceTypeParameters</code> , for example: <code>Patientname=Smith&identifier=13.2.4.11.4%7C847366</code>	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.9. API: read

Both producer and consumer are supported

The `read` API is defined in the syntax as follows:

```
fhir:read/methodName?[parameters]
```

The 2 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
resourceById	Reads a <code>IBaseResource</code> on the server by id
resourceByUrl	Reads a <code>IBaseResource</code> on the server by url

20.5.9.1. Method resourceById

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(Class<org.hl7.fhir.instance.model.api.IBaseResource> resource, Long longId, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(Class<org.hl7.fhir.instance.model.api.IBaseResource> resource, String stringId, String version, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(Class<org.hl7.fhir.instance.model.api.IBaseResource> resource, org.hl7.fhir.instance.model.api.IIdType id, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(String resourceClass, Long longId, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(String resourceClass, String stringId, String ifVersionMatches, String version, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceById(String resourceClass, org.hl7.fhir.instance.model.api.IIdType id, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resourceById` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be <code>NULL</code>	Map
id	The <code>IIdType</code> referencing the resource	<code>IIdType</code>
ifVersionMatches	A version to match against the newest version on the server	String
longId	The resource ID	Long
resource	The resource to read (e.g. Patient)	Class

Parameter	Description	Type
resourceClass	The resource to read (e.g. Patient)	String
returnNull	Return null if version matches	Boolean
returnResource	Return the resource if version matches	IBaseResource
stringId	The resource ID	String
throwError	Throw error if the version matches	Boolean
version	The resource version	String

20.5.9.2. Method `resourceByUrl`

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseResource resourceByUrl(Class<org.hl7.fhir.instance.model.api.IBaseResource> resource, String url, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceByUrl(Class<org.hl7.fhir.instance.model.api.IBaseResource> resource, org.hl7.fhir.instance.model.api.IIdType iUrl, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceByUrl(String resourceClass, String url, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `org.hl7.fhir.instance.model.api.IBaseResource resourceByUrl(String resourceClass, org.hl7.fhir.instance.model.api.IIdType iUrl, String ifVersionMatches, Boolean returnNull, org.hl7.fhir.instance.model.api.IBaseResource returnResource, Boolean throwError, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resourceByUrl` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be <code>NULL</code>	Map
iUrl	The <code>IIdType</code> referencing the resource by absolute url	<code>IIdType</code>

Parameter	Description	Type
ifVersionMatches	A version to match against the newest version on the server	String
resource	The resource to read (e.g. Patient)	Class
resourceClass	The resource to read (e.g. Patient.class)	String
returnNull	Return null if version matches	Boolean
returnResource	Return the resource if version matches	IBaseResource
throwError	Throw error if the version matches	Boolean
url	Referencing the resource by absolute url	String

In addition to the parameters above, the fhir API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.10. API: search

Both producer and consumer are supported

The search API is defined in the syntax as follows:

```
fhir:search/methodName?[parameters]
```

The 1 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
searchByUrl	Perform a search directly by URL

20.5.10.1. Method searchByUrl

Signatures:

- `org.hl7.fhir.instance.model.api.IBaseBundle searchByUrl(String url, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/searchByUrl` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
url	The URL to search for. Note that this URL may be complete (e.g.) in which case the client's base URL will be ignored. Or it can be relative (e.g. Patientname=foo) in which case the client's base URL will be used.	String

In addition to the parameters above, the fhir API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.11. API: transaction

Both producer and consumer are supported

The transaction API is defined in the syntax as follows:

```
fhir:transaction/methodName?[parameters]
```

The 2 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
withBundle	Use the given raw text (should be a Bundle resource) as the transaction input
withResources	Use a list of resources as the transaction input

20.5.11.1. Method withBundle

Signatures:

- String withBundle(String stringBundle, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);
- org.hl7.fhir.instance.model.api.IBaseBundle withBundle(org.hl7.fhir.instance.model.api.IBaseBundle bundle, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);

The fhir/withBundle API method has the parameters listed in the table below:

Parameter	Description	Type
bundle	Bundle to use in the transaction	IBaseBundle
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
stringBundle	Bundle to use in the transaction	String

20.5.11.2. Method withResources

Signatures:

- `java.util.List<org.hl7.fhir.instance.model.api.IBaseResource> withResources(java.util.List<org.hl7.fhir.instance.model.api.IBaseResource> resources, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The fhir/withResources API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See ExtraParameters for a full list of parameters that can be passed, may be NULL	Map
resources	Resources to use in the transaction	List

In addition to the parameters above, the fhir API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.12. API: update

Both producer and consumer are supported

The update API is defined in the syntax as follows:

```
fhir:update/methodName?[parameters]
```

The 2 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
resource	Updates a IBaseResource on the server by id

Method	Description
resourceByUrl	Updates a IBaseResource on the server by search url

20.5.12.1. Method resource

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome resource(String resourceAsString, String stringId, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `ca.uhn.fhir.rest.api.MethodOutcome resource(String resourceAsString, org.hl7.fhir.instance.model.api.IIdType id, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `ca.uhn.fhir.rest.api.MethodOutcome resource(org.hl7.fhir.instance.model.api.IBaseResource resource, String stringId, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`
- `ca.uhn.fhir.rest.api.MethodOutcome resource(org.hl7.fhir.instance.model.api.IBaseResource resource, org.hl7.fhir.instance.model.api.IIdType id, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resource` API method has the parameters listed in the table below:

Parameter	Description	Type
extraParameters	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
id	The <code>IIdType</code> referencing the resource	<code>IIdType</code>
preferReturn	Whether the server include or suppress the resource body as a part of the result	<code>PreferReturnEnum</code>
resource	The resource to update (e.g. Patient)	<code>IBaseResource</code>
resourceAsString	The resource body to update	String
stringId	The ID referencing the resource	String

20.5.12.2. Method resourceByUrl

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome resourceByUrl(String resourceAsString, String url, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

- `ca.uhn.fhir.rest.api.MethodOutcome resourceBySearchUrl(org.hl7.fhir.instance.model.api.IBaseResource resource, String url, ca.uhn.fhir.rest.api.PreferReturnEnum preferReturn, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resourceBySearchUrl` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>preferReturn</code>	Whether the server include or suppress the resource body as a part of the result	PreferReturnEnum
<code>resource</code>	The resource to update (e.g. Patient)	IBaseResource
<code>resourceAsString</code>	The resource body to update	String
<code>url</code>	Specifies that the update should be performed as a conditional create against a given search URL	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.5.13. API: validate

Both producer and consumer are supported

The `validate` API is defined in the syntax as follows:

```
fhir:validate/methodName?[parameters]
```

The 1 method(s) is listed in the table below, followed by detailed syntax for each method. (API methods can have a shorthand alias name which can be used in the syntax instead of the name)

Method	Description
resource	Validates the resource

20.5.13.1. Method resource

Signatures:

- `ca.uhn.fhir.rest.api.MethodOutcome resource(String resourceAsString, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

- `ca.uhn.fhir.rest.api.MethodOutcome resource(org.hl7.fhir.instance.model.api.IBaseResource resource, java.util.Map<org.apache.camel.component.fhir.api.ExtraParameters, Object> extraParameters);`

The `fhir/resource` API method has the parameters listed in the table below:

Parameter	Description	Type
<code>extraParameters</code>	See <code>ExtraParameters</code> for a full list of parameters that can be passed, may be NULL	Map
<code>resource</code>	The <code>IBaseResource</code> to validate	<code>IBaseResource</code>
<code>resourceAsString</code>	Raw resource to validate	String

In addition to the parameters above, the `fhir` API can also use any of the [Query Parameters](#).

Any of the parameters can be provided in either the endpoint URI, or dynamically in a message header. The message header name must be of the format **CamelFhir.parameter**. The **inBody** parameter overrides message header, i.e. the endpoint parameter **inBody=myParameterNameHere** would override a **CamelFhir.myParameterNameHere** header.

20.6. SPRING BOOT AUTO-CONFIGURATION

When using `fhir` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-fhir-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 56 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.fhir.access-token</code>	OAuth access token.		String
<code>camel.component.fhir.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.fhir.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.fhir.client	To use the custom client. The option is a <code>ca.uhn.fhir.rest.client.api.IGenericClient</code> type.		IGenericClient
camel.component.fhir.client-factory	To use the custom client factory. The option is a <code>ca.uhn.fhir.rest.client.api.IRestfulClientFactory</code> type.		IRestfulClientFactory
camel.component.fhir.compress	Compresses outgoing (POST/PUT) contents to the GZIP format.	false	Boolean
camel.component.fhir.configuration	To use the shared configuration. The option is a <code>org.apache.camel.component.fhir.FhirConfiguration</code> type.		FhirConfiguration
camel.component.fhir.connection-timeout	How long to try and establish the initial TCP connection (in ms).	10000	Integer
camel.component.fhir.defer-model-scanning	When this option is set, model classes will not be scanned for children until the child list for the given type is actually accessed.	false	Boolean
camel.component.fhir.enabled	Whether to enable auto configuration of the fhir component. This is enabled by default.		Boolean
camel.component.fhir.encoding	Encoding to use for all request.		String
camel.component.fhir.fhir-context	FhirContext is an expensive object to create. To avoid creating multiple instances, it can be set directly. The option is a <code>ca.uhn.fhir.context.FhirContext</code> type.		FhirContext
camel.component.fhir.fhir-version	The FHIR Version to use.	R4	String

Name	Description	Default	Type
<code>camel.component.fhir.force-conformance-check</code>	Force conformance check.	false	Boolean
<code>camel.component.fhir.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.fhir.log</code>	Will log every requests and responses.	false	Boolean
<code>camel.component.fhir.password</code>	Username to use for basic authentication.		String
<code>camel.component.fhir.pretty-print</code>	Pretty print all request.	false	Boolean
<code>camel.component.fhir.proxy-host</code>	The proxy host.		String
<code>camel.component.fhir.proxy-password</code>	The proxy password.		String
<code>camel.component.fhir.proxy-port</code>	The proxy port.		Integer
<code>camel.component.fhir.proxy-user</code>	The proxy username.		String
<code>camel.component.fhir.server-url</code>	The FHIR server base URL.		String
<code>camel.component.fhir.session-cookie</code>	HTTP session cookie to add to every request.		String

Name	Description	Default	Type
<code>camel.component.fhir.socket-timeout</code>	How long to block for individual read/write operations (in ms).	10000	Integer
<code>camel.component.fhir.summary</code>	Request that the server modify the response using the <code>_summary</code> param.		String
<code>camel.component.fhir.username</code>	Username to use for basic authentication.		String
<code>camel.component.fhir.validation-mode</code>	When should Camel validate the FHIR Server's conformance statement.	ONCE	String
<code>camel.dataformat.fhirjson.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.fhirjson.dont-encode-elements</code>	If provided, specifies the elements which should NOT be encoded. Valid values for this field would include: Patient - Don't encode patient and all its children Patient.name - Don't encode the patient's name Patient.name.family - Don't encode the patient's family name .text - Don't encode the text element on any resource (only the very first position may contain a wildcard) DSTU2 note: Note that values including meta, such as <code>Patient.meta</code> will work for DSTU2 parsers, but values with subelements on meta such as <code>Patient.meta.lastUpdated</code> will only work in DSTU3 mode.		Set
<code>camel.dataformat.fhirjson.dont-strip-versions-from-references-at-paths</code>	If supplied value(s), any resource references at the specified paths will have their resource versions encoded instead of being automatically stripped during the encoding process. This setting has no effect on the parsing process. This method provides a finer-grained level of control than <code>setStripVersionsFromReferences(String)</code> and any paths specified by this method will be encoded even if <code>setStripVersionsFromReferences(String)</code> has been set to true (which is the default).		List
<code>camel.dataformat.fhirjson.enabled</code>	Whether to enable auto configuration of the <code>fhirJson</code> data format. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.dataformat.fhirjson.encode-elements</code>	If provided, specifies the elements which should be encoded, to the exclusion of all others. Valid values for this field would include: Patient - Encode patient and all its children Patient.name - Encode only the patient's name Patient.name.family - Encode only the patient's family name .text - Encode the text element on any resource (only the very first position may contain a wildcard) .(mandatory) - This is a special case which causes any mandatory fields (min 0) to be encoded.		Set
<code>camel.dataformat.fhirjson.encode-elements-applies-to-child-resources-only</code>	If set to true (default is false), the values supplied to <code>setEncodeElements(Set)</code> will not be applied to the root resource (typically a Bundle), but will be applied to any sub-resources contained within it (i.e. search result resources in that bundle).	false	Boolean
<code>camel.dataformat.fhirjson.fhir-version</code>	The version of FHIR to use. Possible values are: DSTU2,DSTU2_HL7ORG,DSTU2_1,DSTU3,R4.	DSTU3	String
<code>camel.dataformat.fhirjson.omit-resource-id</code>	If set to true (default is false) the ID of any resources being encoded will not be included in the output. Note that this does not apply to contained resources, only to root resources. In other words, if this is set to true, contained resources will still have local IDs but the outer/containing ID will not have an ID.	false	Boolean
<code>camel.dataformat.fhirjson.override-resource-id-with-bundle-entry-full-url</code>	If set to true (which is the default), the <code>Bundle.entry.fullUrl</code> will override the <code>Bundle.entry.resource</code> 's resource id if the <code>fullUrl</code> is defined. This behavior happens when parsing the source data into a Bundle object. Set this to false if this is not the desired behavior (e.g. the client code wishes to perform additional validation checks between the <code>fullUrl</code> and the resource id).	false	Boolean
<code>camel.dataformat.fhirjson.pretty-print</code>	Sets the pretty print flag, meaning that the parser will encode resources with human-readable spacing and newlines between elements instead of condensing output as much as possible.	false	Boolean
<code>camel.dataformat.fhirjson.server-base-url</code>	Sets the server's base URL used by this parser. If a value is set, resource references will be turned into relative references if they are provided as absolute URLs but have a base matching the given base.		String

Name	Description	Default	Type
camel.dataformat.fhirjson.strip-versions-from-references	If set to true (which is the default), resource references containing a version will have the version removed when the resource is encoded. This is generally good behaviour because in most situations, references from one resource to another should be to the resource by ID, not by ID and version. In some cases though, it may be desirable to preserve the version in resource links. In that case, this value should be set to false. This method provides the ability to globally disable reference encoding. If finer-grained control is needed, use <code>setDontStripVersionsFromReferencesAtPath(List)</code> .	false	Boolean
camel.dataformat.fhirjson.summary-mode	If set to true (default is false) only elements marked by the FHIR specification as being summary elements will be included.	false	Boolean
camel.dataformat.fhirjson.suppress-narratives	If set to true (default is false), narratives will not be included in the encoded values.	false	Boolean
camel.dataformat.fhirxml.content-type-header	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.	true	Boolean
camel.dataformat.fhirxml.dont-encode-elements	If provided, specifies the elements which should NOT be encoded. Valid values for this field would include: Patient - Don't encode patient and all its children Patient.name - Don't encode the patient's name Patient.name.family - Don't encode the patient's family name .text - Don't encode the text element on any resource (only the very first position may contain a wildcard) DSTU2 note: Note that values including meta, such as <code>Patient.meta</code> will work for DSTU2 parsers, but values with subelements on meta such as <code>Patient.meta.lastUpdated</code> will only work in DSTU3 mode.		Set

Name	Description	Default	Type
<code>camel.dataformat.fhirxml.dont-strip-versions-from-references-at-paths</code>	If supplied value(s), any resource references at the specified paths will have their resource versions encoded instead of being automatically stripped during the encoding process. This setting has no effect on the parsing process. This method provides a finer-grained level of control than <code>setStripVersionsFromReferences(String)</code> and any paths specified by this method will be encoded even if <code>setStripVersionsFromReferences(String)</code> has been set to true (which is the default).		List
<code>camel.dataformat.fhirxml.enabled</code>	Whether to enable auto configuration of the fhirXml data format. This is enabled by default.		Boolean
<code>camel.dataformat.fhirxml.encode-elements</code>	If provided, specifies the elements which should be encoded, to the exclusion of all others. Valid values for this field would include: Patient - Encode patient and all its children Patient.name - Encode only the patient's name Patient.name.family - Encode only the patient's family name .text - Encode the text element on any resource (only the very first position may contain a wildcard) .(mandatory) - This is a special case which causes any mandatory fields (min 0) to be encoded.		Set
<code>camel.dataformat.fhirxml.encode-elements-applies-to-child-resources-only</code>	If set to true (default is false), the values supplied to <code>setEncodeElements(Set)</code> will not be applied to the root resource (typically a Bundle), but will be applied to any sub-resources contained within it (i.e. search result resources in that bundle).	false	Boolean
<code>camel.dataformat.fhirxml.fhir-version</code>	The version of FHIR to use. Possible values are: DSTU2,DSTU2_HL7ORG,DSTU2_1,DSTU3,R4.	DSTU3	String
<code>camel.dataformat.fhirxml.omit-resource-id</code>	If set to true (default is false) the ID of any resources being encoded will not be included in the output. Note that this does not apply to contained resources, only to root resources. In other words, if this is set to true, contained resources will still have local IDs but the outer/containing ID will not have an ID.	false	Boolean

Name	Description	Default	Type
<code>camel.dataformat.fhirxml.override-resource-id-with-bundle-entry-full-url</code>	If set to true (which is the default), the <code>Bundle.entry.fullUrl</code> will override the <code>Bundle.entry.resource</code> 's resource id if the <code>fullUrl</code> is defined. This behavior happens when parsing the source data into a <code>Bundle</code> object. Set this to false if this is not the desired behavior (e.g. the client code wishes to perform additional validation checks between the <code>fullUrl</code> and the resource id).	false	Boolean
<code>camel.dataformat.fhirxml.pretty-print</code>	Sets the pretty print flag, meaning that the parser will encode resources with human-readable spacing and newlines between elements instead of condensing output as much as possible.	false	Boolean
<code>camel.dataformat.fhirxml.server-base-url</code>	Sets the server's base URL used by this parser. If a value is set, resource references will be turned into relative references if they are provided as absolute URLs but have a base matching the given base.		String
<code>camel.dataformat.fhirxml.strip-versions-from-references</code>	If set to true (which is the default), resource references containing a version will have the version removed when the resource is encoded. This is generally good behaviour because in most situations, references from one resource to another should be to the resource by ID, not by ID and version. In some cases though, it may be desirable to preserve the version in resource links. In that case, this value should be set to false. This method provides the ability to globally disable reference encoding. If finer-grained control is needed, use <code>setDontStripVersionsFromReferencesAtPath(List)</code> .	false	Boolean
<code>camel.dataformat.fhirxml.summary-mode</code>	If set to true (default is false) only elements marked by the FHIR specification as being summary elements will be included.	false	Boolean
<code>camel.dataformat.fhirxml.suppress-narratives</code>	If set to true (default is false), narratives will not be included in the encoded values.	false	Boolean

CHAPTER 21. FILE

Both producer and consumer are supported

The File component provides access to file systems, allowing files to be processed by any other Camel Components or messages from other components to be saved to disk.

21.1. URI FORMAT

```
file:directoryName[?options]
```

Where **directoryName** represents the underlying file directory.

Only directories

Camel supports only endpoints configured with a starting directory. So the **directoryName** must be a directory. If you want to consume a single file only, you can use the **fileName** option, e.g. by setting **fileName=thefilename**. Also, the starting directory must not contain dynamic expressions with **\${ }** placeholders. Again use the **fileName** option to specify the dynamic part of the filename.



NOTE

Avoid reading files currently being written by another application

Beware the JDK File IO API is a bit limited in detecting whether another application is currently writing/copying a file. And the implementation can be different depending on OS platform as well. This could lead to that Camel thinks the file is not locked by another process and start consuming it. Therefore you have to do your own investigation what suits your environment. To help with this Camel provides different **readLock** options and **doneFileName** option that you can use. See also the section [Consuming files from folders where others drop files directly](#).

21.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

21.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

21.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

21.3. COMPONENT OPTIONS

The File component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

21.4. ENDPOINT OPTIONS

The File endpoint is configured using URI syntax:

`file:directoryName`

with the following path and query parameters:

21.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
directoryName (common)	Required The starting directory.	t	File

21.4.2. Query Parameters (94 parameters)

Name	Description	Default	Type
charset (common)	This option is used to specify the encoding of the file. You can use this on the consumer, to specify the encodings of the files, which allow Camel to know the charset it should load the file content in case the file content is being accessed. Likewise when writing a file, you can use this option to specify which charset to write the file as well. Do mind that when writing the file Camel may have to read the message content into memory to be able to convert the data into the configured charset, so do not use this if you have big messages.		String
doneFileName (common)	Producer: If provided, then Camel will write a 2nd done file when the original file has been written. The done file will be empty. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file will always be written in the same folder as the original file. Consumer: If provided, Camel will only consume files if a done file exists. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file is always expected in the same folder as the original file. Only <code>file.name</code> and <code>file.name.next</code> is supported as dynamic placeholders.		String

Name	Description	Default	Type
fileName (common)	Use Expression such as File Language to dynamically set the filename. For consumers, it's used as a filename filter. For producers, it's used to evaluate the filename to write. If an expression is set, it take precedence over the CamelFileName header. (Note: The header itself can also be an Expression). The expression options support both String and Expression types. If the expression is a String type, it is always evaluated using the File Language. If the expression is an Expression type, the specified Expression type is used - this allows you, for instance, to use OGNL expressions. For the consumer, you can use it to filter filenames, so you can for instance consume today's file using the File Language syntax: mydata-\${date:now:yyyyMMdd}.txt. The producers support the CamelOverrideFileName header which takes precedence over any existing CamelFileName header; the CamelOverrideFileName is a header that is used only once, and makes it easier as this avoids to temporary store CamelFileName and have to restore it afterwards.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
delete (consumer)	If true, the file will be deleted after it is processed successfully.	false	boolean
moveFailed (consumer)	Sets the move failure expression based on Simple language. For example, to move files into a .error subdirectory use: .error. Note: When moving the files to the fail location Camel will handle the error and will not pick up the file again.		String
noop (consumer)	If true, the file is not moved or deleted in any way. This option is good for readonly data, or for ETL type requirements. If noop=true, Camel will set idempotent=true as well, to avoid consuming the same files over and over again.	false	boolean

Name	Description	Default	Type
preMove (consumer)	Expression (such as File Language) used to dynamically set the filename when moving it before processing. For example to move in-progress files into the order directory set this value to order.		String
preSort (consumer)	When pre-sort is enabled then the consumer will sort the file and directory names during polling, that was retrieved from the file system. You may want to do this in case you need to operate on the files in a sorted order. The pre-sort is executed before the consumer starts to filter, and accept files to process by Camel. This option is default=false meaning disabled.	false	boolean
recursive (consumer)	If a directory, will look for files in all the sub-directories as well.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
directoryMustExist (consumer advanced))	Similar to the startingDirectoryMustExist option but this applies during polling (after starting the consumer).	false	boolean
exceptionHandler (consumer advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
extendedAttributes (consumer advanced))	To define which file attributes of interest. Like posix:permissions,posix:owner,basic:lastAccessTime, it supports basic wildcard like posix:, basic:lastAccessTime.		String

Name	Description	Default	Type
inProgressRepository (consumer (advanced))	A pluggable in-progress repository <code>org.apache.camel.spi.IdempotentRepository</code> . The in-progress repository is used to account the current in progress files being consumed. By default a memory based repository is used.		<code>IdempotentRepository</code>
localWorkDirectory (consumer (advanced))	When consuming, a local work directory can be used to store the remote file content directly in local files, to avoid loading the content into memory. This is beneficial, if you consume a very big remote file and thus can conserve memory.		<code>String</code>
onCompletionExceptionHandler (consumer (advanced))	To use a custom <code>org.apache.camel.spi.ExceptionHandler</code> to handle any thrown exceptions that happens during the file on completion process where the consumer does either a commit or rollback. The default implementation will log any exception at WARN level and ignore.		<code>ExceptionHandler</code>
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>
probeContentType (consumer (advanced))	Whether to enable probing of the content type. If enable then the consumer uses <code>Files#probeContentType(java.nio.file.Path)</code> to determine the content-type of the file, and store that as a header with key <code>Exchange#FILE_CONTENT_TYPE</code> on the Message.	<code>false</code>	<code>boolean</code>
processStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.component.file.GenericFileProcessStrategy</code> allowing you to implement your own <code>readLock</code> option or similar. Can also be used when special conditions must be met before a file can be consumed, such as a special ready file exists. If this option is set then the <code>readLock</code> option does not apply.		<code>GenericFileProcessStrategy</code>
resumeStrategy (consumer (advanced))	Set a resume strategy for files. This makes it possible to define a strategy for resuming reading files after the last point before stopping the application. See the <code>FileConsumerResumeStrategy</code> for implementation details.		<code>FileConsumerResumeStrategy</code>

Name	Description	Default	Type
startingDirectoryMustExist (consumer (advanced))	Whether the starting directory must exist. Mind that the <code>autoCreate</code> option is default enabled, which means the starting directory is normally auto created if it doesn't exist. You can disable <code>autoCreate</code> and enable this to ensure the starting directory must exist. Will thrown an exception if the directory doesn't exist.	false	boolean
startingDirectoryMustHaveAccess (consumer (advanced))	Whether the starting directory has access permissions. Mind that the <code>startingDirectoryMustExist</code> parameter must be set to true in order to verify that the directory exists. Will thrown an exception if the directory doesn't have read and write permissions.	false	boolean
appendChars (producer)	Used to append characters (text) after writing files. This can for example be used to add new lines or other separators when writing and appending new files or existing files. To specify new-line (slash-n or slash-r) or tab (slash-t) characters then escape with an extra slash, eg slash-slash-n.		String

Name	Description	Default	Type
fileExist (producer)	<p>What to do if a file already exists with the same name. Override, which is the default, replaces the existing file. - Append - adds content to the existing file. - Fail - throws a <code>GenericFileOperationException</code>, indicating that there is already an existing file. - Ignore - silently ignores the problem and does not override the existing file, but assumes everything is okay. - Move - option requires to use the <code>moveExisting</code> option to be configured as well. The option <code>eagerDeleteTargetFile</code> can be used to control what to do if an moving the file, and there exists already an existing file, otherwise causing the move operation to fail. The Move option will move any existing files, before writing the target file. - <code>TryRename</code> is only applicable if <code>tempFileName</code> option is in use. This allows to try renaming the file from the temporary name to the actual name, without doing any exists check. This check may be faster on some file systems and especially FTP servers.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● Override ● Append ● Fail ● Ignore ● Move ● TryRename 	Override	GenericFileExist
flatten (producer)	<p>Flatten is used to flatten the file name path to strip any leading paths, so it's just the file name. This allows you to consume recursively into sub-directories, but when you eg write the files to another directory they will be written in a single directory. Setting this to true on the producer enforces that any file name in <code>CamelFileName</code> header will be stripped for any leading paths.</p>	false	boolean
jailStartingDirectory (producer)	<p>Used for jailing (restricting) writing files to the starting directory (and sub) only. This is enabled by default to not allow Camel to write files to outside directories (to be more secured out of the box). You can turn this off to allow writing files to directories outside the starting directory, such as parent or root folders.</p>	true	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
moveExisting (producer)	Expression (such as File Language) used to compute file name to use when fileExist=Move is configured. To move files into a backup subdirectory just enter backup. This option only supports the following File Language tokens: file:name, file:name.ext, file:name.noext, file:onlyname, file:onlyname.noext, file:ext, and file:parent. Notice the file:parent is not supported by the FTP component, as the FTP component can only move any existing files to a relative directory based on current dir as base.		String
tempFileName (producer)	The same as tempPrefix option but offering a more fine grained control on the naming of the temporary filename as it uses the File Language. The location for tempFileName is relative to the final file location in the option 'fileName', not the target directory in the base uri. For example if option fileName includes a directory prefix: dir/finalFilename then tempFileName is relative to that subdirectory dir.		String
tempPrefix (producer)	This option is used to write the file using a temporary name and then, after the write is complete, rename it to the real name. Can be used to identify files being written and also avoid consumers (not using exclusive read locks) reading in progress files. Is often used by FTP when uploading big files.		String
allowNullBody (producer (advanced))	Used to specify if a null body is allowed during file writing. If set to true then an empty file will be created, when set to false, and attempting to send a null body to the file component, a GenericFileWriteException of 'Cannot write null body to file.' will be thrown. If the fileExist option is set to 'Override', then the file will be truncated, and if set to append the file will remain unchanged.	false	boolean

Name	Description	Default	Type
chmod (producer (advanced))	Specify the file permissions which is sent by the producer, the chmod value must be between 000 and 777; If there is a leading digit like in 0755 we will ignore it.		String
chmodDirectory (producer (advanced))	Specify the directory permissions used when the producer creates missing directories, the chmod value must be between 000 and 777; If there is a leading digit like in 0755 we will ignore it.		String
eagerDeleteTargetFile (producer (advanced))	Whether or not to eagerly delete any existing target file. This option only applies when you use <code>fileExists=Override</code> and the <code>tempFileName</code> option as well. You can use this to disable (set it to false) deleting the target file before the temp file is written. For example you may write big files and want the target file to exist during the temp file is being written. This ensure the target file is only deleted until the very last moment, just before the temp file is being renamed to the target filename. This option is also used to control whether to delete any existing files when <code>fileExist=Move</code> is enabled, and an existing file exists. If this option <code>copyAndDeleteOnRenameFails</code> false, then an exception will be thrown if an existing file existed, if its true, then the existing file is deleted before the move operation.	true	boolean
forceWrites (producer (advanced))	Whether to force syncing writes to the file system. You can turn this off if you do not want this level of guarantee, for example if writing to logs / audit logs etc; this would yield better performance.	true	boolean
keepLastModified (producer (advanced))	Will keep the last modified timestamp from the source file (if any). Will use the <code>Exchange.FILE_LAST_MODIFIED</code> header to located the timestamp. This header can contain either a <code>java.util.Date</code> or long with the timestamp. If the timestamp exists and the option is enabled it will set this timestamp on the written file. Note: This option only applies to the file producer. You cannot use this option with any of the ftp producers.	false	boolean
moveExistingFileStrategy (producer (advanced))	Strategy (Custom Strategy) used to move file with special naming token to use when <code>fileExist=Move</code> is configured. By default, there is an implementation used if no custom strategy is provided.		FileMoveExistingStrategy

Name	Description	Default	Type
autoCreate (advanced)	Automatically create missing directories in the file's pathname. For the file consumer, that means creating the starting directory. For the file producer, it means the directory the files should be written to.	true	boolean
bufferSize (advanced)	Buffer size in bytes used for writing files (or in case of FTP for downloading and uploading files).	131072	int
copyAndDeleteOnRenameFail (advanced)	Whether to fallback and do a copy and delete file, in case the file could not be renamed directly. This option is not available for the FTP component.	true	boolean
renameUsingCopy (advanced)	Perform rename operations using a copy and delete strategy. This is primarily used in environments where the regular rename operation is unreliable (e.g. across different file systems or networks). This option takes precedence over the <code>copyAndDeleteOnRenameFail</code> parameter that will automatically fall back to the copy and delete strategy, but only after additional delays.	false	boolean
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
antExclude (filter)	Ant style filter exclusion. If both <code>antInclude</code> and <code>antExclude</code> are used, <code>antExclude</code> takes precedence over <code>antInclude</code> . Multiple exclusions may be specified in comma-delimited format.		String
antFilterCaseSensitive (filter)	Sets case sensitive flag on ant filter.	true	boolean
antInclude (filter)	Ant style filter inclusion. Multiple inclusions may be specified in comma-delimited format.		String
eagerMaxMessagesPerPoll (filter)	Allows for controlling whether the limit from <code>maxMessagesPerPoll</code> is eager or not. If eager then the limit is during the scanning of files. Where as false would scan all files, and then perform sorting. Setting this option to false allows for sorting all files first, and then limit the poll. Mind that this requires a higher memory usage as all file details are in memory to perform the sorting.	true	boolean

Name	Description	Default	Type
exclude (filter)	Is used to exclude files, if filename matches the regex pattern (matching is case in-sensitive). Notice if you use symbols such as plus sign and others you would need to configure this using the RAW() syntax if configuring this as an endpoint uri. See more details at configuring endpoint uris.		String
excludeExt (filter)	Is used to exclude files matching file extension name (case insensitive). For example to exclude bak files, then use excludeExt=bak. Multiple extensions can be separated by comma, for example to exclude bak and dat files, use excludeExt=bak,dat. Note that the file extension includes all parts, for example having a file named mydata.tar.gz will have extension as tar.gz. For more flexibility then use the include/exclude options.		String
filter (filter)	Pluggable filter as a org.apache.camel.component.file.GenericFileFilter class. Will skip files if filter returns false in its accept() method.		GenericFileFilter
filterDirectory (filter)	Filters the directory based on Simple language. For example to filter on current date, you can use a simple date pattern such as <code>\$\$\{date:now:yyyMMdd}</code> .		String
filterFile (filter)	Filters the file based on Simple language. For example to filter on file size, you can use <code>\$\$\{file:size} 5000</code> .		String
idempotent (filter)	Option to use the Idempotent Consumer EIP pattern to let Camel skip already processed files. Will by default use a memory based LRUcache that holds 1000 entries. If noop=true then idempotent will be enabled as well to avoid consuming the same files over and over again.	false	Boolean
idempotentKey (filter)	To use a custom idempotent key. By default the absolute path of the file is used. You can use the File Language, for example to use the file name and file size, you can do: <code>idempotentKey=\$\{file:name}-\$\{file:size}</code> .		String
idempotentRepository (filter)	A pluggable repository org.apache.camel.spi.IdempotentRepository which by default use MemoryIdempotentRepository if none is specified and idempotent is true.		IdempotentRepository

Name	Description	Default	Type
include (filter)	Is used to include files, if filename matches the regex pattern (matching is case in-sensitive). Notice if you use symbols such as plus sign and others you would need to configure this using the RAW() syntax if configuring this as an endpoint uri. See more details at configuring endpoint uris.		String
includeExt (filter)	Is used to include files matching file extension name (case insensitive). For example to include txt files, then use includeExt=txt. Multiple extensions can be separated by comma, for example to include txt and xml files, use includeExt=txt,xml. Note that the file extension includes all parts, for example having a file named mydata.tar.gz will have extension as tar.gz. For more flexibility then use the include/exclude options.		String
maxDepth (filter)	The maximum depth to traverse when recursively processing a directory.	2147483647	int
maxMessagesPerPoll (filter)	To define a maximum messages to gather per poll. By default no maximum is set. Can be used to set a limit of e.g. 1000 to avoid when starting up the server that there are thousands of files. Set a value of 0 or negative to disabled it. Notice: If this option is in use then the File and FTP components will limit before any sorting. For example if you have 100000 files and use maxMessagesPerPoll=500, then only the first 500 files will be picked up, and then sorted. You can use the eagerMaxMessagesPerPoll option and set this to false to allow to scan all files first and then sort afterwards.		int
minDepth (filter)	The minimum depth to start processing when recursively processing a directory. Using minDepth=1 means the base directory. Using minDepth=2 means the first sub directory.		int
move (filter)	Expression (such as Simple Language) used to dynamically set the filename when moving it after processing. To move files into a .done subdirectory just enter .done.		String
exclusiveReadLockStrategy (lock)	Pluggable read-lock as a org.apache.camel.component.file.GenericFileExclusiveReadLockStrategy implementation.		GenericFileExclusiveReadLockStrategy
readLock (lock)	Used by consumer, to only poll the files if it has	none	String

Name	Description	Default	Type
	<p>exclusive read-lock on the file (i.e. the file is not in-progress or being written). Camel will wait until the file lock is granted. This option provides the build in strategies:</p> <ul style="list-style-type: none"> - none - No read lock is in use - markerFile - Camel creates a marker file (fileName.camellLock) and then holds a lock on it. This option is not available for the FTP component - changed - Changed is using file length/modification timestamp to detect whether the file is currently being copied or not. Will at least use 1 sec to determine this, so this option cannot consume files as fast as the others, but can be more reliable as the JDK IO API cannot always determine whether a file is currently being used by another process. The option readLockCheckInterval can be used to set the check frequency. - fileLock - is for using java.nio.channels.FileLock. This option is not avail for Windows OS and the FTP component. This approach should be avoided when accessing a remote file system via a mount/share unless that file system supports distributed file locks. - rename - rename is for using a try to rename the file as a test if we can get exclusive read-lock. - idempotent - (only for file component) idempotent is for using a idempotentRepository as the read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. - idempotent-changed - (only for file component) idempotent-changed is for using a idempotentRepository and changed as the combined read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. - idempotent-rename - (only for file component) idempotent-rename is for using a idempotentRepository and rename as the combined read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. <p>Notice: The various read locks is not all suited to work in clustered mode, where concurrent consumers on different nodes is competing for the same files on a shared file system. The markerFile using a close to atomic operation to create the empty marker file, but its not guaranteed to work in a cluster. The fileLock may work better but then the file system need to support distributed file locks, and so on. Using the idempotent read lock can support clustering if the idempotent repository supports clustering, such as Hazelcast Component or Infinispan.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● none ● markerFile ● fileLock 		

Name	<ul style="list-style-type: none"> ● rename Description <ul style="list-style-type: none"> ● changed 	Default	Type
	<ul style="list-style-type: none"> ● idempotent ● idempotent-changed ● idempotent-rename 		
readLockCheckInterval (lock)	Interval in millis for the read-lock, if supported by the read lock. This interval is used for sleeping between attempts to acquire the read lock. For example when using the changed read lock, you can set a higher interval period to cater for slow writes. The default of 1 sec. may be too fast if the producer is very slow writing the file. Notice: For FTP the default readLockCheckInterval is 5000. The readLockTimeout value must be higher than readLockCheckInterval, but a rule of thumb is to have a timeout that is at least 2 or more times higher than the readLockCheckInterval. This is needed to ensure that ample time is allowed for the read lock process to try to grab the lock before the timeout was hit.	1000	long
readLockDeleteOrphanLockFiles (lock)	Whether or not read lock with marker files should upon startup delete any orphan read lock files, which may have been left on the file system, if Camel was not properly shutdown (such as a JVM crash). If turning this option to false then any orphaned lock file will cause Camel to not attempt to pickup that file, this could also be due another node is concurrently reading files from the same shared directory.	true	boolean
readLockIdempotentReleaseAsync (lock)	Whether the delayed release task should be synchronous or asynchronous. See more details at the readLockIdempotentReleaseDelay option.	false	boolean
readLockIdempotentReleaseAsyncPoolSize (lock)	The number of threads in the scheduled thread pool when using asynchronous release tasks. Using a default of 1 core threads should be sufficient in almost all use-cases, only set this to a higher value if either updating the idempotent repository is slow, or there are a lot of files to process. This option is not in-use if you use a shared thread pool by configuring the readLockIdempotentReleaseExecutorService option. See more details at the readLockIdempotentReleaseDelay option.		int

Name	Description	Default	Type
readLockIdempotentReleaseDelay (lock)	Whether to delay the release task for a period of millis. This can be used to delay the release tasks to expand the window when a file is regarded as read-locked, in an active/active cluster scenario with a shared idempotent repository, to ensure other nodes cannot potentially scan and acquire the same file, due to race-conditions. By expanding the time-window of the release tasks helps prevents these situations. Note delaying is only needed if you have configured <code>readLockRemoveOnCommit</code> to true.		int
readLockIdempotentReleaseExecutorService (lock)	To use a custom and shared thread pool for asynchronous release tasks. See more details at the <code>readLockIdempotentReleaseDelay</code> option.		ScheduledExecutorService
readLockLoggingLevel (lock)	Logging level used when a read lock could not be acquired. By default a DEBUG is logged. You can change this level, for example to OFF to not have any logging. This option is only applicable for readLock of types: <code>changed</code> , <code>fileLock</code> , <code>idempotent</code> , <code>idempotent-changed</code> , <code>idempotent-rename</code> , <code>rename</code> . Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	DEBUG	LogLevel
readLockMarkerFile (lock)	Whether to use marker file with the <code>changed</code> , <code>rename</code> , or <code>exclusive read lock</code> types. By default a marker file is used as well to guard against other processes picking up the same files. This behavior can be turned off by setting this option to false. For example if you do not want to write marker files to the file systems by the Camel application.	true	boolean
readLockMinAge (lock)	This option is applied only for <code>readLock=changed</code> . It allows to specify a minimum age the file must be before attempting to acquire the read lock. For example use <code>readLockMinAge=300s</code> to require the file is at least 5 minutes old. This can speedup the <code>changed read lock</code> as it will only attempt to acquire files which are at least that given age.	0	long

Name	Description	Default	Type
readLockMinLength (lock)	This option is applied only for readLock=changed. It allows you to configure a minimum file length. By default Camel expects the file to contain data, and thus the default value is 1. You can set this option to zero, to allow consuming zero-length files.	1	long
readLockRemoveOnCommit (lock)	This option is applied only for readLock=idempotent. It allows to specify whether to remove the file name entry from the idempotent repository when processing the file is succeeded and a commit happens. By default the file is not removed which ensures that any race-condition do not occur so another active node may attempt to grab the file. Instead the idempotent repository may support eviction strategies that you can configure to evict the file name entry after X minutes - this ensures no problems with race conditions. See more details at the readLockIdempotentReleaseDelay option.	false	boolean
readLockRemoveOnRollback (lock)	This option is applied only for readLock=idempotent. It allows to specify whether to remove the file name entry from the idempotent repository when processing the file failed and a rollback happens. If this option is false, then the file name entry is confirmed (as if the file did a commit).	true	boolean
readLockTimeout (lock)	Optional timeout in millis for the read-lock, if supported by the read-lock. If the read-lock could not be granted and the timeout triggered, then Camel will skip the file. At next poll Camel, will try the file again, and this time maybe the read-lock could be granted. Use a value of 0 or lower to indicate forever. Currently fileLock, changed and rename support the timeout. Notice: For FTP the default readLockTimeout value is 20000 instead of 10000. The readLockTimeout value must be higher than readLockCheckInterval, but a rule of thumb is to have a timeout that is at least 2 or more times higher than the readLockCheckInterval. This is needed to ensure that ample time is allowed for the read lock process to try to grab the lock before the timeout was hit.	10000	long
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultipler should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultipler should kick-in.		int

Name	Description	Default	Type
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object

Name	Description	Default	Type
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> • NANoseconds • MICROseconds • MILLIseconds • SECONDS • MINUTES • HOURS • DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
shuffle (sort)	To shuffle the list of files (sort in random order).	false	boolean
sortBy (sort)	Built-in sort by using the File Language. Supports nested sorts, so you can have a sort by file name and as a 2nd group sort by modified date.		String
sorter (sort)	Pluggable sorter as a java.util.Comparator class.		Comparator

**NOTE****Default behavior for file producer**

By default it will override any existing file, if one exist with the same name.

21.5. MOVE AND DELETE OPERATIONS

Any move or delete operations is executed after (post command) the routing has completed; so during processing of the **Exchange** the file is still located in the inbox folder.

Lets illustrate this with an example:

```
from("file://inbox?move=.done").to("bean:handleOrder");
```

When a file is dropped in the **inbox** folder, the file consumer notices this and creates a new **FileExchange** that is routed to the **handleOrder** bean. The bean then processes the **File** object. At this point in time the file is still located in the **inbox** folder. After the bean completes, and thus the route is completed, the file consumer will perform the move operation and move the file to the **.done** sub-folder.

The **move** and the **preMove** options are considered as a directory name (though if you use an expression such as [File Language](#), or [Simple](#) then the result of the expression evaluation is the file name to be used. For example, if you set:

```
move=../backup/copy-of-${file:name}
```

then that's using the [File](#) language which we use return the file name to be used), which can be either relative or absolute. If relative, the directory is created as a sub-folder from within the folder where the file was consumed.

By default, Camel will move consumed files to the **.camel** sub-folder relative to the directory where the file was consumed.

If you want to delete the file after processing, the route should be:

```
from("file://inbox?delete=true").to("bean:handleOrder");
```

We have introduced a **pre** move operation to move files **before** they are processed. This allows you to mark which files have been scanned as they are moved to this sub folder before being processed.

```
from("file://inbox?preMove=inprogress").to("bean:handleOrder");
```

You can combine the **pre** move and the regular move:

```
from("file://inbox?preMove=inprogress&move=.done").to("bean:handleOrder");
```

So in this situation, the file is in the **inprogress** folder when being processed and after it's processed, it's moved to the **.done** folder.

21.6. FINE GRAINED CONTROL OVER MOVE AND PREMOVE OPTION

The **move** and **preMove** options are Expression-based, so we have the full power of the [File Language](#) to do advanced configuration of the directory and name pattern.

Camel will, in fact, internally convert the directory name you enter into a [File Language](#) expression. So when we enter **move=.done** Camel will convert this into: **\${file:parent}.done/\${file:onlyname}**. This is only done if Camel detects that you have not provided a `$$` in the option value yourself. So when you enter a `$$` Camel will **not** convert it and thus you have the full power.

So if we want to move the file into a backup folder with today's date as the pattern, we can do:

```
move=backup/${date:now:yyyyMMdd}/${file:name}
```

21.7. ABOUT MOVEFAILED

The **moveFailed** option allows you to move files that **could not** be processed successfully to another location such as an error folder of your choice. For example to move the files in an error folder with a

timestamp you can use `moveFailed=/error/${file:name.noext}-${date:now:yyyyMMddHHmmssSSS}.${file:ext}`.

See more examples at

21.8. MESSAGE HEADERS

The following headers are supported by this component:

21.8.1. File producer only

Header	Description
CamelFileName	Specifies the name of the file to write (relative to the endpoint directory). This name can be a String ; a String with a File Language or Simple language expression; or an Expression object. If it's null then Camel will auto-generate a filename based on the message unique ID.
CamelFileNameProduced	The actual absolute filepath (path + name) for the output file that was written. This header is set by Camel and its purpose is providing end-users with the name of the file that was written.
CamelOverruleFileName	Is used for overruling CamelFileName header and use the value instead (but only once, as the producer will remove this header after writing the file). The value can be only be a String. Notice that if the option fileName has been configured, then this is still being evaluated.

21.8.2. File consumer only

Header	Description
CamelFileName	Name of the consumed file as a relative file path with offset from the starting directory configured on the endpoint.
CamelFileNameOnly	Only the file name (the name with no leading paths).
CamelFileAbsolute	A boolean option specifying whether the consumed file denotes an absolute path or not. Should normally be false for relative paths. Absolute paths should normally not be used but we added to the move option to allow moving files to absolute paths. But can be used elsewhere as well.
CamelFileAbsolutePath	The absolute path to the file. For relative files this path holds the relative path instead.
CamelFilePath	The file path. For relative files this is the starting directory + the relative filename. For absolute files this is the absolute path.
CamelFileRelativePath	The relative path.

Header	Description
CamelFileParent	The parent path.
CamelFileLength	A long value containing the file size.
CamelFileLastModified	A Long value containing the last modified timestamp of the file.

21.9. BATCH CONSUMER

This component implements the Batch Consumer.

21.10. EXCHANGE PROPERTIES, FILE CONSUMER ONLY

As the file consumer implements the **BatchConsumer** it supports batching the files it polls. By batching we mean that Camel will add the following additional properties to the Exchange, so you know the number of files polled, the current index, and whether the batch is already completed.

Property	Description
CamelBatchSize	The total number of files that was polled in this batch.
CamelBatchIndex	The current index of the batch. Starts from 0.
CamelBatchComplete	A boolean value indicating the last Exchange in the batch. Is only true for the last entry.

This allows you for instance to know how many files exist in this batch and for instance let the `Aggregator2` aggregate this number of files.

21.11. USING CHARSET

The `charset` option allows for configuring an encoding of the files on both the consumer and producer endpoints. For example if you read utf-8 files, and want to convert the files to iso-8859-1, you can do:

```
from("file:inbox?charset=utf-8")
  .to("file:outbox?charset=iso-8859-1")
```

You can also use the **convertBodyTo** in the route. In the example below we have still input files in utf-8 format, but we want to convert the file content to a byte array in iso-8859-1 format. And then let a bean process the data. Before writing the content to the outbox folder using the current charset.

```
from("file:inbox?charset=utf-8")
  .convertBodyTo(byte[].class, "iso-8859-1")
  .to("bean:myBean")
  .to("file:outbox");
```

If you omit the charset on the consumer endpoint, then Camel does not know the charset of the file, and would by default use "UTF-8". However you can configure a JVM system property to override and use a different default encoding with the key **org.apache.camel.default.charset**.

In the example below this could be a problem if the files is not in UTF-8 encoding, which would be the default encoding for read the files.

In this example when writing the files, the content has already been converted to a byte array, and thus would write the content directly as is (without any further encodings).

```
from("file:inbox")
  .convertBodyTo(byte[].class, "iso-8859-1")
  .to("bean:myBean")
  .to("file:outbox");
```

You can also override and control the encoding dynamic when writing files, by setting a property on the exchange with the key **Exchange.CHARSET_NAME**. For example in the route below we set the property with a value from a message header.

```
from("file:inbox")
  .convertBodyTo(byte[].class, "iso-8859-1")
  .to("bean:myBean")
  .setProperty(Exchange.CHARSET_NAME, header("someCharsetHeader"))
  .to("file:outbox");
```

We suggest to keep things simpler, so if you pickup files with the same encoding, and want to write the files in a specific encoding, then favor to use the **charset** option on the endpoints.

Notice that if you have explicit configured a **charset** option on the endpoint, then that configuration is used, regardless of the **Exchange.CHARSET_NAME** property.

If you have some issues then you can enable DEBUG logging on **org.apache.camel.component.file**, and Camel logs when it reads/write a file using a specific charset.

For example the route below will log the following:

```
from("file:inbox?charset=utf-8")
  .to("file:outbox?charset=iso-8859-1")
```

And the logs:

```
DEBUG GenericFileConverter      - Read file /Users/davsclaus/workspace/camel/camel-core/target/charset/input/input.txt with charset utf-8
DEBUG FileOperations            - Using Reader to write file: target/charset/output.txt with charset: iso-8859-1
```

21.12. COMMON GOTCHAS WITH FOLDER AND FILENAMES

When Camel is producing files (writing files) there are a few gotchas affecting how to set a filename of your choice. By default, Camel will use the message ID as the filename, and since the message ID is normally a unique generated ID, you will end up with filenames such as: **ID-MACHINENAME-2443-1211718892437-1-0**. If such a filename is not desired, then you must provide a filename in the **CamelFileName** message header. The constant, **Exchange.FILE_NAME**, can also be used.

The sample code below produces files using the message ID as the filename:

```
from("direct:report").to("file:target/reports");
```

To use **report.txt** as the filename you have to do:

```
from("direct:report").setHeader(Exchange.FILE_NAME, constant("report.txt")).to("file:target/reports");
```

- the same as above, but with **CamelFileName**:

```
from("direct:report").setHeader("CamelFileName", constant("report.txt")).to("file:target/reports");
```

And a syntax where we set the filename on the endpoint with the **fileName** URI option.

```
from("direct:report").to("file:target/reports/?fileName=report.txt");
```

21.13. FILENAME EXPRESSION

Filename can be set either using the **expression** option or as a string-based [File](#) language expression in the **CamelFileName** header. See the [File](#) language for syntax and samples.

21.14. CONSUMING FILES FROM FOLDERS WHERE OTHERS DROP FILES DIRECTLY

Beware if you consume files from a folder where other applications write files to directly. Take a look at the different `readLock` options to see what suits your use cases. The best approach is however to write to another folder and after the write move the file in the drop folder. However if you write files directly to the drop folder then the option `changed` could better detect whether a file is currently being written/copied as it uses a file changed algorithm to see whether the file size / modification changes over a period of time. The other `readLock` options rely on Java File API that sadly is not always very good at detecting this. You may also want to look at the `doneFileName` option, which uses a marker file (done file) to signal when a file is done and ready to be consumed.

21.15. USING DONE FILES

See also section *writing done files* below.

If you want only to consume files when a done file exists, then you can use the **doneFileName** option on the endpoint.

```
from("file:bar?doneFileName=done");
```

Will only consume files from the bar folder, if a done *file* exists in the same directory as the target files. Camel will automatically delete the *done file* when it's done consuming the files. Camel does not delete automatically the *done file* if **noop=true** is configured.

However it is more common to have one *done file* per target file. This means there is a 1:1 correlation. To do this you must use dynamic placeholders in the **doneFileName** option. Currently Camel supports the following two dynamic tokens: **file:name** and **file:name.noext** which must be enclosed in `${ }`. The consumer only supports the static part of the *done file* name as either prefix or suffix (not both).

```
from("file:bar?doneFileName=${file:name}.done");
```

In this example only files will be polled if there exists a done file with the name *file name.done*. For example

- **hello.txt** - is the file to be consumed
- **hello.txt.done** - is the associated done file

You can also use a prefix for the done file, such as:

```
from("file:bar?doneFileName=ready-${file:name}");
```

- **hello.txt** - is the file to be consumed
- **ready-hello.txt** - is the associated done file

21.16. WRITING DONE FILES

After you have written a file you may want to write an additional *donefile* as a kind of marker, to indicate to others that the file is finished and has been written. To do that you can use the **doneFileName** option on the file producer endpoint.

```
.to("file:bar?doneFileName=done");
```

Will simply create a file named **done** in the same directory as the target file.

However it is more common to have one done file per target file. This means there is a 1:1 correlation. To do this you must use dynamic placeholders in the **doneFileName** option. Currently Camel supports the following two dynamic tokens: **file:name** and **file:name.noext** which must be enclosed in $\${}$.

```
.to("file:bar?doneFileName=done-${file:name}");
```

Will for example create a file named **done-foo.txt** if the target file was **foo.txt** in the same directory as the target file.

```
.to("file:bar?doneFileName=${file:name}.done");
```

Will for example create a file named **foo.txt.done** if the target file was **foo.txt** in the same directory as the target file.

```
.to("file:bar?doneFileName=${file:name.noext}.done");
```

Will for example create a file named **foo.done** if the target file was **foo.txt** in the same directory as the target file.

21.17. SAMPLES

21.17.1. Read from a directory and write to another directory

```
from("file://inputdir/?delete=true").to("file://outputdir")
```

21.17.2. Read from a directory and write to another directory using an overrule dynamic name

```
from("file://inputdir/?delete=true").to("file://outputdir?overruleFile=copy-of-${file:name}")
```

Listen on a directory and create a message for each file dropped there. Copy the contents to the **outputdir** and delete the file in the **inputdir**.

21.17.3. Reading recursively from a directory and writing to another

```
from("file://inputdir/?recursive=true&delete=true").to("file://outputdir")
```

Listen on a directory and create a message for each file dropped there. Copy the contents to the **outputdir** and delete the file in the **inputdir**. Will scan recursively into sub-directories. Will lay out the files in the same directory structure in the **outputdir** as the **inputdir**, including any sub-directories.

```
inputdir/foo.txt
inputdir/sub/bar.txt
```

Will result in the following output layout:

```
outputdir/foo.txt
outputdir/sub/bar.txt
```

21.18. USING FLATTEN

If you want to store the files in the outputdir directory in the same directory, disregarding the source directory layout (e.g. to flatten out the path), you just add the **flatten=true** option on the file producer side:

```
from("file://inputdir/?recursive=true&delete=true").to("file://outputdir?flatten=true")
```

Will result in the following output layout:

```
outputdir/foo.txt
outputdir/bar.txt
```

21.19. READING FROM A DIRECTORY AND THE DEFAULT MOVE OPERATION

Camel will by default move any processed file into a **.camel** subdirectory in the directory the file was consumed from.

```
from("file://inputdir/?recursive=true&delete=true").to("file://outputdir")
```

Affects the layout as follows:
before

```
inputdir/foo.txt
inputdir/sub/bar.txt
```

after

```
inputdir/.camel/foo.txt
inputdir/sub/.camel/bar.txt
outputdir/foo.txt
outputdir/sub/bar.txt
```

21.20. READ FROM A DIRECTORY AND PROCESS THE MESSAGE IN JAVA

```
from("file://inputdir/").process(new Processor() {
    public void process(Exchange exchange) throws Exception {
        Object body = exchange.getIn().getBody();
        // do some business logic with the input body
    }
});
```

The body will be a **File** object that points to the file that was just dropped into the **inputdir** directory.

21.21. WRITING TO FILES

Camel is of course also able to write files, i.e. produce files. In the sample below we receive some reports on the SEDA queue that we process before they are being written to a directory.

21.21.1. Write to subdirectory using `Exchange.FILE_NAME`

Using a single route, it is possible to write a file to any number of subdirectories. If you have a route setup as such:

```
<route>
  <from uri="bean:myBean"/>
  <to uri="file:/rootDirectory"/>
</route>
```

You can have **myBean** set the header **Exchange.FILE_NAME** to values such as:

```
Exchange.FILE_NAME = hello.txt => /rootDirectory/hello.txt
Exchange.FILE_NAME = foo/bye.txt => /rootDirectory/foo/bye.txt
```

This allows you to have a single route to write files to multiple destinations.

21.21.2. Writing file through the temporary directory relative to the final destination

Sometime you need to temporarily write the files to some directory relative to the destination directory. Such situation usually happens when some external process with limited filtering capabilities is reading from the directory you are writing to. In the example below files will be written to the **/var/myapp/filesInProgress** directory and after data transfer is done, they will be atomically moved to the ``/var/myapp/finalDirectory`` directory.

```
from("direct:start").
  to("file:///var/myapp/finalDirectory?tempPrefix=./filesInProgress/");
```

21.22. USING EXPRESSION FOR FILENAMES

In this sample we want to move consumed files to a backup folder using today's date as a sub-folder name:

```
from("file://inbox?move=backup/${date:now:yyyyMMdd}/${file:name}").to("...");
```

See [File language](#) for more samples.

21.23. AVOIDING READING THE SAME FILE MORE THAN ONCE (IDEMPOTENT CONSUMER)

Camel supports Idempotent Consumer directly within the component so it will skip already processed files. This feature can be enabled by setting the **idempotent=true** option.

```
from("file://inbox?idempotent=true").to("...");
```

Camel uses the absolute file name as the idempotent key, to detect duplicate files. You can customize this key by using an expression in the idempotentKey option. For example to use both the name and the file size as the key

```
<route>
  <from uri="file://inbox?idempotent=true&idempotentKey=${file:name}-${file:size}"/>
  <to uri="bean:processInbox"/>
</route>
```

By default Camel uses a in memory based store for keeping track of consumed files, it uses a least recently used cache holding up to 1000 entries. You can plugin your own implementation of this store by using the **idempotentRepository** option using the # sign in the value to indicate it's a referring to a bean in the Registry with the specified **id**.

```
<!-- define our store as a plain spring bean -->
<bean id="myStore" class="com.mycompany.MyIdempotentStore"/>

<route>
  <from uri="file://inbox?idempotent=true&idempotentRepository=#myStore"/>
  <to uri="bean:processInbox"/>
</route>
```

Camel will log at **DEBUG** level if it skips a file because it has been consumed before:

```
DEBUG FileConsumer is idempotent and the file has been consumed before. Will skip this file:
target\idempotent\report.txt
```

21.24. USING A FILE BASED IDEMPOTENT REPOSITORY

In this section we will use the file based idempotent repository **org.apache.camel.processor.idempotent.FileIdempotentRepository** instead of the in-memory based that is used as default.

This repository uses a 1st level cache to avoid reading the file repository. It will only use the file repository to store the content of the 1st level cache. Thereby the repository can survive server restarts. It will load the content of the file into the 1st level cache upon startup. The file structure is very simple as

it stores the key in separate lines in the file. By default, the file store has a size limit of 1mb. When the file grows larger Camel will truncate the file store, rebuilding the content by flushing the 1st level cache into a fresh empty file.

We configure our repository using Spring XML creating our file idempotent repository and define our file consumer to use our repository with the **idempotentRepository** using # sign to indicate Registry lookup:

21.25. USING A JPA BASED IDEMPOTENT REPOSITORY

In this section we will use the JPA based idempotent repository instead of the in-memory based that is used as default.

First we need a persistence-unit in **META-INF/persistence.xml** where we need to use the class **org.apache.camel.processor.idempotent.jpa.MessageProcessed** as model.

```
<persistence-unit name="idempotentDb" transaction-type="RESOURCE_LOCAL">
  <class>org.apache.camel.processor.idempotent.jpa.MessageProcessed</class>

  <properties>
    <property name="openjpa.ConnectionURL" value="jdbc:derby:target/idempotentTest;create=true"/>
    <property name="openjpa.ConnectionDriverName"
value="org.apache.derby.jdbc.EmbeddedDriver"/>
    <property name="openjpa.jdbc.SynchronizeMappings" value="buildSchema"/>
    <property name="openjpa.Log" value="DefaultLevel=WARN, Tool=INFO"/>
    <property name="openjpa.Multithreaded" value="true"/>
  </properties>
</persistence-unit>
```

Next, we can create our JPA idempotent repository in the spring XML file as well:

```
<!-- we define our jpa based idempotent repository we want to use in the file consumer -->
<bean id="jpaStore" class="org.apache.camel.processor.idempotent.jpa.JpaMessageIdRepository">
  <!-- Here we refer to the entityManagerFactory -->
  <constructor-arg index="0" ref="entityManagerFactory"/>
  <!-- This 2nd parameter is the name (= a category name).
      You can have different repositories with different names -->
  <constructor-arg index="1" value="FileConsumer"/>
</bean>
```

And yes then we just need to refer to the **jpaStore** bean in the file consumer endpoint using the **idempotentRepository** using the # syntax option:

```
<route>
  <from uri="file://inbox?idempotent=true&idempotentRepository=#jpaStore"/>
  <to uri="bean:processInbox"/>
</route>
```

21.26. FILTER USING

ORG.APACHE.CAMEL.COMPONENT.FILE.GENERICFILEFILTER

Camel supports pluggable filtering strategies. You can then configure the endpoint with such a filter to skip certain files being processed.

In the sample we have built our own filter that skips files starting with **skip** in the filename:

And then we can configure our route using the **filter** attribute to reference our filter (using # notation) that we have defined in the spring XML file:

```
<!-- define our filter as a plain spring bean -->
<bean id="myFilter" class="com.mycompany.MyFileFilter"/>

<route>
  <from uri="file://inbox?filter=#myFilter"/>
  <to uri="bean:processInbox"/>
</route>
```

21.27. FILTERING USING ANT PATH MATCHER

The ANT path matcher is based on [AntPathMatcher](#).

The file paths is matched with the following rules:

- **?** matches one character
- ***** matches zero or more characters
- ****** matches zero or more directories in a path

The **antInclude** and **antExclude** options make it easy to specify ANT style include/exclude without having to define the filter. See the URI options above for more information.

The sample below demonstrates how to use it:

21.27.1. Sorting using Comparator

Camel supports pluggable sorting strategies. This strategy it to use the build in **java.util.Comparator** in Java. You can then configure the endpoint with such a comparator and have Camel sort the files before being processed.

In the sample we have built our own comparator that just sorts by file name:

And then we can configure our route using the **sorter** option to reference to our sorter (**mySorter**) we have defined in the spring XML file:

```
<!-- define our sorter as a plain spring bean -->
<bean id="mySorter" class="com.mycompany.MyFileSorter"/>

<route>
  <from uri="file://inbox?sorter=#mySorter"/>
  <to uri="bean:processInbox"/>
</route>
```



NOTE

URI options can reference beans using the # syntax

In the Spring DSL route above notice that we can refer to beans in the Registry by prefixing the id with #. So writing **sorter=#mySorter**, will instruct Camel to go look in the Registry for a bean with the ID, **mySorter**.

21.27.2. Sorting using sortBy

Camel supports pluggable sorting strategies. This strategy it to use the [File](#) language to configure the sorting. The **sortBy** option is configured as follows:

```
sortBy=group 1;group 2;group 3;...
```

Where each group is separated with semi colon. In the simple situations you just use one group, so a simple example could be:

```
sortBy=file:name
```

This will sort by file name, you can reverse the order by prefixing **reverse:** to the group, so the sorting is now Z..A:

```
sortBy=reverse:file:name
```

As we have the full power of [File](#) language we can use some of the other parameters, so if we want to sort by file size we do:

```
sortBy=file:length
```

You can configure to ignore the case, using **ignoreCase:** for string comparison, so if you want to use file name sorting but to ignore the case then we do:

```
sortBy=ignoreCase:file:name
```

You can combine ignore case and reverse, however reverse must be specified first:

```
sortBy=reverse:ignoreCase:file:name
```

In the sample below we want to sort by last modified file, so we do:

```
sortBy=file:modified
```

And then we want to group by name as a 2nd option so files with same modification is sorted by name:

```
sortBy=file:modified;file:name
```

Now there is an issue here, can you spot it? Well the modified timestamp of the file is too fine as it will be in milliseconds, but what if we want to sort by date only and then subgroup by name?

Well as we have the true power of [File](#) language we can use its date command that supports patterns. So this can be solved as:

```
sortBy=date:file:yyyyMMdd;file:name
```

Yeah, that is pretty powerful, oh by the way you can also use reverse per group, so we could reverse the file names:

```
sortBy=date:file:yyyyMMdd;reverse:file:name
```

21.28. USING GENERICFILEPROCESSSTRATEGY

The option **processStrategy** can be used to use a custom **GenericFileProcessStrategy** that allows you to implement your own *begin*, *commit* and *rollback* logic.

For instance lets assume a system writes a file in a folder you should consume. But you should not start consuming the file before another *ready* file has been written as well.

So by implementing our own **GenericFileProcessStrategy** we can implement this as:

- In the **begin()** method we can test whether the special *ready* file exists. The begin method returns a **boolean** to indicate if we can consume the file or not.
- In the **abort()** method special logic can be executed in case the **begin** operation returned **false**, for example to cleanup resources etc.
- in the **commit()** method we can move the actual file and also delete the *ready* file.

21.29. USING FILTER

The **filter** option allows you to implement a custom filter in Java code by implementing the **org.apache.camel.component.file.GenericFileFilter** interface. This interface has an **accept** method that returns a boolean. Return **true** to include the file, and **false** to skip the file. There is a **isDirectory** method on **GenericFile** whether the file is a directory. This allows you to filter unwanted directories, to avoid traversing down unwanted directories.

For example to skip any directories which starts with **"skip"** in the name, can be implemented as follows:

21.30. USING BRIDGEERRORHANDLER

If you want to use the Camel Error Handler to deal with any exception occurring in the file consumer, then you can enable the **bridgeErrorHandler** option as shown below:

```
// to handle any IOException being thrown
onException(IOException.class)
    .handled(true)
    .log("IOException occurred due: ${exception.message}")
    .transform().simple("Error ${exception.message}")
    .to("mock:error");

// this is the file route that pickup files, notice how we bridge the consumer to use the Camel routing
// error handler
// the exclusiveReadLockStrategy is only configured because this is from an unit test, so we use that
// to simulate exceptions
from("file:target/nospace?bridgeErrorHandler=true")
    .convertBodyTo(String.class)
    .to("mock:result");
```

So all you have to do is to enable this option, and the error handler in the route will take it from there.



IMPORTANT

When using **bridgeErrorHandler**

When using `bridgeErrorHandler`, then interceptors, `OnCompletions` does **not** apply. The Exchange is processed directly by the Camel Error Handler, and does not allow prior actions such as interceptors, `onCompletion` to take action.

21.31. DEBUG LOGGING

This component has log level **TRACE** that can be helpful if you have problems.

21.32. SPRING BOOT AUTO-CONFIGURATION

When using file with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-file-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 11 options, which are listed below.

Name	Description	Default	Type
<code>camel.cluster.file.acquire-lock-delay</code>	The time to wait before starting to try to acquire lock.		String
<code>camel.cluster.file.acquire-lock-interval</code>	The time to wait between attempts to try to acquire lock.		String
<code>camel.cluster.file.attributes</code>	Custom service attributes.		Map
<code>camel.cluster.file.enabled</code>	Sets if the file cluster service should be enabled or not, default is false.	false	Boolean
<code>camel.cluster.file.id</code>	Cluster Service ID.		String
<code>camel.cluster.file.order</code>	Service lookup order/priority.		Integer
<code>camel.cluster.file.root</code>	The root path.		String

Name	Description	Default	Type
camel.component.file.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.file.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.file.enabled	Whether to enable auto configuration of the file component. This is enabled by default.		Boolean
camel.component.file.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 22. FTP

Both producer and consumer are supported

This component provides access to remote file systems over the FTP and SFTP protocols.

When consuming from remote FTP server make sure you read the section titled *Default when consuming files* further below for details related to consuming files.

Absolute path is **not** supported. Camel translates absolute path to relative by trimming all leading slashes from **directoryname**. There'll be WARN message printed in the logs.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-ftp</artifactId>
  <version>3.14.5.redhat-00018</version>See the documentation of the Apache Commons
  <!-- use the same version as your Camel core version -->
</dependency>
```

22.1. URI FORMAT

```
ftp://[username@]hostname[:port]/directoryname[?options]
sftp://[username@]hostname[:port]/directoryname[?options]
ftps://[username@]hostname[:port]/directoryname[?options]
```

Where **directoryname** represents the underlying directory. The directory name is a relative path. Absolute path's is **not** supported. The relative path can contain nested folders, such as /inbox/us.

The **autoCreate** option is supported. When consumer starts, before polling is scheduled, there's additional FTP operation performed to create the directory configured for endpoint. The default value for **autoCreate** is **true**.

If no **username** is provided, then **anonymous** login is attempted using no password.

If no **port** number is provided, Camel will provide default values according to the protocol (ftp = 21, sftp = 22, ftps = 2222).

You can append query options to the URI in the following format, **?option=value&option=value&...**

This component uses two different libraries for the actual FTP work. FTP and FTPS uses [Apache Commons Net](#) while SFTP uses [JCraft JSCH](#).

FTPS (also known as FTP Secure) is an extension to FTP that adds support for the Transport Layer Security (TLS) and the Secure Sockets Layer (SSL) cryptographic protocols.

22.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

22.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

22.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

22.3. COMPONENT OPTIONS

The FTP component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

22.4. ENDPOINT OPTIONS

The FTP endpoint is configured using URI syntax:

```
ftp:host:port/directoryName
```

with the following path and query parameters:

22.4.1. Path Parameters (3 parameters)

Name	Description	Default	Type
host (common)	Required Hostname of the FTP server.		String
port (common)	Port of the FTP server.		int
directoryName (common)	The starting directory.		String

22.4.2. Query Parameters (111 parameters)

Name	Description	Default	Type
binary (common)	Specifies the file transfer mode, BINARY or ASCII. Default is ASCII (false).	false	boolean
charset (common)	This option is used to specify the encoding of the file. You can use this on the consumer, to specify the encodings of the files, which allow Camel to know the charset it should load the file content in case the file content is being accessed. Likewise when writing a file, you can use this option to specify which charset to write the file as well. Do mind that when writing the file Camel may have to read the message content into memory to be able to convert the data into the configured charset, so do not use this if you have big messages.		String
disconnect (common)	Whether or not to disconnect from remote FTP server right after use. Disconnect will only disconnect the current connection to the FTP server. If you have a consumer which you want to stop, then you need to stop the consumer/route instead.	false	boolean
doneFileName (common)	Producer: If provided, then Camel will write a 2nd done file when the original file has been written. The done file will be empty. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file will always be written in the same folder as the original file. Consumer: If provided, Camel will only consume files if a done file exists. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file is always expected in the same folder as the original file. Only <code>file.name</code> and <code>file.name.next</code> is supported as dynamic placeholders.		String

Name	Description	Default	Type
fileName (common)	Use Expression such as File Language to dynamically set the filename. For consumers, it's used as a filename filter. For producers, it's used to evaluate the filename to write. If an expression is set, it take precedence over the CamelFileName header. (Note: The header itself can also be an Expression). The expression options support both String and Expression types. If the expression is a String type, it is always evaluated using the File Language. If the expression is an Expression type, the specified Expression type is used - this allows you, for instance, to use OGNL expressions. For the consumer, you can use it to filter filenames, so you can for instance consume today's file using the File Language syntax: mydata-\${date:now:yyyyMMdd}.txt. The producers support the CamelOverruleFileName header which takes precedence over any existing CamelFileName header; the CamelOverruleFileName is a header that is used only once, and makes it easier as this avoids to temporary store CamelFileName and have to restore it afterwards.		String
passiveMode (common)	Sets passive mode connections. Default is active mode connections.	false	boolean
separator (common)	Sets the path separator to be used. UNIX = Uses unix style path separator Windows = Uses windows style path separator Auto = (is default) Use existing path separator in file name. Enum values: <ul style="list-style-type: none"> ● UNIX ● Windows ● Auto 	UNIX	PathSeparator
transferLoggingIntervalSeconds (common)	Configures the interval in seconds to use when logging the progress of upload and download operations that are in-flight. This is used for logging progress when operations takes longer time.	5	int

Name	Description	Default	Type
transferLoggingLevel (common)	<p>Configure the logging level to use when logging the progress of upload and download operations.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	DEBUG	LogLevel
transferLoggingVerbose (common)	Configures whether the perform verbose (fine grained) logging of the progress of upload and download operations.	false	boolean
fastExistsCheck (common (advanced))	If set this option to be true, camel-ftp will use the list file directly to check if the file exists. Since some FTP server may not support to list the file directly, if the option is false, camel-ftp will use the old way to list the directory and check if the file exists. This option also influences readLock=changed to control whether it performs a fast check to update file information or not. This can be used to speed up the process if the FTP server has a lot of files.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
delete (consumer)	If true, the file will be deleted after it is processed successfully.	false	boolean
moveFailed (consumer)	Sets the move failure expression based on Simple language. For example, to move files into a .error subdirectory use: .error. Note: When moving the files to the fail location Camel will handle the error and will not pick up the file again.		String

Name	Description	Default	Type
noop (consumer)	If true, the file is not moved or deleted in any way. This option is good for readonly data, or for ETL type requirements. If <code>noop=true</code> , Camel will set <code>idempotent=true</code> as well, to avoid consuming the same files over and over again.	false	boolean
preMove (consumer)	Expression (such as File Language) used to dynamically set the filename when moving it before processing. For example to move in-progress files into the order directory set this value to <code>order</code> .		String
preSort (consumer)	When pre-sort is enabled then the consumer will sort the file and directory names during polling, that was retrieved from the file system. You may want to do this in case you need to operate on the files in a sorted order. The pre-sort is executed before the consumer starts to filter, and accept files to process by Camel. This option is default=false meaning disabled.	false	boolean
recursive (consumer)	If a directory, will look for files in all the sub-directories as well.	false	boolean
resumeDownload (consumer)	Configures whether resume download is enabled. This must be supported by the FTP server (almost all FTP servers support it). In addition the options <code>localWorkDirectory</code> must be configured so downloaded files are stored in a local directory, and the option <code>binary</code> must be enabled, which is required to support resuming of downloads.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
streamDownload (consumer)	Sets the download method to use when not using a local working directory. If set to true, the remote files are streamed to the route as they are read. When set to false, the remote files are loaded into memory before being sent into the route. If enabling this option then you must set <code>stepwise=false</code> as both cannot be enabled at the same time.	false	boolean
download (consumer (advanced))	Whether the FTP consumer should download the file. If this option is set to false, then the message body will be null, but the consumer will still trigger a Camel Exchange that has details about the file such as file name, file size, etc. It's just that the file will not be downloaded.	false	boolean

Name	Description	Default	Type
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
handleDirectoryParserAbsoluteResult (consumer (advanced))	Allows you to set how the consumer will handle subfolders and files in the path if the directory parser results in with absolute paths The reason for this is that some FTP servers may return file names with absolute paths, and if so then the FTP component needs to handle this by converting the returned path into a relative path.	false	boolean
ignoreFileNotFoundOrPermissionError (consumer (advanced))	Whether to ignore when (trying to list files in directories or when downloading a file), which does not exist or due to permission error. By default when a directory or file does not exists or insufficient permission, then an exception is thrown. Setting this option to true allows to ignore that instead.	false	boolean
inProgressRepository (consumer (advanced))	A pluggable in-progress repository org.apache.camel.spi.IdempotentRepository. The in-progress repository is used to account the current in progress files being consumed. By default a memory based repository is used.		IdempotentRepository
localWorkDirectory (consumer (advanced))	When consuming, a local work directory can be used to store the remote file content directly in local files, to avoid loading the content into memory. This is beneficial, if you consume a very big remote file and thus can conserve memory.		String

Name	Description	Default	Type
onCompletionExceptionHandler (consumer (advanced))	To use a custom <code>org.apache.camel.spi.ExceptionHandler</code> to handle any thrown exceptions that happens during the file on completion process where the consumer does either a commit or rollback. The default implementation will log any exception at WARN level and ignore.		ExceptionHandler
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
processStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.component.file.GenericFileProcessStrategy</code> allowing you to implement your own readLock option or similar. Can also be used when special conditions must be met before a file can be consumed, such as a special ready file exists. If this option is set then the readLock option does not apply.		GenericFileProcessStrategy
useList (consumer (advanced))	Whether to allow using LIST command when downloading a file. Default is true. In some use cases you may want to download a specific file and are not allowed to use the LIST command, and therefore you can set this option to false. Notice when using this option, then the specific file to download does not include meta-data information such as file size, timestamp, permissions etc, because those information is only possible to retrieve when LIST command is in use.	true	boolean

Name	Description	Default	Type
fileExist (producer)	<p>What to do if a file already exists with the same name. Override, which is the default, replaces the existing file. - Append - adds content to the existing file. - Fail - throws a <code>GenericFileOperationException</code>, indicating that there is already an existing file. - Ignore - silently ignores the problem and does not override the existing file, but assumes everything is okay. - Move - option requires to use the <code>moveExisting</code> option to be configured as well. The option <code>eagerDeleteTargetFile</code> can be used to control what to do if an moving the file, and there exists already an existing file, otherwise causing the move operation to fail. The Move option will move any existing files, before writing the target file. - TryRename is only applicable if <code>tempFileName</code> option is in use. This allows to try renaming the file from the temporary name to the actual name, without doing any exists check. This check may be faster on some file systems and especially FTP servers.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● Override ● Append ● Fail ● Ignore ● Move ● TryRename 	Override	GenericFileExist
flatten (producer)	<p>Flatten is used to flatten the file name path to strip any leading paths, so it's just the file name. This allows you to consume recursively into sub-directories, but when you eg write the files to another directory they will be written in a single directory. Setting this to true on the producer enforces that any file name in <code>CamelFileName</code> header will be stripped for any leading paths.</p>	false	boolean
jailStartingDirectory (producer)	<p>Used for jailing (restricting) writing files to the starting directory (and sub) only. This is enabled by default to not allow Camel to write files to outside directories (to be more secured out of the box). You can turn this off to allow writing files to directories outside the starting directory, such as parent or root folders.</p>	true	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
moveExisting (producer)	Expression (such as File Language) used to compute file name to use when fileExist=Move is configured. To move files into a backup subdirectory just enter backup. This option only supports the following File Language tokens: file:name, file:name.ext, file:name.noext, file:onlyname, file:onlyname.noext, file:ext, and file:parent. Notice the file:parent is not supported by the FTP component, as the FTP component can only move any existing files to a relative directory based on current dir as base.		String
tempFileName (producer)	The same as tempPrefix option but offering a more fine grained control on the naming of the temporary filename as it uses the File Language. The location for tempFileName is relative to the final file location in the option 'fileName', not the target directory in the base uri. For example if option fileName includes a directory prefix: dir/finalFilename then tempFileName is relative to that subdirectory dir.		String
tempPrefix (producer)	This option is used to write the file using a temporary name and then, after the write is complete, rename it to the real name. Can be used to identify files being written and also avoid consumers (not using exclusive read locks) reading in progress files. Is often used by FTP when uploading big files.		String
allowNullBody (producer (advanced))	Used to specify if a null body is allowed during file writing. If set to true then an empty file will be created, when set to false, and attempting to send a null body to the file component, a GenericFileWriteException of 'Cannot write null body to file.' will be thrown. If the fileExist option is set to 'Override', then the file will be truncated, and if set to append the file will remain unchanged.	false	boolean
chmod (producer (advanced))	Allows you to set chmod on the stored file. For example chmod=640.		String

Name	Description	Default	Type
disconnectOnBatchComplete (producer (advanced))	Whether or not to disconnect from remote FTP server right after a Batch upload is complete. <code>disconnectOnBatchComplete</code> will only disconnect the current connection to the FTP server.	false	boolean
eagerDeleteTargetFile (producer (advanced))	Whether or not to eagerly delete any existing target file. This option only applies when you use <code>fileExists=Override</code> and the <code>tempFileName</code> option as well. You can use this to disable (set it to false) deleting the target file before the temp file is written. For example you may write big files and want the target file to exist during the temp file is being written. This ensure the target file is only deleted until the very last moment, just before the temp file is being renamed to the target filename. This option is also used to control whether to delete any existing files when <code>fileExist=Move</code> is enabled, and an existing file exists. If this option <code>copyAndDeleteOnRenameFails</code> false, then an exception will be thrown if an existing file existed, if its true, then the existing file is deleted before the move operation.	true	boolean
keepLastModified (producer (advanced))	Will keep the last modified timestamp from the source file (if any). Will use the <code>Exchange.FILE_LAST_MODIFIED</code> header to located the timestamp. This header can contain either a <code>java.util.Date</code> or long with the timestamp. If the timestamp exists and the option is enabled it will set this timestamp on the written file. Note: This option only applies to the file producer. You cannot use this option with any of the ftp producers.	false	boolean
moveExistingFileStrategy (producer (advanced))	Strategy (Custom Strategy) used to move file with special naming token to use when <code>fileExist=Move</code> is configured. By default, there is an implementation used if no custom strategy is provided.		FileMoveExistingStrategy
sendNoop (producer (advanced))	Whether to send a noop command as a pre-write check before uploading files to the FTP server. This is enabled by default as a validation of the connection is still valid, which allows to silently re-connect to be able to upload the file. However if this causes problems, you can turn this option off.	true	boolean

Name	Description	Default	Type
activePortRange (advanced)	Set the client side port range in active mode. The syntax is: minPort-maxPort Both port numbers are inclusive, eg 10000-19999 to include all lxxx ports.		String
autoCreate (advanced)	Automatically create missing directories in the file's pathname. For the file consumer, that means creating the starting directory. For the file producer, it means the directory the files should be written to.	true	boolean
bufferSize (advanced)	Buffer size in bytes used for writing files (or in case of FTP for downloading and uploading files).	131072	int
connectTimeout (advanced)	Sets the connect timeout for waiting for a connection to be established Used by both FTPClient and JSCH.	10000	int
ftpClient (advanced)	To use a custom instance of FTPClient.		FTPClient
ftpClientConfig (advanced)	To use a custom instance of FTPClientConfig to configure the FTP client the endpoint should use.		FTPClientConfig
ftpClientConfigParameters (advanced)	Used by FtpComponent to provide additional parameters for the FTPClientConfig.		Map
ftpClientParameters (advanced)	Used by FtpComponent to provide additional parameters for the FTPClient.		Map
maximumReconnectAttempts (advanced)	Specifies the maximum reconnect attempts Camel performs when it tries to connect to the remote FTP server. Use 0 to disable this behavior.		int
reconnectDelay (advanced)	Delay in millis Camel will wait before performing a reconnect attempt.	1000	long
siteCommand (advanced)	Sets optional site command(s) to be executed after successful login. Multiple site commands can be separated using a new line character.		String
soTimeout (advanced)	Sets the so timeout FTP and FTPS Is the SocketOptions.SO_TIMEOUT value in millis. Recommended option is to set this to 300000 so as not have a hanged connection. On SFTP this option is set as timeout on the JSCH Session instance.	300000	int

Name	Description	Default	Type
stepwise (advanced)	Sets whether we should stepwise change directories while traversing file structures when downloading files, or as well when uploading a file to a directory. You can disable this if you for example are in a situation where you cannot change directory on the FTP server due security reasons. Stepwise cannot be used together with streamDownload.	true	boolean
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
throwExceptionOnConnectFailed (advanced)	Should an exception be thrown if connection failed (exhausted)By default exception is not thrown and a WARN is logged. You can use this to enable exception being thrown and handle the thrown exception from the <code>org.apache.camel.spi.PollingConsumerPollStrategy</code> rollback method.	false	boolean
timeout (advanced)	Sets the data timeout for waiting for reply Used only by FTPClient.	30000	int
antExclude (filter)	Ant style filter exclusion. If both antInclude and antExclude are used, antExclude takes precedence over antInclude. Multiple exclusions may be specified in comma-delimited format.		String
antFilterCaseSensitive (filter)	Sets case sensitive flag on ant filter.	true	boolean
antInclude (filter)	Ant style filter inclusion. Multiple inclusions may be specified in comma-delimited format.		String
eagerMaxMessagesPerPoll (filter)	Allows for controlling whether the limit from <code>maxMessagesPerPoll</code> is eager or not. If eager then the limit is during the scanning of files. Where as false would scan all files, and then perform sorting. Setting this option to false allows for sorting all files first, and then limit the poll. Mind that this requires a higher memory usage as all file details are in memory to perform the sorting.	true	boolean

Name	Description	Default	Type
exclude (filter)	Is used to exclude files, if filename matches the regex pattern (matching is case in-sensitive). Notice if you use symbols such as plus sign and others you would need to configure this using the RAW() syntax if configuring this as an endpoint uri. See more details at configuring endpoint uris.		String
excludeExt (filter)	Is used to exclude files matching file extension name (case insensitive). For example to exclude bak files, then use excludeExt=bak. Multiple extensions can be separated by comma, for example to exclude bak and dat files, use excludeExt=bak,dat. Note that the file extension includes all parts, for example having a file named mydata.tar.gz will have extension as tar.gz. For more flexibility then use the include/exclude options.		String
filter (filter)	Pluggable filter as a org.apache.camel.component.file.GenericFileFilter class. Will skip files if filter returns false in its accept() method.		GenericFileFilter
filterDirectory (filter)	Filters the directory based on Simple language. For example to filter on current date, you can use a simple date pattern such as <code>\$\$\{date:now:yyMMdd}</code> .		String
filterFile (filter)	Filters the file based on Simple language. For example to filter on file size, you can use <code>\$\$\{file:size} 5000</code> .		String
idempotent (filter)	Option to use the Idempotent Consumer EIP pattern to let Camel skip already processed files. Will by default use a memory based LRUcache that holds 1000 entries. If noop=true then idempotent will be enabled as well to avoid consuming the same files over and over again.	false	Boolean
idempotentKey (filter)	To use a custom idempotent key. By default the absolute path of the file is used. You can use the File Language, for example to use the file name and file size, you can do: <code>idempotentKey=\$\{file:name}-\$\{file:size}</code> .		String
idempotentRepository (filter)	A pluggable repository org.apache.camel.spi.IdempotentRepository which by default use MemoryIdempotentRepository if none is specified and idempotent is true.		IdempotentRepository

Name	Description	Default	Type
include (filter)	Is used to include files, if filename matches the regex pattern (matching is case in-sensitive). Notice if you use symbols such as plus sign and others you would need to configure this using the RAW() syntax if configuring this as an endpoint uri. See more details at configuring endpoint uris.		String
includeExt (filter)	Is used to include files matching file extension name (case insensitive). For example to include txt files, then use includeExt=txt. Multiple extensions can be separated by comma, for example to include txt and xml files, use includeExt=txt,xml. Note that the file extension includes all parts, for example having a file named mydata.tar.gz will have extension as tar.gz. For more flexibility then use the include/exclude options.		String
maxDepth (filter)	The maximum depth to traverse when recursively processing a directory.	2147483647	int
maxMessagesPerPoll (filter)	To define a maximum messages to gather per poll. By default no maximum is set. Can be used to set a limit of e.g. 1000 to avoid when starting up the server that there are thousands of files. Set a value of 0 or negative to disabled it. Notice: If this option is in use then the File and FTP components will limit before any sorting. For example if you have 100000 files and use maxMessagesPerPoll=500, then only the first 500 files will be picked up, and then sorted. You can use the eagerMaxMessagesPerPoll option and set this to false to allow to scan all files first and then sort afterwards.		int
minDepth (filter)	The minimum depth to start processing when recursively processing a directory. Using minDepth=1 means the base directory. Using minDepth=2 means the first sub directory.		int
move (filter)	Expression (such as Simple Language) used to dynamically set the filename when moving it after processing. To move files into a .done subdirectory just enter .done.		String
exclusiveReadLockStrategy (lock)	Pluggable read-lock as a org.apache.camel.component.file.GenericFileExclusiveReadLockStrategy implementation.		GenericFileExclusiveReadLockStrategy
readLock (lock)	Used by consumer, to only poll the files if it has	none	String

Name	Description	Default	Type
	<p>exclusive read-lock on the file (i.e. the file is not in progress of being written). Camel will wait until the file lock is granted. This option provides the build in strategies:</p> <ul style="list-style-type: none"> - none - No read lock is in use - markerFile - Camel creates a marker file (fileName.camellLock) and then holds a lock on it. This option is not available for the FTP component - changed - Changed is using file length/modification timestamp to detect whether the file is currently being copied or not. Will at least use 1 sec to determine this, so this option cannot consume files as fast as the others, but can be more reliable as the JDK IO API cannot always determine whether a file is currently being used by another process. The option readLockCheckInterval can be used to set the check frequency. - fileLock - is for using java.nio.channels.FileLock. This option is not avail for Windows OS and the FTP component. This approach should be avoided when accessing a remote file system via a mount/share unless that file system supports distributed file locks. - rename - rename is for using a try to rename the file as a test if we can get exclusive read-lock. - idempotent - (only for file component) idempotent is for using a idempotentRepository as the read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. - idempotent-changed - (only for file component) idempotent-changed is for using a idempotentRepository and changed as the combined read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. - idempotent-rename - (only for file component) idempotent-rename is for using a idempotentRepository and rename as the combined read-lock. This allows to use read locks that supports clustering if the idempotent repository implementation supports that. <p>Notice: The various read locks is not all suited to work in clustered mode, where concurrent consumers on different nodes is competing for the same files on a shared file system. The markerFile using a close to atomic operation to create the empty marker file, but its not guaranteed to work in a cluster. The fileLock may work better but then the file system need to support distributed file locks, and so on. Using the idempotent read lock can support clustering if the idempotent repository supports clustering, such as Hazelcast Component or Infinispan.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● none ● markerFile ● fileLock 		

Name	<ul style="list-style-type: none"> ● rename ● changed Description	Default	Type
	<ul style="list-style-type: none"> ● idempotent ● idempotent-changed ● idempotent-rename 		
readLockCheckInterval (lock)	Interval in millis for the read-lock, if supported by the read lock. This interval is used for sleeping between attempts to acquire the read lock. For example when using the changed read lock, you can set a higher interval period to cater for slow writes. The default of 1 sec. may be too fast if the producer is very slow writing the file. Notice: For FTP the default readLockCheckInterval is 5000. The readLockTimeout value must be higher than readLockCheckInterval, but a rule of thumb is to have a timeout that is at least 2 or more times higher than the readLockCheckInterval. This is needed to ensure that ample time is allowed for the read lock process to try to grab the lock before the timeout was hit.	1000	long
readLockDeleteOrphanLockFiles (lock)	Whether or not read lock with marker files should upon startup delete any orphan read lock files, which may have been left on the file system, if Camel was not properly shutdown (such as a JVM crash). If turning this option to false then any orphaned lock file will cause Camel to not attempt to pickup that file, this could also be due another node is concurrently reading files from the same shared directory.	true	boolean
readLockLoggingLevel (lock)	<p>Logging level used when a read lock could not be acquired. By default a DEBUG is logged. You can change this level, for example to OFF to not have any logging. This option is only applicable for readLock of types: changed, fileLock, idempotent, idempotent-changed, idempotent-rename, rename.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	DEBUG	LoggingLevel

Name	Description	Default	Type
readLockMarkerFile (lock)	Whether to use marker file with the changed, rename, or exclusive read lock types. By default a marker file is used as well to guard against other processes picking up the same files. This behavior can be turned off by setting this option to false. For example if you do not want to write marker files to the file systems by the Camel application.	true	boolean
readLockMinAge (lock)	This option is applied only for readLock=changed. It allows to specify a minimum age the file must be before attempting to acquire the read lock. For example use readLockMinAge=300s to require the file is at least 5 minutes old. This can speedup the changed read lock as it will only attempt to acquire files which are at least that given age.	0	long
readLockMinLength (lock)	This option is applied only for readLock=changed. It allows you to configure a minimum file length. By default Camel expects the file to contain data, and thus the default value is 1. You can set this option to zero, to allow consuming zero-length files.	1	long
readLockRemoveOnCommit (lock)	This option is applied only for readLock=idempotent. It allows to specify whether to remove the file name entry from the idempotent repository when processing the file is succeeded and a commit happens. By default the file is not removed which ensures that any race-condition do not occur so another active node may attempt to grab the file. Instead the idempotent repository may support eviction strategies that you can configure to evict the file name entry after X minutes - this ensures no problems with race conditions. See more details at the readLockIdempotentReleaseDelay option.	false	boolean
readLockRemoveOnRollback (lock)	This option is applied only for readLock=idempotent. It allows to specify whether to remove the file name entry from the idempotent repository when processing the file failed and a rollback happens. If this option is false, then the file name entry is confirmed (as if the file did a commit).	true	boolean

Name	Description	Default	Type
readLockTimeout (lock)	Optional timeout in millis for the read-lock, if supported by the read-lock. If the read-lock could not be granted and the timeout triggered, then Camel will skip the file. At next poll Camel, will try the file again, and this time maybe the read-lock could be granted. Use a value of 0 or lower to indicate forever. Currently fileLock, changed and rename support the timeout. Notice: For FTP the default readLockTimeout value is 20000 instead of 10000. The readLockTimeout value must be higher than readLockCheckInterval, but a rule of thumb is to have a timeout that is at least 2 or more times higher than the readLockCheckInterval. This is needed to ensure that ample time is allowed for the read lock process to try to grab the lock before the timeout was hit.	10000	long
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long

Name	Description	Default	Type
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	<p>Time unit for initialDelay and delay options.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit

Name	Description	Default	Type
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
account (security)	Account to use for login.		String
password (security)	Password to use for login.		String
username (security)	Username to use for login.		String
shuffle (sort)	To shuffle the list of files (sort in random order).	false	boolean
sortBy (sort)	Built-in sort by using the File Language. Supports nested sorts, so you can have a sort by file name and as a 2nd group sort by modified date.		String
sorter (sort)	Pluggable sorter as a java.util.Comparator class.		Comparator

22.5. FTPS COMPONENT DEFAULT TRUST STORE

When using the **ftpClient**, properties related to SSL with the FTPS component, the trust store accept all certificates. If you only want trust selective certificates, you have to configure the trust store with the **ftpClient.trustStore.xxx** options or by configuring a custom **ftpClient**.

When using **sslContextParameters**, the trust store is managed by the configuration of the provided SSLContextParameters instance.

You can configure additional options on the **ftpClient** and **ftpClientConfig** from the URI directly by using the **ftpClient.** or **ftpClientConfig.** prefix.

For example to set the **setDataTimeout** on the **FTPClient** to 30 seconds you can do:

```
from("ftp://foo@myserver?password=secret&ftpClient.dataTimeout=30000").to("bean:foo");
```

You can mix and match and have use both prefixes, for example to configure date format or timezones.

```
from("ftp://foo@myserver?password=secret&ftpClient.dataTimeout=30000&ftpClientConfig.serverLanguageCode=fr").to("bean:foo");
```

You can have as many of these options as you like.

See the documentation of the Apache Commons FTP FTPClientConfig for possible options and more details. And as well for Apache Commons FTP FTPClient.

If you do not like having many and long configuration in the url you can refer to the **ftpClient** or **ftpClientConfig** to use by letting Camel lookup in the Registry for it.

For example:

```
<bean id="myConfig" class="org.apache.commons.net.ftp.FTPClientConfig">
  <property name="lenientFutureDates" value="true"/>
  <property name="serverLanguageCode" value="fr"/>
</bean>
```

And then let Camel lookup this bean when you use the # notation in the url.

```
from("ftp://foo@myserver?password=secret&ftpClientConfig=#myConfig").to("bean:foo");
```

22.6. EXAMPLES

```
ftp://someone@someftpserver.com/public/upload/images/holiday2008?password=secret&binary=true
```

```
ftp://someoneelse@someotherftpserver.co.uk:12049/reports/2008/password=secret&binary=false
```

```
ftp://publicftpserver.com/download
```

22.7. CONCURRENCY

FTP Consumer does not support concurrency

The FTP consumer (with the same endpoint) does not support concurrency (the backing FTP client is not thread safe).

You can use multiple FTP consumers to poll from different endpoints. It is only a single endpoint that does not support concurrent consumers.

The FTP producer does **not** have this issue, it supports concurrency.

22.8. MORE INFORMATION

This component is an extension of the File component. So there are more samples and details on the File component page.

22.9. DEFAULT WHEN CONSUMING FILES

The FTP consumer will by default leave the consumed files untouched on the remote FTP server. You have to configure it explicitly if you want it to delete the files or move them to another location. For example you can use **delete=true** to delete the files, or use **move=.done** to move the files into a hidden done sub directory.

The regular File consumer is different as it will by default move files to a **.camel** sub directory. The reason Camel does **not** do this by default for the FTP consumer is that it may lack permissions by default to be able to move or delete files.

22.9.1. limitations

The option **readLock** can be used to force Camel **not** to consume files that is currently in the progress of being written. However, this option is turned off by default, as it requires that the user has write access. See the options table at File2 for more details about read locks.

There are other solutions to avoid consuming files that are currently being written over FTP; for instance, you can write to a temporary destination and move the file after it has been written.

When moving files using **move** or **preMove** option the files are restricted to the FTP_ROOT folder. That prevents you from moving files outside the FTP area. If you want to move files to another area you can use soft links and move files into a soft linked folder.

22.10. MESSAGE HEADERS

The following message headers can be used to affect the behavior of the component

Header	Description
CamelFileName	Specifies the output file name (relative to the endpoint directory) to be used for the output message when sending to the endpoint. If this is not present and no expression either, then a generated message ID is used as the filename instead.
CamelFileNameProduced	The actual filepath (path + name) for the output file that was written. This header is set by Camel and its purpose is providing end-users the name of the file that was written.
CamelFileNameConsumed	The file name of the file consumed
CamelFileHost	The remote hostname.
CamelFileLocalWorkPath	Path to the local work file, if local work directory is used.

In addition the FTP/FTPS consumer and producer will enrich the Camel **Message** with the following headers

Header	Description
CamelFtpReplyCode	The FTP client reply code (the type is a integer)
CamelFtpReplyString	The FTP client reply string

22.10.1. Exchange Properties

Camel sets the following exchange properties

Header	Description
CamelBatchIndex	Current index out of total number of files being consumed in this batch.
CamelBatchSize	Total number of files being consumed in this batch.

Header	Description
CamelBatchComplete	True if there are no more files in this batch.

22.11. ABOUT TIMEOUTS

The two set of libraries (see top) has different API for setting timeout. You can use the **connectTimeout** option for both of them to set a timeout in millis to establish a network connection. An individual **soTimeout** can also be set on the FTP/FTPS, which corresponds to using **ftpClient.soTimeout**. Notice SFTP will automatically use **connectTimeout** as its **soTimeout**. The **timeout** option only applies for FTP/FTPS as the data timeout, which corresponds to the **ftpClient.dataTimeout** value. All timeout values are in millis.

22.12. USING LOCAL WORK DIRECTORY

Camel supports consuming from remote FTP servers and downloading the files directly into a local work directory. This avoids reading the entire remote file content into memory as it is streamed directly into the local file using **FileOutputStream**.

Camel will store to a local file with the same name as the remote file, though with **.inprogress** as extension while the file is being downloaded. Afterwards, the file is renamed to remove the **.inprogress** suffix. And finally, when the Exchange is complete the local file is deleted.

So if you want to download files from a remote FTP server and store it as files then you need to route to a file endpoint such as:

```
from("ftp://someone@someserver.com?password=secret&localWorkDirectory=/tmp").to("file://inbox");
```



NOTE

The route above is ultra efficient as it avoids reading the entire file content into memory. It will download the remote file directly to a local file stream. The **java.io.File** handle is then used as the Exchange body. The file producer leverages this fact and can work directly on the work file **java.io.File** handle and perform a **java.io.File.rename** to the target filename. As Camel knows it's a local work file, it can optimize and use a rename instead of a file copy, as the work file is meant to be deleted anyway.

22.13. STEPWISE CHANGING DIRECTORIES

Camel FTP can operate in two modes in terms of traversing directories when consuming files (eg downloading) or producing files (eg uploading)

- stepwise
- not stepwise

You may want to pick either one depending on your situation and security issues. Some Camel end users can only download files if they use stepwise, while others can only download if they do not.

You can use the **stepwise** option to control the behavior.

Note that stepwise changing of directory will in most cases only work when the user is confined to its home directory and when the home directory is reported as `"/"`.

The difference between the two of them is best illustrated with an example. Suppose we have the following directory structure on the remote FTP server we need to traverse and download files:

```

/
/one
/one/two
/one/two/sub-a
/one/two/sub-b

```

And that we have a file in each of sub-a (a.txt) and sub-b (b.txt) folder.

22.14. USING STEPWISE=TRUE (DEFAULT MODE)

```

TYPE A
200 Type set to A
PWD
257 "/" is current directory.
CWD one
250 CWD successful. "/one" is current directory.
CWD two
250 CWD successful. "/one/two" is current directory.
SYST
215 UNIX emulated by FileZilla
PORT 127,0,0,1,17,94
200 Port command successful
LIST
150 Opening data channel for directory list.
226 Transfer OK
CWD sub-a
250 CWD successful. "/one/two/sub-a" is current directory.
PORT 127,0,0,1,17,95
200 Port command successful
LIST
150 Opening data channel for directory list.
226 Transfer OK
CDUP
200 CDUP successful. "/one/two" is current directory.
CWD sub-b
250 CWD successful. "/one/two/sub-b" is current directory.
PORT 127,0,0,1,17,96
200 Port command successful
LIST
150 Opening data channel for directory list.
226 Transfer OK
CDUP
200 CDUP successful. "/one/two" is current directory.
CWD /
250 CWD successful. "/" is current directory.
PWD
257 "/" is current directory.
CWD one
250 CWD successful. "/one" is current directory.

```



```

CWD two
250 CWD successful. "/one/two" is current directory.
PORT 127,0,0,1,17,97
200 Port command successful
RETR foo.txt
150 Opening data channel for file transfer.
226 Transfer OK
CWD /
250 CWD successful. "/" is current directory.
PWD
257 "/" is current directory.
CWD one
250 CWD successful. "/one" is current directory.
CWD two
250 CWD successful. "/one/two" is current directory.
CWD sub-a
250 CWD successful. "/one/two/sub-a" is current directory.
PORT 127,0,0,1,17,98
200 Port command successful
RETR a.txt
150 Opening data channel for file transfer.
226 Transfer OK
CWD /
250 CWD successful. "/" is current directory.
PWD
257 "/" is current directory.
CWD one
250 CWD successful. "/one" is current directory.
CWD two
250 CWD successful. "/one/two" is current directory.
CWD sub-b
250 CWD successful. "/one/two/sub-b" is current directory.
PORT 127,0,0,1,17,99
200 Port command successful
RETR b.txt
150 Opening data channel for file transfer.
226 Transfer OK
CWD /
250 CWD successful. "/" is current directory.
QUIT
221 Goodbye
disconnected.

```

As you can see when stepwise is enabled, it will traverse the directory structure using CD xxx.

22.15. USING STEPWISE=FALSE

```

230 Logged on
TYPE A
200 Type set to A
SYST
215 UNIX emulated by FileZilla
PORT 127,0,0,1,4,122
200 Port command successful
LIST one/two

```

```

150 Opening data channel for directory list
226 Transfer OK
PORT 127,0,0,1,4,123
200 Port command successful
LIST one/two/sub-a
150 Opening data channel for directory list
226 Transfer OK
PORT 127,0,0,1,4,124
200 Port command successful
LIST one/two/sub-b
150 Opening data channel for directory list
226 Transfer OK
PORT 127,0,0,1,4,125
200 Port command successful
RETR one/two/foo.txt
150 Opening data channel for file transfer.
226 Transfer OK
PORT 127,0,0,1,4,126
200 Port command successful
RETR one/two/sub-a/a.txt
150 Opening data channel for file transfer.
226 Transfer OK
PORT 127,0,0,1,4,127
200 Port command successful
RETR one/two/sub-b/b.txt
150 Opening data channel for file transfer.
226 Transfer OK
QUIT
221 Goodbye
disconnected.

```

As you can see when not using stepwise, there are no CD operation invoked at all.

22.16. SAMPLES

In the sample below we set up Camel to download all the reports from the FTP server once every hour (60 min) as BINARY content and store it as files on the local file system.

And the route using XML DSL:

```

<route>
  <from uri="ftp://scott@localhost/public/reports?
password=tiger&binary=true&delay=60000"/>
  <to uri="file://target/test-reports"/>
</route>

```

22.16.1. Consuming a remote FTPS server (implicit SSL) and client authentication

```

from("ftps://admin@localhost:2222/public/camel?
password=admin&securityProtocol=SSL&implicit=true
  &ftpClient.keyStore.file=./src/test/resources/server.jks
  &ftpClient.keyStore.password=password&ftpClient.keyStore.keyPassword=password")
.to("bean:foo");

```

22.16.2. Consuming a remote FTPS server (explicit TLS) and a custom trust store configuration

```
from("ftps://admin@localhost:2222/public/camel?
password=admin&ftpClient.trustStore.file=./src/test/resources/server.jks&ftpClient.trustStore.password=
password")
.to("bean:foo");
```

22.17. CUSTOM FILTERING

Camel supports pluggable filtering strategies. This strategy it to use the build in **org.apache.camel.component.file.GenericFileFilter** in Java. You can then configure the endpoint with such a filter to skip certain filters before being processed.

In the sample we have built our own filter that only accepts files starting with report in the filename.

And then we can configure our route using the **filter** attribute to reference our filter (using # notation) that we have defined in the spring XML file:

```
<!-- define our sorter as a plain spring bean -->
<bean id="myFilter" class="com.mycompany.MyFileFilter"/>

<route>
  <from uri="ftp://someuser@someftpsserver.com?password=secret&filter=#myFilter"/>
  <to uri="bean:processInbox"/>
</route>
```

22.18. FILTERING USING ANT PATH MATCHER

The ANT path matcher is a filter that is shipped out-of-the-box in the **camel-spring** jar. So you need to depend on **camel-spring** if you are using Maven.

The reason is that we leverage Spring's [AntPathMatcher](#) to do the actual matching.

The file paths are matched with the following rules:

- **?** matches one character
- ***** matches zero or more characters
- ****** matches zero or more directories in a path

The sample below demonstrates how to use it:

22.19. USING A PROXY WITH SFTP

To use an HTTP proxy to connect to your remote host, you can configure your route in the following way:

```
<!-- define our sorter as a plain spring bean -->
<bean id="proxy" class="com.jcraft.jsch.ProxyHTTP">
  <constructor-arg value="localhost"/>
  <constructor-arg value="7777"/>
</bean>
```

```
<route>
  <from uri="sftp://localhost:9999/root?username=admin&password=admin&proxy=#proxy"/>
  <to uri="bean:processFile"/>
</route>
```

You can also assign a user name and password to the proxy, if necessary. Please consult the documentation for **com.jcraft.jsch.Proxy** to discover all options.

22.20. SETTING PREFERRED SFTP AUTHENTICATION METHOD

If you want to explicitly specify the list of authentication methods that should be used by **sftp** component, use **preferredAuthentications** option. If for example you would like Camel to attempt to authenticate with private/public SSH key and fallback to user/password authentication in the case when no public key is available, use the following route configuration:

```
from("sftp://localhost:9999/root?
username=admin&password=admin&preferredAuthentications=publickey,password").
to("bean:processFile");
```

22.21. CONSUMING A SINGLE FILE USING A FIXED NAME

When you want to download a single file and knows the file name, you can use **fileName=myFileName.txt** to tell Camel the name of the file to download. By default the consumer will still do a FTP LIST command to do a directory listing and then filter these files based on the **fileName** option. Though in this use-case it may be desirable to turn off the directory listing by setting **useList=false**. For example the user account used to login to the FTP server may not have permission to do a FTP LIST command. So you can turn off this with **useList=false**, and then provide the fixed name of the file to download with **fileName=myFileName.txt**, then the FTP consumer can still download the file. If the file for some reason does not exist, then Camel will by default throw an exception, you can turn this off and ignore this by setting **ignoreFileNotFoundOrPermissionError=true**.

For example to have a Camel route that pickup a single file, and delete it after use you can do

```
from("ftp://admin@localhost:21/nolist/?
password=admin&stepwise=false&useList=false&ignoreFileNotFoundOrPermissionError=true&fileName
=report.txt&delete=true")
.to("activemq:queue:report");
```

Notice that we have used all the options we talked above.

You can also use this with **ConsumerTemplate**. For example to download a single file (if it exists) and grab the file content as a String type:

```
String data = template.retrieveBodyNoWait("ftp://admin@localhost:21/nolist/?
password=admin&stepwise=false&useList=false&ignoreFileNotFoundOrPermissionError=true&fileName
=report.txt&delete=true", String.class);
```

22.22. DEBUG LOGGING

This component has log level **TRACE** that can be helpful if you have problems.

22.23. SPRING BOOT AUTO-CONFIGURATION

When using ftp with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-ftp-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 13 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.ftp.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.ftp.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.ftp.enabled</code>	Whether to enable auto configuration of the ftp component. This is enabled by default.		Boolean
<code>camel.component.ftp.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

Name	Description	Default	Type
<code>camel.component.ftp.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.ftp.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.ftp.enabled</code>	Whether to enable auto configuration of the ftps component. This is enabled by default.		Boolean
<code>camel.component.ftp.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.ftp.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	false	Boolean
<code>camel.component.sftp.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.sftp.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.sftp.enabled	Whether to enable auto configuration of the sftp component. This is enabled by default.		Boolean
camel.component.sftp.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 23. HTTP

Only producer is supported

The HTTP component provides HTTP based endpoints for calling external HTTP resources (as a client to call external servers using HTTP).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-http</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

23.1. URI FORMAT

```
http:hostname[:port][/resourceUri][?options]
```

Will by default use port 80 for HTTP and 443 for HTTPS.

23.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

23.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

23.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

23.3. COMPONENT OPTIONS

The HTTP component supports 37 options, which are listed below.

Name	Description	Default	Type
cookieStore (producer)	To use a custom <code>org.apache.http.client.CookieStore</code> . By default the <code>org.apache.http.impl.client.BasicCookieStore</code> is used which is an in-memory only cookie store. Notice if <code>bridgeEndpoint=true</code> then the cookie store is forced to be a noop cookie store as cookie shouldn't be stored as we are just bridging (eg acting as a proxy).		CookieStore
copyHeaders (producer)	If this option is true then IN exchange headers will be copied to OUT exchange headers according to copy strategy. Setting this to false, allows to only include the headers from the HTTP response (not propagating IN headers).	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
responsePayloadStreamingThreshold (producer)	This threshold in bytes controls whether the response payload should be stored in memory as a byte array or be streaming based. Set this to -1 to always use streaming mode.	8192	int
skipRequestHeaders (producer (advanced))	Whether to skip mapping all the Camel headers as HTTP request headers. If there are no data from Camel headers needed to be included in the HTTP request then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	boolean

Name	Description	Default	Type
skipResponseHeaders (producer (advanced))	Whether to skip mapping all the HTTP response headers to Camel headers. If there are no data needed from HTTP headers then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	boolean
allowJavaSerializedObject (advanced)	Whether to allow java serialization when a request uses context-type=application/x-java-serialized-object. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	boolean
authCachingDisabled (advanced)	Disables authentication scheme caching.	false	boolean
automaticRetriesDisabled (advanced)	Disables automatic request recovery and re-execution.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
clientConnectionManager (advanced)	To use a custom and shared HttpClientConnectionManager to manage connections. If this has been configured then this is always used for all endpoints created by this component.		HttpClientConnectionManager
connectionsPerRoute (advanced)	The maximum number of connections per route.	20	int
connectionStateDisabled (advanced)	Disables connection state tracking.	false	boolean
connectionTimeToLive (advanced)	The time for connection to live, the time unit is millisecond, the default value is always keep alive.		long
contentCompressionDisabled (advanced)	Disables automatic content decompression.	false	boolean

Name	Description	Default	Type
cookieManagementDisabled (advanced)	Disables state (cookie) management.	false	boolean
defaultUserAgentDisabled (advanced)	Disables the default user agent set by this builder if none has been provided by the user.	false	boolean
httpBinding (advanced)	To use a custom <code>HttpBinding</code> to control the mapping between Camel message and <code>HttpClient</code> .		<code>HttpBinding</code>
httpClientConfigurer (advanced)	To use the custom <code>HttpClientConfigurer</code> to perform configuration of the <code>HttpClient</code> that will be used.		<code>HttpClientConfigurer</code>
httpConfiguration (advanced)	To use the shared <code>HttpConfiguration</code> as base configuration.		<code>HttpConfiguration</code>
httpContext (advanced)	To use a custom <code>org.apache.http.protocol.HttpContext</code> when executing requests.		<code>HttpContext</code>
maxTotalConnections (advanced)	The maximum number of connections.	200	int
redirectHandlingDisabled (advanced)	Disables automatic redirect handling.	false	boolean
headerFilterStrategy (filter)	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message.		<code>HeaderFilterStrategy</code>
proxyAuthDomain (proxy)	Proxy authentication domain to use.		String
proxyAuthHost (proxy)	Proxy authentication host.		String
proxyAuthMethod (proxy)	Proxy authentication method to use. Enum values: <ul style="list-style-type: none"> ● Basic ● Digest ● NTLM 		String

Name	Description	Default	Type
proxyAuthNtHost (proxy)	Proxy authentication domain (workstation name) to use with NTLM.		String
proxyAuthPassword (proxy)	Proxy authentication password.		String
proxyAuthPort (proxy)	Proxy authentication port.		Integer
proxyAuthUsername (proxy)	Proxy authentication username.		String
sslContextParameters (security)	To configure security using SSLContextParameters. Important: Only one instance of org.apache.camel.support.jsse.SSLContextParameters is supported per HttpComponent. If you need to use 2 or more different instances, you need to define a new HttpComponent per instance you need.		SSLContextParameters
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean
x509HostnameVerifier (security)	To use a custom X509HostnameVerifier such as DefaultHostnameVerifier or NoopHostnameVerifier.		HostnameVerifier
connectionRequestTimeout (timeout)	The timeout in milliseconds used when requesting a connection from the connection manager. A timeout value of zero is interpreted as an infinite timeout. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	int
connectTimeout (timeout)	Determines the timeout in milliseconds until a connection is established. A timeout value of zero is interpreted as an infinite timeout. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	int
socketTimeout (timeout)	Defines the socket timeout in milliseconds, which is the timeout for waiting for data or, put differently, a maximum period inactivity between two consecutive data packets). A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	int

23.4. ENDPOINT OPTIONS

The HTTP endpoint is configured using URI syntax:

```
http://httpUri
```

with the following path and query parameters:

23.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
httpUri (common)	Required The url of the HTTP endpoint to call.		URI

23.4.2. Query Parameters (51 parameters)

Name	Description	Default	Type
chunked (producer)	If this option is false the Servlet will disable the HTTP streaming and set the content-length header on the response.	true	boolean
disableStreamCache (common)	Determines whether or not the raw input stream from Servlet is cached or not (Camel will read the stream into a in memory/overflow to file, Stream caching) cache. By default Camel will cache the Servlet input stream to support reading it multiple times to ensure it Camel can retrieve all data from the stream. However you can set this option to true when you for example need to access the raw stream, such as streaming it directly to a file or other persistent store. DefaultHttpBinding will copy the request input stream into a stream cache and put it into message body if this option is false to support reading the stream multiple times. If you use Servlet to bridge/proxy an endpoint then consider enabling this option to improve performance, in case you do not need to read the message payload multiple times. The http producer will by default cache the response body stream. If setting this option to true, then the producers will not cache the response body stream but use the response stream as-is as the message body.	false	boolean
headerFilterStrategy (common)	To use a custom HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
httpBinding (common (advanced))	To use a custom HttpBinding to control the mapping between Camel message and HttpClient.		HttpBinding

Name	Description	Default	Type
bridgeEndpoint (producer)	If the option is true, HttpProducer will ignore the Exchange.HTTP_URI header, and use the endpoint's URI for request. You may also set the option <code>throwExceptionOnFailure</code> to be false to let the HttpProducer send all the fault response back.	false	boolean
clearExpiredCookies (producer)	Whether to clear expired cookies before sending the HTTP request. This ensures the cookies store does not keep growing by adding new cookies which is newer removed when they are expired. If the component has disabled cookie management then this option is disabled too.	true	boolean
connectionClose (producer)	Specifies whether a Connection Close header must be added to HTTP Request. By default <code>connectionClose</code> is false.	false	boolean
copyHeaders (producer)	If this option is true then IN exchange headers will be copied to OUT exchange headers according to copy strategy. Setting this to false, allows to only include the headers from the HTTP response (not propagating IN headers).	true	boolean
customHostHeader (producer)	To use custom host header for producer. When not set in query will be ignored. When set will override host header derived from url.		String
httpMethod (producer)	Configure the HTTP method to use. The <code>HttpMethod</code> header cannot override this option if set. Enum values: <ul style="list-style-type: none"> ● GET ● POST ● PUT ● DELETE ● HEAD ● OPTIONS ● TRACE ● PATCH 		HttpMethods
ignoreResponseBody (producer)	If this option is true, The http producer won't read response body and cache the input stream.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
preserveHostHeader (producer)	If the option is true, HttpProducer will set the Host header to the value contained in the current exchange Host header, useful in reverse proxy applications where you want the Host header received by the downstream server to reflect the URL called by the upstream client, this allows applications which use the Host header to generate accurate URL's for a proxied service.	false	boolean
throwExceptionOnFailure (producer)	Option to disable throwing the <code>HttpOperationFailedException</code> in case of failed responses from the remote server. This allows you to get all responses regardless of the HTTP status code.	true	boolean
transferException (producer)	If enabled and an Exchange failed processing on the consumer side, and if the caused Exception was send back serialized in the response as a <code>application/x-java-serialized-object</code> content type. On the producer side the exception will be deserialized and thrown as is, instead of the <code>HttpOperationFailedException</code> . The caused exception is required to be serialized. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	boolean
cookieHandler (producer (advanced))	Configure a cookie handler to maintain a HTTP session.		CookieHandler

Name	Description	Default	Type
cookieStore (producer (advanced))	To use a custom CookieStore. By default the BasicCookieStore is used which is an in-memory only cookie store. Notice if bridgeEndpoint=true then the cookie store is forced to be a noop cookie store as cookie shouldn't be stored as we are just bridging (eg acting as a proxy). If a cookieHandler is set then the cookie store is also forced to be a noop cookie store as cookie handling is then performed by the cookieHandler.		CookieStore
deleteWithBody (producer (advanced))	Whether the HTTP DELETE should include the message body or not. By default HTTP DELETE do not include any HTTP body. However in some rare cases users may need to be able to include the message body.	false	boolean
getWithBody (producer (advanced))	Whether the HTTP GET should include the message body or not. By default HTTP GET do not include any HTTP body. However in some rare cases users may need to be able to include the message body.	false	boolean
okStatusCodeRange (producer (advanced))	The status codes which are considered a success response. The values are inclusive. Multiple ranges can be defined, separated by comma, e.g. 200-204,209,301-304. Each range must be a single number or from-to with the dash included.	200-299	String
skipRequestHeaders (producer (advanced))	Whether to skip mapping all the Camel headers as HTTP request headers. If there are no data from Camel headers needed to be included in the HTTP request then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	boolean
skipResponseHeaders (producer (advanced))	Whether to skip mapping all the HTTP response headers to Camel headers. If there are no data needed from HTTP headers then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	boolean
userAgent (producer (advanced))	To set a custom HTTP User-Agent request header.		String
clientBuilder (advanced)	Provide access to the http client request parameters used on new RequestConfig instances used by producers or consumers of this endpoint.		HttpClientBuilder

Name	Description	Default	Type
clientConnectionManager (advanced)	To use a custom <code>HttpClientConnectionManager</code> to manage connections.		<code>HttpClientConnectionManager</code>
connectionsPerRoute (advanced)	The maximum number of connections per route.	20	int
httpClient (advanced)	Sets a custom <code>HttpClient</code> to be used by the producer.		<code>HttpClient</code>
httpClientConfigurer (advanced)	Register a custom configuration strategy for new <code>HttpClient</code> instances created by producers or consumers such as to configure authentication mechanisms etc.		<code>HttpClientConfigurer</code>
httpClientOptions (advanced)	To configure the <code>HttpClient</code> using the key/values from the Map.		Map
httpClient (advanced)	To use a custom <code>HttpContext</code> instance.		<code>HttpContext</code>
maxTotalConnections (advanced)	The maximum number of connections.	200	int
useSystemProperties (advanced)	To use System Properties as fallback for configuration.	false	boolean
proxyAuthDomain (proxy)	Proxy authentication domain to use with NTLM.		String
proxyAuthHost (proxy)	Proxy authentication host.		String
proxyAuthMethod (proxy)	Proxy authentication method to use. Enum values: <ul style="list-style-type: none"> • Basic • Digest • NTLM 		String
proxyAuthNtHost (proxy)	Proxy authentication domain (workstation name) to use with NTLM.		String

Name	Description	Default	Type
proxyAuthPassword (proxy)	Proxy authentication password.		String
proxyAuthPort (proxy)	Proxy authentication port.		int
proxyAuthScheme (proxy)	Proxy authentication scheme to use. Enum values: <ul style="list-style-type: none"> • http • https 		String
proxyAuthUsername (proxy)	Proxy authentication username.		String
proxyHost (proxy)	Proxy hostname to use.		String
proxyPort (proxy)	Proxy port to use.		int
authDomain (security)	Authentication domain to use with NTLM.		String
authenticationPreemptive (security)	If this option is true, camel-http sends preemptive basic authentication to the server.	false	boolean
authHost (security)	Authentication host to use with NTLM.		String
authMethod (security)	Authentication methods allowed to use as a comma separated list of values Basic, Digest or NTLM.		String
authMethodPriority (security)	Which authentication method to prioritize to use, either as Basic, Digest or NTLM. Enum values: <ul style="list-style-type: none"> • Basic • Digest • NTLM 		String
authPassword (security)	Authentication password.		String

Name	Description	Default	Type
authUsername (security)	Authentication username.		String
sslContextParameters (security)	To configure security using SSLContextParameters. Important: Only one instance of org.apache.camel.util.jsse.SSLContextParameters is supported per HttpComponent. If you need to use 2 or more different instances, you need to define a new HttpComponent per instance you need.		SSLContextParameters
x509HostnameVerifier (security)	To use a custom X509HostnameVerifier such as DefaultHostnameVerifier or NoopHostnameVerifier.		HostnameVerifier

23.5. MESSAGE HEADERS

Name	Type	Description
Exchange.HTTP_URI	String	URI to call. Will override existing URI set directly on the endpoint. This uri is the uri of the http server to call. Its not the same as the Camel endpoint uri, where you can configure endpoint options such as security etc. This header does not support that, its only the uri of the http server.
Exchange.HTTP_PATH	String	Request URI's path, the header will be used to build the request URI with the HTTP_URI.
Exchange.HTTP_QUERY	String	URI parameters. Will override existing URI parameters set directly on the endpoint.
Exchange.HTTP_RESPONSE_CODE	int	The HTTP response code from the external server. Is 200 for OK.
Exchange.HTTP_RESPONSE_TEXT	String	The HTTP response text from the external server.
Exchange.HTTP_CHARACTER_ENCODING	String	Character encoding.
Exchange.CONTENT_TYPE	String	The HTTP content type. Is set on both the IN and OUT message to provide a content type, such as text/html .
Exchange.CONTENT_ENCODING	String	The HTTP content encoding. Is set on both the IN and OUT message to provide a content encoding, such as gzip .

23.6. MESSAGE BODY

Camel will store the HTTP response from the external server on the OUT body. All headers from the IN message will be copied to the OUT message, so headers are preserved during routing. Additionally Camel will add the HTTP response headers as well to the OUT message headers.

23.7. USING SYSTEM PROPERTIES

When setting `useSystemProperties` to true, the HTTP Client will look for the following System Properties and it will use it:

- `ssl.TrustManagerFactory.algorithm`
- `javax.net.ssl.trustStoreType`
- `javax.net.ssl.trustStore`
- `javax.net.ssl.trustStoreProvider`
- `javax.net.ssl.trustStorePassword`
- `java.home`
- `ssl.KeyManagerFactory.algorithm`
- `javax.net.ssl.keyStoreType`
- `javax.net.ssl.keyStore`
- `javax.net.ssl.keyStoreProvider`
- `javax.net.ssl.keyStorePassword`
- `http.proxyHost`
- `http.proxyPort`
- `http.nonProxyHosts`
- `http.keepAlive`
- `http.maxConnections`

23.8. RESPONSE CODE

Camel will handle according to the HTTP response code:

- Response code is in the range 100..299, Camel regards it as a success response.
- Response code is in the range 300..399, Camel regards it as a redirection response and will throw a **HttpOperationFailedException** with the information.
- Response code is 400+, Camel regards it as an external server failure and will throw a **HttpOperationFailedException** with the information.

throwExceptionOnFailure The option, **throwExceptionOnFailure**, can be set to **false** to prevent the

HttpOperationFailedException from being thrown for failed response codes. This allows you to get any response from the remote server.

There is a sample below demonstrating this.

23.9. EXCEPTIONS

HttpOperationFailedException exception contains the following information:

- The HTTP status code
- The HTTP status line (text of the status code)
- Redirect location, if server returned a redirect
- Response body as a **java.lang.String**, if server provided a body as response

23.10. WHICH HTTP METHOD WILL BE USED

The following algorithm is used to determine what HTTP method should be used:

1. Use method provided as endpoint configuration (**httpMethod**).
2. Use method provided in header (**Exchange.HTTP_METHOD**).
3. **GET** if query string is provided in header.
4. **GET** if endpoint is configured with a query string.
5. **POST** if there is data to send (body is not **null**).
6. **GET** otherwise.

23.11. HOW TO GET ACCESS TO HTTPSERVLETREQUEST AND HTTPSERVLETRESPONSE

You can get access to these two using the Camel type converter system using

```
HttpServletRequest request = exchange.getIn().getBody(HttpServletRequest.class);
HttpServletResponse response = exchange.getIn().getBody(HttpServletResponse.class);
```



NOTE

You can get the request and response not just from the processor after the camel-jetty or camel-cxf endpoint.

23.12. CONFIGURING URI TO CALL

You can set the HTTP producer's URI directly from the endpoint URI. In the route below, Camel will call out to the external server, **oldhost**, using HTTP.

```
from("direct:start")
    .to("http://oldhost");
```

And the equivalent Spring sample:

```
<camelContext xmlns="http://activemq.apache.org/camel/schema/spring">
  <route>
    <from uri="direct:start"/>
```

```
<to uri="http://oldhost"/>
</route>
</camelContext>
```

You can override the HTTP endpoint URI by adding a header with the key, **Exchange.HTTP_URI**, on the message.

```
from("direct:start")
  .setHeader(Exchange.HTTP_URI, constant("http://newhost"))
  .to("http://oldhost");
```

In the sample above Camel will call the <http://newhost/> despite the endpoint is configured with <http://oldhost/>.

If the http endpoint is working in bridge mode, it will ignore the message header of **Exchange.HTTP_URI**.

23.13. CONFIGURING URI PARAMETERS

The **http** producer supports URI parameters to be sent to the HTTP server. The URI parameters can either be set directly on the endpoint URI or as a header with the key **Exchange.HTTP_QUERY** on the message.

```
from("direct:start")
  .to("http://oldhost?order=123&detail=short");
```

Or options provided in a header:

```
from("direct:start")
  .setHeader(Exchange.HTTP_QUERY, constant("order=123&detail=short"))
  .to("http://oldhost");
```

23.14. HOW TO SET THE HTTP METHOD (GET/PATCH/POST/PUT/DELETE/HEAD/OPTIONS/TRACE) TO THE HTTP PRODUCER

The HTTP component provides a way to set the HTTP request method by setting the message header. Here is an example:

```
from("direct:start")
  .setHeader(Exchange.HTTP_METHOD,
    constant(org.apache.camel.component.http.HttpMethods.POST))
  .to("http://www.google.com")
  .to("mock:results");
```

The method can be written a bit shorter using the string constants:

```
.setHeader("CamelHttpMethod", constant("POST"))
```

And the equivalent Spring sample:

```
<camelContext xmlns="http://activemq.apache.org/camel/schema/spring">
  <route>
```

```

<from uri="direct:start"/>
<setHeader name="CamelHttpMethod">
  <constant>POST</constant>
</setHeader>
<to uri="http://www.google.com"/>
<to uri="mock:results"/>
</route>
</camelContext>

```

23.15. USING CLIENT TIMEOUT - SO_TIMEOUT

See the [HttpSOTimeoutTest](#) unit test.

23.16. CONFIGURING A PROXY

The HTTP component provides a way to configure a proxy.

```

from("direct:start")
.to("http://oldhost?proxyAuthHost=www.myproxy.com&proxyAuthPort=80");

```

There is also support for proxy authentication via the **proxyAuthUsername** and **proxyAuthPassword** options.

23.16.1. Using proxy settings outside of URI

To avoid System properties conflicts, you can set proxy configuration only from the CamelContext or URI.

Java DSL :

```

context.getGlobalOptions().put("http.proxyHost", "172.168.18.9");
context.getGlobalOptions().put("http.proxyPort", "8080");

```

Spring XML

```

<camelContext>
  <properties>
    <property key="http.proxyHost" value="172.168.18.9"/>
    <property key="http.proxyPort" value="8080"/>
  </properties>
</camelContext>

```

Camel will first set the settings from Java System or CamelContext Properties and then the endpoint proxy options if provided.

So you can override the system properties with the endpoint options.

There is also a **http.proxyScheme** property you can set to explicit configure the scheme to use.

23.17. CONFIGURING CHARSET

If you are using **POST** to send data you can configure the **charset** using the **Exchange** property:

```

exchange.setProperty(Exchange.CHARSET_NAME, "ISO-8859-1");

```

23.17.1. Sample with scheduled poll

This sample polls the Google homepage every 10 seconds and write the page to the file **message.html**:

```
from("timer://foo?fixedRate=true&delay=0&period=10000")
  .to("http://www.google.com")
  .setHeader(FileComponent.HEADER_FILE_NAME, "message.html")
  .to("file:target/google");
```

23.17.2. URI Parameters from the endpoint URI

In this sample we have the complete URI endpoint that is just what you would have typed in a web browser. Multiple URI parameters can of course be set using the **&** character as separator, just as you would in the web browser. Camel does no tricks here.

```
// we query for Camel at the Google page
template.sendBody("http://www.google.com/search?q=Camel", null);
```

23.17.3. URI Parameters from the Message

```
Map headers = new HashMap();
headers.put(Exchange.HTTP_QUERY, "q=Camel&lr=lang_en");
// we query for Camel and English language at Google
template.sendBody("http://www.google.com/search", null, headers);
```

In the header value above notice that it should **not** be prefixed with **?** and you can separate parameters as usual with the **&** char.

23.17.4. Getting the Response Code

You can get the HTTP response code from the HTTP component by getting the value from the Out message header with **Exchange.HTTP_RESPONSE_CODE**.

```
Exchange exchange = template.send("http://www.google.com/search", new Processor() {
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setHeader(Exchange.HTTP_QUERY, constant("hl=en&q=activemq"));
    }
});
Message out = exchange.getOut();
int responseCode = out.getHeader(Exchange.HTTP_RESPONSE_CODE, Integer.class);
```

23.18. DISABLING COOKIES

To disable cookies you can set the HTTP Client to ignore cookies by adding this URI option: **httpClient.cookieSpec=ignoreCookies**

23.19. BASIC AUTH WITH THE STREAMING MESSAGE BODY

In order to avoid the **NonRepeatableRequestException**, you need to do the Preemptive Basic Authentication by adding the option: **authenticationPreemptive=true**

23.20. ADVANCED USAGE

If you need more control over the HTTP producer you should use the **HttpComponent** where you can set various classes to give you custom behavior.

23.20.1. Setting up SSL for HTTP Client

Using the JSSE Configuration Utility

The HTTP component supports SSL/TLS configuration through the [Camel JSSE Configuration Utility](#). This utility greatly decreases the amount of component specific code you need to write and is configurable at the endpoint and component levels. The following examples demonstrate how to use the utility with the HTTP component.

Programmatic configuration of the component

```

KeyStoreParameters ksp = new KeyStoreParameters();
ksp.setResource("/users/home/server/keystore.jks");
ksp.setPassword("keystorePassword");

KeyManagersParameters kmp = new KeyManagersParameters();
kmp.setKeyStore(ksp);
kmp.setKeyPassword("keyPassword");

SSLContextParameters scp = new SSLContextParameters();
scp.setKeyManagers(kmp);

HttpComponent httpComponent = getContext().getComponent("https", HttpComponent.class);
httpComponent.setSslContextParameters(scp);

```

Spring DSL based configuration of endpoint

```

<camel:sslContextParameters
  id="sslContextParameters">
  <camel:keyManagers
    keyPassword="keyPassword">
    <camel:keyStore
      resource="/users/home/server/keystore.jks"
      password="keystorePassword"/>
    </camel:keyManagers>
  </camel:sslContextParameters>

  <to uri="https://127.0.0.1/mail/?sslContextParameters=#sslContextParameters"/>

```

Configuring Apache HTTP Client Directly

Basically camel-http component is built on the top of [Apache HttpClient](#). Please refer to [SSL/TLS customization](#) for details or have a look into the

org.apache.camel.component.http.HttpsServerTestSupport unit test base class.

You can also implement a custom **org.apache.camel.component.http.HttpClientConfigurer** to do some configuration on the http client if you need full control of it.

However if you *just* want to specify the keystore and truststore you can do this with Apache HTTP **HttpClientConfigurer**, for example:

```

KeyStore keystore = ...;
KeyStore truststore = ...;

SchemeRegistry registry = new SchemeRegistry();
registry.register(new Scheme("https", 443, new SSLSocketFactory(keystore, "mypassword",
truststore)));

```

And then you need to create a class that implements **HttpClientConfigurer**, and registers https protocol providing a keystore or truststore per example above. Then, from your camel route builder class you can hook it up like so:

```

HttpClientComponent httpComponent = getContext().getComponent("http", HttpClientComponent.class);
httpComponent.setHttpClientConfigurer(new MyHttpClientConfigurer());

```

If you are doing this using the Spring DSL, you can specify your **HttpClientConfigurer** using the URI. For example:

```

<bean id="myHttpClientConfigurer"
class="my.https.HttpClientConfigurer">
</bean>

<to uri="https://myhostname.com:443/myURL?httpClientConfigurer=myHttpClientConfigurer"/>

```

As long as you implement the `HttpClientConfigurer` and configure your keystore and truststore as described above, it will work fine.

Using HTTPS to authenticate gotchas

An end user reported that he had problem with authenticating with HTTPS. The problem was eventually resolved by providing a custom configured **org.apache.http.protocol.HttpContext**:

- 1. Create a (Spring) factory for `HttpContext`s:

```

public class HttpContextFactory {

    private String httpHost = "localhost";
    private String httpPort = "9001";

    private BasicHttpContext httpContext = new BasicHttpContext();
    private BasicAuthCache authCache = new BasicAuthCache();
    private BasicScheme basicAuth = new BasicScheme();

    public HttpContext getObject() {
        authCache.put(new HttpHost(httpHost, httpPort), basicAuth);

        httpContext.setAttribute(ClientContext.AUTH_CACHE, authCache);

        return httpContext;
    }

    // getter and setter
}

```

- 2. Declare an `HttpContext` in the Spring application context file:

```
<bean id="myHttpContext" factory-bean="httpContextFactory" factory-method="getObject"/>
```

- 3. Reference the context in the http URL:

```
<to uri="https://myhostname.com:443/myURL?httpContext=myHttpContext"/>
```

Using different SSLContextParameters

The [HTTP](#) component only support one instance of **org.apache.camel.support.jsse.SSLContextParameters** per component. If you need to use 2 or more different instances, then you need to setup multiple [HTTP](#) components as shown below. Where we have 2 components, each using their own instance of **sslContextParameters** property.

```
<bean id="http-foo" class="org.apache.camel.component.http.HttpComponent">
  <property name="sslContextParameters" ref="sslContextParams1"/>
  <property name="x509HostnameVerifier" ref="hostnameVerifier"/>
</bean>

<bean id="http-bar" class="org.apache.camel.component.http.HttpComponent">
  <property name="sslContextParameters" ref="sslContextParams2"/>
  <property name="x509HostnameVerifier" ref="hostnameVerifier"/>
</bean>
```

23.21. SPRING BOOT AUTO-CONFIGURATION

When using http with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-http-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 38 options, which are listed below.

Name	Description	Default	Type
camel.component.http.allow-java-serialized-object	Whether to allow java serialization when a request uses context-type=application/x-java-serialized-object. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	Boolean
camel.component.http.auth-caching-disabled	Disables authentication scheme caching.	false	Boolean

Name	Description	Default	Type
<code>camel.component.http.automatic-retries-disabled</code>	Disables automatic request recovery and re-execution.	false	Boolean
<code>camel.component.http.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.http.client-connection-manager</code>	To use a custom and shared <code>HttpClientConnectionManager</code> to manage connections. If this has been configured then this is always used for all endpoints created by this component. The option is a <code>org.apache.http.conn.HttpClientConnectionManager</code> type.		<code>HttpClientConnectionManager</code>
<code>camel.component.http.connect-timeout</code>	Determines the timeout in milliseconds until a connection is established. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	Integer
<code>camel.component.http.connection-request-timeout</code>	The timeout in milliseconds used when requesting a connection from the connection manager. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	Integer
<code>camel.component.http.connection-state-disabled</code>	Disables connection state tracking.	false	Boolean
<code>camel.component.http.connection-time-to-live</code>	The time for connection to live, the time unit is millisecond, the default value is always keep alive.		Long
<code>camel.component.http.connections-per-route</code>	The maximum number of connections per route.	20	Integer

Name	Description	Default	Type
<code>camel.component.http.content-compression-disabled</code>	Disables automatic content decompression.	false	Boolean
<code>camel.component.http.cookie-management-disabled</code>	Disables state (cookie) management.	false	Boolean
<code>camel.component.http.cookie-store</code>	To use a custom <code>org.apache.http.client.CookieStore</code> . By default the <code>org.apache.http.impl.client.BasicCookieStore</code> is used which is an in-memory only cookie store. Notice if <code>bridgeEndpoint=true</code> then the cookie store is forced to be a noop cookie store as cookie shouldn't be stored as we are just bridging (eg acting as a proxy). The option is a <code>org.apache.http.client.CookieStore</code> type.		<code>CookieStore</code>
<code>camel.component.http.copy-headers</code>	If this option is true then IN exchange headers will be copied to OUT exchange headers according to copy strategy. Setting this to false, allows to only include the headers from the HTTP response (not propagating IN headers).	true	Boolean
<code>camel.component.http.default-user-agent-disabled</code>	Disables the default user agent set by this builder if none has been provided by the user.	false	Boolean
<code>camel.component.http.enabled</code>	Whether to enable auto configuration of the http component. This is enabled by default.		Boolean
<code>camel.component.http.header-filter-strategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		<code>HeaderFilterStrategy</code>
<code>camel.component.http.http-binding</code>	To use a custom <code>HttpBinding</code> to control the mapping between Camel message and <code>HttpClient</code> . The option is a <code>org.apache.camel.http.common.HttpBinding</code> type.		<code>HttpBinding</code>

Name	Description	Default	Type
<code>camel.component.http.http-client-configurer</code>	To use the custom <code>HttpClientConfigurer</code> to perform configuration of the <code>HttpClient</code> that will be used. The option is a <code>org.apache.camel.component.http.HttpClientConfigurer</code> type.		<code>HttpClientConfigurer</code>
<code>camel.component.http.http-configuration</code>	To use the shared <code>HttpConfiguration</code> as base configuration. The option is a <code>org.apache.camel.http.common.HttpConfiguration</code> type.		<code>HttpConfiguration</code>
<code>camel.component.http.http-context</code>	To use a custom <code>org.apache.http.protocol.HttpContext</code> when executing requests. The option is a <code>org.apache.http.protocol.HttpContext</code> type.		<code>HttpContext</code>
<code>camel.component.http.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.http.max-total-connections</code>	The maximum number of connections.	200	Integer
<code>camel.component.http.proxy-auth-domain</code>	Proxy authentication domain to use.		String
<code>camel.component.http.proxy-auth-host</code>	Proxy authentication host.		String
<code>camel.component.http.proxy-auth-method</code>	Proxy authentication method to use.		String
<code>camel.component.http.proxy-auth-nt-host</code>	Proxy authentication domain (workstation name) to use with NTLM.		String

Name	Description	Default	Type
<code>camel.component.http.proxy-auth-password</code>	Proxy authentication password.		String
<code>camel.component.http.proxy-auth-port</code>	Proxy authentication port.		Integer
<code>camel.component.http.proxy-auth-username</code>	Proxy authentication username.		String
<code>camel.component.http.redirect-handling-disabled</code>	Disables automatic redirect handling.	false	Boolean
<code>camel.component.http.response-payload-streaming-threshold</code>	This threshold in bytes controls whether the response payload should be stored in memory as a byte array or be streaming based. Set this to -1 to always use streaming mode.	8192	Integer
<code>camel.component.http.skip-request-headers</code>	Whether to skip mapping all the Camel headers as HTTP request headers. If there are no data from Camel headers needed to be included in the HTTP request then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	Boolean
<code>camel.component.http.skip-response-headers</code>	Whether to skip mapping all the HTTP response headers to Camel headers. If there are no data needed from HTTP headers then this can avoid parsing overhead with many object allocations for the JVM garbage collector.	false	Boolean
<code>camel.component.http.socket-timeout</code>	Defines the socket timeout in milliseconds, which is the timeout for waiting for data or, put differently, a maximum period inactivity between two consecutive data packets). A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	Integer

Name	Description	Default	Type
camel.component.http.ssl-context-parameters	To configure security using SSLContextParameters. Important: Only one instance of org.apache.camel.support.jsse.SSLContextParameters is supported per HttpComponent. If you need to use 2 or more different instances, you need to define a new HttpComponent per instance you need. The option is a org.apache.camel.support.jsse.SSLContextParameters type.		SSLContextParameters
camel.component.http.use-global-ssl-context-parameters	Enable usage of global SSL context parameters.	false	Boolean
camel.component.http.x509-hostname-verifier	To use a custom X509HostnameVerifier such as DefaultHostnameVerifier or NoopHostnameVerifier. The option is a javax.net.ssl.HostnameVerifier type.		HostnameVerifier

CHAPTER 24. INFINISPAN

Both producer and consumer are supported

This component allows you to interact with [Infinispan](#) distributed data grid/cache using the Hot Rod protocol. Infinispan is an extremely scalable, highly available key/value data store and data grid platform written in Java.

If you use Maven, you must add the following dependency to your **pom.xml**:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-infinispan</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

24.1. URI FORMAT

```
infinispan://cacheName?[options]
```

The producer allows sending messages to a remote cache using the HotRod protocol. The consumer allows listening for events from a remote cache using the HotRod protocol.

24.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

24.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

24.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

24.3. COMPONENT OPTIONS

The Infinispan component supports 26 options, which are listed below.

Name	Description	Default	Type
configuration (common)	Component configuration.		InfinispanRemote Configuration
hosts (common)	Specifies the host of the cache on Infinispan instance.		String
queryBuilder (common)	Specifies the query builder.		InfinispanQueryBuilder
secure (common)	Define if we are connecting to a secured Infinispan instance.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
customListener (consumer)	Returns the custom listener in use, if provided.		InfinispanRemote CustomListener
eventTypes (consumer)	Specifies the set of event types to register by the consumer. Multiple event can be separated by comma. The possible event types are: CLIENT_CACHE_ENTRY_CREATED, CLIENT_CACHE_ENTRY_MODIFIED, CLIENT_CACHE_ENTRY_REMOVED, CLIENT_CACHE_ENTRY_EXPIRED, CLIENT_CACHE_FAILOVER.		String
defaultValue (producer)	Set a specific default value for some producer operations.		Object
key (producer)	Set a specific key for producer operations.		Object

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
oldValue (producer)	Set a specific old value for some producer operations.		Object

Name	Description	Default	Type
operation (producer)	<p>The operation to perform.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● PUT ● PUTASYNC ● PUTALL ● PUTALLASYNC ● PUTIFABSENT ● PUTIFABSENTASYNC ● GET ● GETORDEFAULT ● CONTAINSKEY ● CONTAINSVALUE ● REMOVE ● REMOVEASYNC ● REPLACE ● REPLACEASYNC ● SIZE ● CLEAR ● CLEARASYNC ● QUERY ● STATS ● COMPUTE ● COMPUTEASYNC 	PUT	InfinispanOperation
value (producer)	Set a specific value for producer operations.		Object
password (security)	Define the password to access the infinispan instance.		String
saslMechanism (security)	Define the SASL Mechanism to access the infinispan instance.		String

Name	Description	Default	Type
securityRealm (security)	Define the security realm to access the infinispn instance.		String
securityServerName (security)	Define the security server name to access the infinispn instance.		String
username (security)	Define the username to access the infinispn instance.		String
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
cacheContainer (advanced)	Autowired Specifies the cache Container to connect.		RemoteCacheManager
cacheContainerConfiguration (advanced)	Autowired The CacheContainer configuration. Used if the cacheContainer is not defined.		Configuration
configurationProperties (advanced)	Implementation specific properties for the CacheManager.		Map
configurationUri (advanced)	An implementation specific URI for the CacheManager.		String
flags (advanced)	A comma separated list of org.infinispn.client.hotrod.Flag to be applied by default on each cache invocation.		String
remappingFunction (advanced)	Set a specific remappingFunction to use in a compute operation.		BiFunction
resultHeader (advanced)	Store the operation result in a header instead of the message body. By default, resultHeader == null and the query result is stored in the message body, any existing content in the message body is discarded. If resultHeader is set, the value is used as the name of the header to store the query result and the original message body is preserved. This value can be overridden by an in message header named: CamelInfinispnOperationResultHeader.		String

24.4. ENDPOINT OPTIONS

The Infinispan endpoint is configured using URI syntax:

```
infinispan:cacheName
```

with the following path and query parameters:

24.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
cacheName (common)	Required The name of the cache to use. Use current to use the existing cache name from the currently configured cached manager. Or use default for the default cache manager name.		String

24.4.2. Query Parameters (26 parameters)

Name	Description	Default	Type
hosts (common)	Specifies the host of the cache on Infinispan instance.		String
queryBuilder (common)	Specifies the query builder.		InfinispanQueryBuilder
secure (common)	Define if we are connecting to a secured Infinispan instance.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
customListener (consumer)	Returns the custom listener in use, if provided.		InfinispanRemoteCustomListener

Name	Description	Default	Type
eventTypes (consumer)	Specifies the set of event types to register by the consumer. Multiple event can be separated by comma. The possible event types are: CLIENT_CACHE_ENTRY_CREATED, CLIENT_CACHE_ENTRY_MODIFIED, CLIENT_CACHE_ENTRY_REMOVED, CLIENT_CACHE_ENTRY_EXPIRED, CLIENT_CACHE_FAILOVER.		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
defaultValue (producer)	Set a specific default value for some producer operations.		Object
key (producer)	Set a specific key for producer operations.		Object
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
oldValue (producer)	Set a specific old value for some producer operations.		Object

Name	Description	Default	Type
operation (producer)	<p>The operation to perform.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● PUT ● PUTASYNC ● PUTALL ● PUTALLASYNC ● PUTIFABSENT ● PUTIFABSENTASYNC ● GET ● GETORDEFAULT ● CONTAINSKEY ● CONTAINSVALUE ● REMOVE ● REMOVEASYNC ● REPLACE ● REPLACEASYNC ● SIZE ● CLEAR ● CLEARASYNC ● QUERY ● STATS ● COMPUTE ● COMPUTEASYNC 	PUT	InfinispanOperation
value (producer)	Set a specific value for producer operations.		Object
password (security)	Define the password to access the infinispan instance.		String
saslMechanism (security)	Define the SASL Mechanism to access the infinispan instance.		String

Name	Description	Default	Type
securityRealm (security)	Define the security realm to access the infinispn instance.		String
securityServerName (security)	Define the security server name to access the infinispn instance.		String
username (security)	Define the username to access the infinispn instance.		String
cacheContainer (advanced)	Autowired Specifies the cache Container to connect.		RemoteCacheManager
cacheContainerConfiguration (advanced)	Autowired The CacheContainer configuration. Used if the cacheContainer is not defined.		Configuration
configurationProperties (advanced)	Implementation specific properties for the CacheManager.		Map
configurationUri (advanced)	An implementation specific URI for the CacheManager.		String
flags (advanced)	A comma separated list of org.infinispn.client.hotrod.Flag to be applied by default on each cache invocation.		String
remappingFunction (advanced)	Set a specific remappingFunction to use in a compute operation.		BiFunction
resultHeader (advanced)	Store the operation result in a header instead of the message body. By default, resultHeader == null and the query result is stored in the message body, any existing content in the message body is discarded. If resultHeader is set, the value is used as the name of the header to store the query result and the original message body is preserved. This value can be overridden by an in message header named: CamelInfinispnOperationResultHeader.		String

24.5. CAMEL OPERATIONS

This section lists all available operations, along with their header information.

Table 24.1. Table 1. Put Operations

Operation Name	Description
InfinispanOperation.PUT	Puts a key/value pair in the cache, optionally with expiration
InfinispanOperation.PUTASYNC	Asynchronously puts a key/value pair in the cache, optionally with expiration
InfinispanOperation.PUTIFABSENT	Puts a key/value pair in the cache if it did not exist, optionally with expiration
InfinispanOperation.PUTIFABSENTASYNC	Asynchronously puts a key/value pair in the cache if it did not exist, optionally with expiration

- **Required Headers:**
 - CamelInfinispanKey
 - CamelInfinispanValue
- **Optional Headers:**
 - CamelInfinispanLifespanTime
 - CamelInfinispanLifespanTimeUnit
 - CamelInfinispanMaxIdleTime
 - CamelInfinispanMaxIdleTimeUnit
- **Result Header:**
 - CamelInfinispanOperationResult

Table 24.2. Table 2. Put All Operations

Operation Name	Description
InfinispanOperation.PUTALL	Adds multiple entries to a cache, optionally with expiration
CamelInfinispanOperation.PUTALLASYNC	Asynchronously adds multiple entries to a cache, optionally with expiration

- **Required Headers:**
 - CamelInfinispanMap
- **Optional Headers:**
 - CamelInfinispanLifespanTime
 - CamelInfinispanLifespanTimeUnit
 - CamelInfinispanMaxIdleTime

- `CamelInfinispanMaxIdleTimeUnit`

Table 24.3. Table 3. Get Operations

Operation Name	Description
<code>InfinispanOperation.GET</code>	Retrieves the value associated with a specific key from the cache
<code>InfinispanOperation.GETORDEFAULT</code>	Retrieves the value, or default value, associated with a specific key from the cache

- **Required Headers:**
 - `CamelInfinispanKey`

Table 24.4. Table 4. Contains Key Operation

Operation Name	Description
<code>InfinispanOperation.CONTAINSKEY</code>	Determines whether a cache contains a specific key

- **Required Headers**
 - `CamelInfinispanKey`
- **Result Header**
 - `CamelInfinispanOperationResult`

Table 24.5. Table 5. Contains Value Operation

Operation Name	Description
<code>InfinispanOperation.CONTAINSVALUE</code>	Determines whether a cache contains a specific value

- **Required Headers:**
 - `CamelInfinispanKey`

Table 24.6. Table 6. Remove Operations

Operation Name	Description
<code>InfinispanOperation.REMOVE</code>	Removes an entry from a cache, optionally only if the value matches a given one
<code>InfinispanOperation.REMOVEASYNC</code>	Asynchronously removes an entry from a cache, optionally only if the value matches a given one

- **Required Headers:**

- CamelInfinispanKey
- **Optional Headers:**
 - CamelInfinispanValue
- **Result Header:**
 - CamelInfinispanOperationResult

Table 24.7. Table 7. Replace Operations

Operation Name	Description
InfinispanOperation.REPLACE	Conditionally replaces an entry in the cache, optionally with expiration
InfinispanOperation.REPLACEASYNC	Asynchronously conditionally replaces an entry in the cache, optionally with expiration

- **Required Headers:**
 - CamelInfinispanKey
 - CamelInfinispanValue
 - CamelInfinispanOldValue
- **Optional Headers:**
 - CamelInfinispanLifespanTime
 - CamelInfinispanLifespanTimeUnit
 - CamelInfinispanMaxIdleTime
 - CamelInfinispanMaxIdleTimeUnit
- **Result Header:**
 - CamelInfinispanOperationResult

Table 24.8. Table 8. Clear Operations

Operation Name	Description
InfinispanOperation.CLEAR	Clears the cache
InfinispanOperation.CLEARASYNC	Asynchronously clears the cache

Table 24.9. Table 9. Size Operation

Operation Name	Description
InfinispanOperation.SIZE	Returns the number of entries in the cache

- **Result Header**
 - CamelInfinispanOperationResult

Table 24.10. Table 10. Stats Operation

Operation Name	Description
InfinispanOperation.STATS	Returns statistics about the cache

- **Result Header:**
 - CamelInfinispanOperationResult

Table 24.11. Table 11. Query Operation

Operation Name	Description
InfinispanOperation.QUERY	Executes a query on the cache

- **Required Headers:**
 - CamelInfinispanQueryBuilder
- **Result Header:**
 - CamelInfinispanOperationResult

**NOTE**

Write methods like `put(key, value)` and `remove(key)` do not return the previous value by default.

24.6. MESSAGE HEADERS

Name	Default Value	Type	Context	Description
CamelInfinispanCacheName	null	String	Shared	The cache participating in the operation or event.
CamelInfinispanOperation	PUT	InfinispanOperation	Producer	The operation to perform.

Name	Default Value	Type	Context	Description
CamelInfinispanMap	null	Map	Producer	A Map to use in case of CamelInfinispanOperationPutAll operation
CamelInfinispanKey	null	Object	Shared	The key to perform the operation to or the key generating the event.
CamelInfinispanValue	null	Object	Producer	The value to use for the operation.
CamelInfinispanEventType	null	String	Consumer	The type of the received event.
CamelInfinispanLifespanTime	null	long	Producer	The Lifespan time of a value inside the cache. Negative values are interpreted as infinity.
CamelInfinispanTimeUnit	null	String	Producer	The Time Unit of an entry Lifespan Time.
CamelInfinispanMaxIdleTime	null	long	Producer	The maximum amount of time an entry is allowed to be idle for before it is considered as expired.
CamelInfinispanMaxIdleTimeUnit	null	String	Producer	The Time Unit of an entry Max Idle Time.
CamelInfinispanQueryBuilder	null	InfinispanQueryBuilder	Producer	The QueryBuilde to use for QUERY command, if not present the command defaults to InifinispanConfiguration's one
CamelInfinispanOperationResultHeader	null	String	Producer	Store the operation result in a header instead of the message body

24.7. EXAMPLES

- Put a key/value into a named cache:

```
from("direct:start")
  .setHeader(InfinispanConstants.OPERATION).constant(InfinispanOperation.PUT) (1)
  .setHeader(InfinispanConstants.KEY).constant("123") (2)
  .to("infinispan:myCacheName&cacheContainer=#cacheContainer"); (3)
```

Where,

- 1 - Set the operation to perform
- 2 - Set the key used to identify the element in the cache
- 3 - Use the configured cache manager **cacheContainer** from the registry to put an element to the cache named **myCacheName**

It is possible to configure the lifetime and/or the idle time before the entry expires and gets evicted from the cache, as example:

```
from("direct:start")
  .setHeader(InfinispanConstants.OPERATION).constant(InfinispanOperation.GET)
  .setHeader(InfinispanConstants.KEY).constant("123")
  .setHeader(InfinispanConstants.LIFESPAN_TIME).constant(100L) (1)

  .setHeader(InfinispanConstants.LIFESPAN_TIME_UNIT.constant(TimeUnit.MILLISECONDS.t
oString())) (2)
  .to("infinispan:myCacheName");
```

where,

- 1 - Set the lifespan of the entry
- 2 - Set the time unit for the lifespan

Queries

```
from("direct:start")
  .setHeader(InfinispanConstants.OPERATION, InfinispanConstants.QUERY)
  .setHeader(InfinispanConstants.QUERY_BUILDER, new InfinispanQueryBuilder() {
    @Override
    public Query build(QueryFactory<Query> qf) {
      return qf.from(User.class).having("name").like("%abc%").build();
    }
  })
  .to("infinispan:myCacheName?cacheContainer=#cacheManager") ;
```



NOTE

The .proto descriptors for domain objects must be registered with the remote Data Grid server, see [Remote Query Example](#) in the official Infinispan documentation.

Custom Listeners

```
from("infinispan://?cacheContainer=#cacheManager&customListener=#myCustomListener")
  .to("mock:result");
```

The instance of **myCustomListener** must exist and Camel should be able to look it up from the **Registry**. Users are encouraged to extend the **org.apache.camel.component.infinispan.remote.InfinispanRemoteCustomListener** class and annotate the resulting class with **@ClientListener** which can be found in package **org.infinispan.client.hotrod.annotation**.

24.8. USING THE INFINISPAN BASED IDEMPOTENT REPOSITORY

In this section we will use the Infinispan based idempotent repository.

Java Example

```
InfinispanRemoteConfiguration conf = new InfinispanRemoteConfiguration(); (1)
conf.setHosts("localhost:11222")

InfinispanRemoteldempotentRepository repo = new
InfinispanRemoteldempotentRepository("idempotent"); (2)
repo.setConfiguration(conf);

context.addRoutes(new RouteBuilder() {
    @Override
    public void configure() {
        from("direct:start")
            .idempotentConsumer(header("MessageID"), repo) (3)
            .to("mock:result");
    }
});
```

where,

- 1 - Configure the cache
- 2 - Configure the repository bean
- 3 - Set the repository to the route

XML Example

```
<bean id="infinispanRepo"
class="org.apache.camel.component.infinispan.remote.InfinispanRemoteldempotentRepository"
destroy-method="stop">
    <constructor-arg value="idempotent"/> (1)
    <property name="configuration"> (2)
        <bean class="org.apache.camel.component.infinispan.remote.InfinispanRemoteConfiguration">
            <property name="hosts" value="localhost:11222"/>
        </bean>
    </property>
</bean>

<camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start" />
        <idempotentConsumer messageIdRepositoryRef="infinispanRepo"> (3)
            <header>MessageID</header>
            <to uri="mock:result" />
        </idempotentConsumer>
    </route>
</camelContext>
```

where,

- 1 - Set the name of the cache that will be used by the repository
- 2 - Configure the repository bean

- 3 - Set the repository to the route

24.9. USING THE INFINISPAN BASED AGGREGATION REPOSITORY

In this section we will use the Infinispan based aggregation repository.

Java Example

```
InfinispanRemoteConfiguration conf = new InfinispanRemoteConfiguration(); (1)
conf.setHosts("localhost:11222")

InfinispanRemoteAggregationRepository repo = new InfinispanRemoteAggregationRepository(); (2)
repo.setCacheName("aggregation");
repo.setConfiguration(conf);

context.addRoutes(new RouteBuilder() {
    @Override
    public void configure() {
        from("direct:start")
            .aggregate(header("MessageID"))
            .completionSize(3)
            .aggregationRepository(repo) (3)
            .aggregationStrategyRef("myStrategy")
            .to("mock:result");
    }
});
```

where,

- 1 - Configure the cache
- 2 - Create the repository bean
- 3 - Set the repository to the route

XML Example

```
<bean id="infinispanRepo"
class="org.apache.camel.component.infinispan.remote.InfinispanRemoteAggregationRepository"
destroy-method="stop">
    <constructor-arg value="aggregation"/> (1)
    <property name="configuration"> (2)
        <bean class="org.apache.camel.component.infinispan.remote.InfinispanRemoteConfiguration">
            <property name="hosts" value="localhost:11222"/>
        </bean>
    </property>
</bean>

<camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start" />
        <aggregate strategyRef="myStrategy"
            completionSize="3"
            aggregationRepositoryRef="infinispanRepo"> (3)
            <correlationExpression>
```

```

    <header>MessageID</header>
  </correlationExpression>
  <to uri="mock:result"/>
</aggregate>
</route>
</camelContext>

```

where,

- 1 - Set the name of the cache that will be used by the repository
- 2 - Configure the repository bean
- 3 - Set the repository to the route



NOTE

With the release of Infinispan 11, it is required to set the encoding configuration on any cache created. This is critical for consuming events too. For more information have a look at [Data Encoding and MediaTypes](#) in the official Infinispan documentation.

24.10. SPRING BOOT AUTO-CONFIGURATION

When using infinispan with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-infinispan-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 23 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.infinispan.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.infinispan.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.infinispan.cache-container</code>	Specifies the cache Container to connect. The option is a <code>org.infinispan.client.hotrod.RemoteCacheManager</code> type.		RemoteCacheManager
<code>camel.component.infinispan.cache-container-configuration</code>	The CacheContainer configuration. Used if the cacheContainer is not defined. The option is a <code>org.infinispan.client.hotrod.configuration.Configuration</code> type.		Configuration
<code>camel.component.infinispan.configuration</code>	Component configuration. The option is a <code>org.apache.camel.component.infinispan.remote.InfinispanRemoteConfiguration</code> type.		InfinispanRemoteConfiguration
<code>camel.component.infinispan.configuration-properties</code>	Implementation specific properties for the CacheManager.		Map
<code>camel.component.infinispan.configuration-uri</code>	An implementation specific URI for the CacheManager.		String
<code>camel.component.infinispan.custom-listener</code>	Returns the custom listener in use, if provided. The option is a <code>org.apache.camel.component.infinispan.remote.InfinispanRemoteCustomListener</code> type.		InfinispanRemoteCustomListener
<code>camel.component.infinispan.enabled</code>	Whether to enable auto configuration of the infinispan component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.infinispan.event-types</code>	Specifies the set of event types to register by the consumer. Multiple event can be separated by comma. The possible event types are: CLIENT_CACHE_ENTRY_CREATED, CLIENT_CACHE_ENTRY_MODIFIED, CLIENT_CACHE_ENTRY_REMOVED, CLIENT_CACHE_ENTRY_EXPIRED, CLIENT_CACHE_FAILOVER.		String
<code>camel.component.infinispan.flags</code>	A comma separated list of org.infinispan.client.hotrod.Flag to be applied by default on each cache invocation.		String
<code>camel.component.infinispan.hosts</code>	Specifies the host of the cache on Infinispan instance.		String
<code>camel.component.infinispan.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.infinispan.operation</code>	The operation to perform.		InfinispanOperation
<code>camel.component.infinispan.password</code>	Define the password to access the infinispan instance.		String
<code>camel.component.infinispan.query-builder</code>	Specifies the query builder. The option is a org.apache.camel.component.infinispan.InfinispanQueryBuilder type.		InfinispanQueryBuilder
<code>camel.component.infinispan.remapping-function</code>	Set a specific remappingFunction to use in a compute operation. The option is a java.util.function.BiFunction type.		BiFunction

Name	Description	Default	Type
camel.component.infinispan.result-header	Store the operation result in a header instead of the message body. By default, resultHeader == null and the query result is stored in the message body, any existing content in the message body is discarded. If resultHeader is set, the value is used as the name of the header to store the query result and the original message body is preserved. This value can be overridden by an in message header named: CamelInfinispanOperationResultHeader.		String
camel.component.infinispan.sasl-mechanism	Define the SASL Mechanism to access the infinispan instance.		String
camel.component.infinispan.secure	Define if we are connecting to a secured Infinispan instance.	false	Boolean
camel.component.infinispan.security-realm	Define the security realm to access the infinispan instance.		String
camel.component.infinispan.security-server-name	Define the security server name to access the infinispan instance.		String
camel.component.infinispan.username	Define the username to access the infinispan instance.		String

CHAPTER 25. JIRA

Both producer and consumer are supported

The JIRA component interacts with the JIRA API by encapsulating Atlassian's [REST Java Client for JIRA](#). It currently provides polling for new issues and new comments. It is also able to create new issues, add comments, change issues, add/remove watchers, add attachment and transition the state of an issue.

Rather than webhooks, this endpoint relies on simple polling. Reasons include:

- Concern for reliability/stability
- The types of payloads we're polling aren't typically large (plus, paging is available in the API)
- The need to support apps running somewhere not publicly accessible where a webhook would fail

Note that the JIRA API is fairly expansive. Therefore, this component could be easily expanded to provide additional interactions.

Maven users will need to add the following dependency to their pom.xml for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jira</artifactId>
  <version>${camel-version}</version>
</dependency>
```

25.1. URI FORMAT

```
jira://type[?options]
```

The Jira type accepts the following operations:

For consumers:

- newIssues: retrieve only new issues after the route is started
- newComments: retrieve only new comments after the route is started
- watchUpdates: retrieve only updated fields/issues based on provided jql

For producers:

- addIssue: add an issue
- addComment: add a comment on a given issue
- attach: add an attachment on a given issue
- deleteIssue: delete a given issue
- updateIssue: update fields of a given issue

- `transitionIssue`: transition a status of a given issue
- `watchers`: add/remove watchers of a given issue

As Jira is fully customizable, you must assure the fields IDs exists for the project and workflow, as they can change between different Jira servers.

25.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

25.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (`application.properties|yaml`), or directly with Java code.

25.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

25.3. COMPONENT OPTIONS

The Jira component supports 12 options, which are listed below.

Name	Description	Default	Type
<code>delay</code> (common)	Time in milliseconds to elapse for the next poll.	6000	Integer
<code>jiraUrl</code> (common)	Required The Jira server url, example: .		String

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
configuration (advanced)	To use a shared base jira configuration.		JiraConfiguration
accessToken (security)	(OAuth only) The access token generated by the Jira server.		String
consumerKey (security)	(OAuth only) The consumer key from Jira settings.		String
password (security)	(Basic authentication only) The password to authenticate to the Jira server. Use only if username basic authentication is used.		String
privateKey (security)	(OAuth only) The private key generated by the client to encrypt the conversation to the server.		String

Name	Description	Default	Type
username (security)	(Basic authentication only) The username to authenticate to the Jira server. Use only if OAuth is not enabled on the Jira server. Do not set the username and OAuth token parameter, if they are both set, the username basic authentication takes precedence.		String
verificationCode (security)	(OAuth only) The verification code from Jira generated in the first step of the authorization process.		String

25.4. ENDPOINT OPTIONS

The Jira endpoint is configured using URI syntax:

```
jira:type
```

with the following path and query parameters:

25.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
type (common)	<p>Required Operation to perform. Consumers: NewIssues, NewComments. Producers: AddIssue, AttachFile, DeleteIssue, TransitionIssue, UpdateIssue, Watchers. See this class javadoc description for more information.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● ADDCOMMENT ● ADDISSUE ● ATTACH ● DELETEISSUE ● NEWISSUES ● NEWCOMMENTS ● WATCHUPDATES ● UPDATEISSUE ● TRANSITIONISSUE ● WATCHERS ● ADDISSUELINK ● ADDWORKLOG ● FETCHISSUE ● FETCHCOMMENTS 		JiraType

25.4.2. Query Parameters (16 parameters)

Name	Description	Default	Type
delay (common)	Time in milliseconds to elapse for the next poll.	6000	Integer
jiraUrl (common)	Required The Jira server url, example: .		String

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
jql (consumer)	JQL is the query language from JIRA which allows you to retrieve the data you want. For example <code>jql=project=MyProject</code> Where <code>MyProject</code> is the product key in Jira. It is important to use the <code>RAW()</code> and set the JQL inside it to prevent camel parsing it, example: <code>RAW(project in (MYP, COM) AND resolution = Unresolved)</code> .		String
maxResults (consumer)	Max number of issues to search for.	50	Integer
sendOnlyUpdatedField (consumer)	Indicator for sending only changed fields in exchange body or issue object. By default consumer sends only changed fields.	true	boolean
watchedFields (consumer)	Comma separated list of fields to watch for changes. <code>Status,Priority</code> are the defaults.	Status, Priority	String
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
accessToken (security)	(OAuth only) The access token generated by the Jira server.		String
consumerKey (security)	(OAuth only) The consumer key from Jira settings.		String
password (security)	(Basic authentication only) The password to authenticate to the Jira server. Use only if username basic authentication is used.		String
privateKey (security)	(OAuth only) The private key generated by the client to encrypt the conversation to the server.		String
username (security)	(Basic authentication only) The username to authenticate to the Jira server. Use only if OAuth is not enabled on the Jira server. Do not set the username and OAuth token parameter, if they are both set, the username basic authentication takes precedence.		String
verificationCode (security)	(OAuth only) The verification code from Jira generated in the first step of the authorization process.		String

25.5. CLIENT FACTORY

You can bind the **JiraRestClientFactory** with name **JiraRestClientFactory** in the registry to have it automatically set in the Jira endpoint.

25.6. AUTHENTICATION

Camel-jira supports [Basic Authentication](#) and [OAuth 3 legged authentication](#).

We recommend to use OAuth whenever possible, as it provides the best security for your users and system.

25.6.1. Basic authentication requirements:

- An username and password

25.6.2. OAuth authentication requirements:

Follow the tutorial in [Jira OAuth documentation](#) to generate the client private key, consumer key, verification code and access token.

- a private key, generated locally on your system.
- A verification code, generated by Jira server.
- The consumer key, set in the Jira server settings.
- An access token, generated by Jira server.

25.7. JQL

The JQL URI option is used by both consumer endpoints. Theoretically, items like "project key", etc. could be URI options themselves. However, by requiring the use of JQL, the consumers become much more flexible and powerful.

At the bare minimum, the consumers will require the following:

```
jira://[type]?[required options]&jql=project=[project key]
```

One important thing to note is that the newIssues consumer will automatically set the JQL as:

- append **ORDER BY key desc** to your JQL
- prepend **id > latestIssueId** to retrieve issues added after the camel route was started.

This is in order to optimize startup processing, rather than having to index every single issue in the project.

Another note is that, similarly, the newComments consumer will have to index every single issue **and** comment in the project. Therefore, for large projects, it's **vital** to optimize the JQL expression as much as possible. For example, the JIRA Toolkit Plugin includes a "Number of comments" custom field – use "'Number of comments' > 0" in your query. Also try to minimize based on state (status=Open), increase the polling delay, etc. Example:

```
jira://[type]?[required options]&jql=RAW(project=[project key] AND status in (Open, \"Coding In Progress\") AND \"Number of comments\">0)"
```

25.8. OPERATIONS

See a list of required headers to set when using the Jira operations. The author field for the producers is automatically set to the authenticated user in the Jira side.

If any required field is not set, then an `IllegalArgumentExpection` is throw.

There are operations that requires **id** for fields suchs as: issue type, priority, transition. Check the valid **id** on your jira project as they may differ on a jira installation and project workflow.

25.9. ADDISSUE

Required:

- **ProjectKey**: The project key, example: CAMEL, HHH, MYP.
- **IssueTypeId** or **IssueTypeName**: The **id** of the issue type or the name of the issue type, you can see the valid list in http://jira_server/rest/api/2/issue/createmeta?projectKeys=SAMPLE_KEY.
- **IssueSummary**: The summary of the issue.

Optional:

- **IssueAssignee**: the assignee user
- **IssuePriorityId** or **IssuePriorityName**: The priority of the issue, you can see the valid list in http://jira_server/rest/api/2/priority.
- **IssueComponents**: A list of string with the valid component names.
- **IssueWatchersAdd**: A list of strings with the usernames to add to the watcher list.
- **IssueDescription**: The description of the issue.

25.10. ADDCOMMENT

Required:

- **IssueKey**: The issue key identifier.
- body of the exchange is the description.

25.11. ATTACH

Only one file should attach per invocation.

Required:

- **IssueKey**: The issue key identifier.
- body of the exchange should be of type **File**

25.12. DELETEISSUE

Required:

- **IssueKey**: The issue key identifier.

25.13. TRANSITIONISSUE

Required:

- **IssueKey:** The issue key identifier.
- **IssueTransitionId:** The issue transition **id**.
- body of the exchange is the description.

25.14. UPDATEISSUE

- **IssueKey:** The issue key identifier.
- **IssueTypeId** or **IssueTypeName:** The **id** of the issue type or the name of the issue type, you can see the valid list in http://jira_server/rest/api/2/issue/createmeta?projectKeys=SAMPLE_KEY.
- **IssueSummary:** The summary of the issue.
- **IssueAssignee:** the assignee user
- **IssuePriorityId** or **IssuePriorityName:** The priority of the issue, you can see the valid list in http://jira_server/rest/api/2/priority.
- **IssueComponents:** A list of string with the valid component names.
- **IssueDescription:** The description of the issue.

25.15. WATCHER

- **IssueKey:** The issue key identifier.
- **IssueWatchersAdd:** A list of strings with the usernames to add to the watcher list.
- **IssueWatchersRemove:** A list of strings with the usernames to remove from the watcher list.

25.16. WATCHUPDATES (CONSUMER)

- **watchedFields** Comma separated list of fields to watch for changes i.e **Status,Priority,Assignee,Components** etc.
- **sendOnlyUpdatedField** By default only changed field is send as the body.

All messages also contain following headers that add additional info about the change:

- **issueKey:** Key of the updated issue
- **changed:** name of the updated field (i.e Status)
- **watchedIssues:** list of all issue keys that are watched in the time of update

25.17. SPRING BOOT AUTO-CONFIGURATION

When using jira with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
<groupId>org.apache.camel.springboot</groupId>
```

```

<artifactId>camel-jira-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 13 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.jira.access-token</code>	(OAuth only) The access token generated by the Jira server.		String
<code>camel.component.jira.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.jira.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.jira.configuration</code>	To use a shared base jira configuration. The option is a <code>org.apache.camel.component.jira.JiraConfiguration</code> type.		JiraConfiguration
<code>camel.component.jira.consumer-key</code>	(OAuth only) The consumer key from Jira settings.		String
<code>camel.component.jira.delay</code>	Time in milliseconds to elapse for the next poll.	6000	Integer
<code>camel.component.jira.enabled</code>	Whether to enable auto configuration of the jira component. This is enabled by default.		Boolean
<code>camel.component.jira.jira-url</code>	The Jira server url, example: http://my_jira.com:8081/ .		String

Name	Description	Default	Type
camel.component.jira.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.jira.password	(Basic authentication only) The password to authenticate to the Jira server. Use only if username basic authentication is used.		String
camel.component.jira.private-key	(OAuth only) The private key generated by the client to encrypt the conversation to the server.		String
camel.component.jira.username	(Basic authentication only) The username to authenticate to the Jira server. Use only if OAuth is not enabled on the Jira server. Do not set the username and OAuth token parameter, if they are both set, the username basic authentication takes precedence.		String
camel.component.jira.verification-code	(OAuth only) The verification code from Jira generated in the first step of the authorization process.		String

CHAPTER 26. JMS

Both producer and consumer are supported

This component allows messages to be sent to (or consumed from) a [JMS](#) Queue or Topic. It uses Spring's JMS support for declarative transactions, including Spring's **JmsTemplate** for sending and a **MessageListenerContainer** for consuming.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jms</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```



NOTE

Using ActiveMQ

If you are using [Apache ActiveMQ](#), you should prefer the ActiveMQ component as it has been optimized for ActiveMQ. All of the options and samples on this page are also valid for the ActiveMQ component.



NOTE

Transacted and caching

See section [Transactions and Cache Levels](#) below if you are using transactions with [JMS](#) as it can impact performance.



NOTE

Request/Reply over JMS

Make sure to read the section *Request-reply over JMS* further below on this page for important notes about request/reply, as Camel offers a number of options to configure for performance, and clustered environments.

26.1. URI FORMAT

```
jms:[queue:|topic:]destinationName[?options]
```

Where **destinationName** is a JMS queue or topic name. By default, the **destinationName** is interpreted as a queue name. For example, to connect to the queue, **FOO.BAR** use:

```
jms:FOO.BAR
```

You can include the optional **queue:** prefix, if you prefer:

```
jms:queue:FOO.BAR
```

To connect to a topic, you *must* include the **topic:** prefix. For example, to connect to the topic, **Stocks.Prices**, use:

jms:topic:Stocks.Prices

You append query options to the URI by using the following format,

?option=value&option=value&...

26.1.1. Using ActiveMQ

The JMS component reuses Spring 2's **JmsTemplate** for sending messages. This is not ideal for use in a non-J2EE container and typically requires some caching in the JMS provider to avoid [poor performance](#).

If you intend to use [Apache ActiveMQ](#) as your message broker, the recommendation is that you do one of the following:

- Use the ActiveMQ component, which is already optimized to use ActiveMQ efficiently
- Use the **PoolingConnectionFactory** in ActiveMQ.

26.1.2. Transactions and Cache Levels

If you are consuming messages and using transactions (**transacted=true**) then the default settings for cache level can impact performance.

If you are using XA transactions then you cannot cache as it can cause the XA transaction to not work properly.

If you are **not** using XA, then you should consider caching as it speeds up performance, such as setting **cacheLevelName=CACHE_CONSUMER**.

The default setting for **cacheLevelName** is **CACHE_AUTO**. This default auto detects the mode and sets the cache level accordingly to:

- **CACHE_CONSUMER** if **transacted=false**
- **CACHE_NONE** if **transacted=true**

So you can say the default setting is conservative. Consider using **cacheLevelName=CACHE_CONSUMER** if you are using non-XA transactions.

26.1.3. Durable Subscriptions

If you wish to use durable topic subscriptions, you need to specify both **clientId** and **durableSubscriptionName**. The value of the **clientId** must be unique and can only be used by a single JMS connection instance in your entire network. You may prefer to use [Virtual Topics](#) instead to avoid this limitation. More background on durable messaging [here](#).

26.1.4. Message Header Mapping

When using message headers, the JMS specification states that header names must be valid Java identifiers. So try to name your headers to be valid Java identifiers. One benefit of doing this is that you can then use your headers inside a JMS Selector (whose SQL92 syntax mandates Java identifier syntax for headers).

A simple strategy for mapping header names is used by default. The strategy is to replace any dots and

hyphens in the header name as shown below and to reverse the replacement when the header name is restored from a JMS message sent over the wire. What does this mean? No more losing method names to invoke on a bean component, no more losing the filename header for the File Component, and so on.

The current header name strategy for accepting header names in Camel is as follows:

- Dots are replaced by ``DOT`` and the replacement is reversed when Camel consume the message
- Hyphen is replaced by ``HYPHEN`` and the replacement is reversed when Camel consumes the message

You can configure many different properties on the JMS endpoint, which map to properties on the **JMSConfiguration** object.



NOTE

Mapping to Spring JMS

Many of these properties map to properties on Spring JMS, which Camel uses for sending and receiving messages. So you can get more information about these properties by consulting the relevant Spring documentation.

26.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

26.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

26.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

26.3. COMPONENT OPTIONS

The JMS component supports 98 options, which are listed below.

Name	Description	Default	Type
clientId (common)	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.		String
connectionFactory (common)	The connection factory to be use. A connection factory must be configured either on the component or endpoint.		ConnectionFactory
disableReplyTo (common)	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	boolean
durableSubscriptionName (common)	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.		String

Name	Description	Default	Type
jmsMessageType (common)	<p>Allows you to force the use of a specific <code>javax.jms.Message</code> implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • Bytes • Map • Object • Stream • Text 		JmsMessageType
replyTo (common)	Provides an explicit ReplyTo destination (overrides any incoming value of <code>Message.getJMSReplyTo()</code> in consumer).		String
testConnectionOnStartup (common)	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	boolean
acknowledgmentModeName (consumer)	<p>The JMS acknowledgement name, which is one of: <code>SESSION_TRANSACTED</code>, <code>CLIENT_ACKNOWLEDGE</code>, <code>AUTO_ACKNOWLEDGE</code>, <code>DUPS_OK_ACKNOWLEDGE</code>.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • <code>SESSION_TRANSACTED</code> • <code>CLIENT_ACKNOWLEDGE</code> • <code>AUTO_ACKNOWLEDGE</code> • <code>DUPS_OK_ACKNOWLEDGE</code> 	AUTO_ACKNOWLEDGE	String

Name	Description	Default	Type
artemisConsumerPriority (consumer)	<p>Consumer priorities allow you to ensure that high priority consumers receive messages while they are active. Normally, active consumers connected to a queue receive messages from it in a round-robin fashion. When consumer priorities are in use, messages are delivered round-robin if multiple active consumers exist with the same high priority. Messages will only go to lower priority consumers when the high priority consumers do not have credit available to consume the message, or those high priority consumers have declined to accept the message (for instance because it does not meet the criteria of any selectors associated with the consumer).</p>		int
asyncConsumer (consumer)	<p>Whether the JmsConsumer processes the Exchange asynchronously. If enabled then the JmsConsumer may pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). This means that messages may be processed not 100% strictly in order. If disabled (as default) then the Exchange is fully processed before the JmsConsumer will pickup the next message from the JMS queue. Note if transacted has been enabled, then <code>asyncConsumer=true</code> does not run asynchronously, as transaction must be executed synchronously (Camel 3.0 may support async transactions).</p>	false	boolean
autoStartup (consumer)	<p>Specifies whether the consumer container should auto-startup.</p>	true	boolean
cacheLevel (consumer)	<p>Sets the cache level by ID for the underlying JMS resources. See <code>cacheLevelName</code> option for more details.</p>		int

Name	Description	Default	Type
cacheLevelName (consumer)	<p>Sets the cache level by name for the underlying JMS resources. Possible values are: <code>CACHE_AUTO</code>, <code>CACHE_CONNECTION</code>, <code>CACHE_CONSUMER</code>, <code>CACHE_NONE</code>, and <code>CACHE_SESSION</code>. The default setting is <code>CACHE_AUTO</code>. See the Spring documentation and Transactions Cache Levels for more information.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● <code>CACHE_AUTO</code> ● <code>CACHE_CONNECTION</code> ● <code>CACHE_CONSUMER</code> ● <code>CACHE_NONE</code> ● <code>CACHE_SESSION</code> 	<code>CACHE_AUTO</code>	String
concurrentConsumers (consumer)	Specifies the default number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	1	int
maxConcurrentConsumers (consumer)	Specifies the maximum number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToMaxConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.		int
replyToDeliveryPersistent (consumer)	Specifies whether to use persistent delivery by default for replies.	true	boolean
selector (consumer)	Sets the JMS selector to use.		String

Name	Description	Default	Type
subscriptionDurable (consumer)	Set whether to make the subscription durable. The durable subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a durable subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well.	false	boolean
subscriptionName (consumer)	Set the name of a subscription to create. To be applied in case of a topic (pub-sub domain) with a shared or durable subscription. The subscription name needs to be unique within this client's JMS client id. Default is the class name of the specified message listener. Note: Only 1 concurrent consumer (which is the default of this message listener container) is allowed for each subscription, except for a shared subscription (which requires JMS 2.0).		String
subscriptionShared (consumer)	Set whether to make the subscription shared. The shared subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a shared subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Note that shared subscriptions may also be durable, so this flag can (and often will) be combined with <code>subscriptionDurable</code> as well. Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well. Requires a JMS 2.0 compatible message broker.	false	boolean
acceptMessagesWhileStopping (consumer (advanced))	Specifies whether the consumer accept messages while it is stopping. You may consider enabling this option, if you start and stop JMS routes at runtime, while there are still messages enqueued on the queue. If this option is false, and you stop the JMS route, then messages may be rejected, and the JMS broker would have to attempt redeliveries, which yet again may be rejected, and eventually the message may be moved at a dead letter queue on the JMS broker. To avoid this its recommended to enable this option.	false	boolean

Name	Description	Default	Type
allowReplyManagerQuickStop (consumer (advanced))	Whether the DefaultMessageListenerContainer used in the reply managers for request-reply messaging allow the DefaultMessageListenerContainer.runningAllowed flag to quick stop in case JmsConfiguration#isAcceptMessagesWhileStopping is enabled, and org.apache.camel.CamelContext is currently being stopped. This quick stop ability is enabled by default in the regular JMS consumers but to enable for reply managers you must enable this flag.	false	boolean
consumerType (consumer (advanced))	The consumer type to use, which can be one of: Simple, Default, or Custom. The consumer type determines which Spring JMS listener to use. Default will use org.springframework.jms.listener.DefaultMessageListenerContainer, Simple will use org.springframework.jms.listener.SimpleMessageListenerContainer. When Custom is specified, the MessageListenerContainerFactory defined by the messageListenerContainerFactory option will determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use. Enum values: <ul style="list-style-type: none"> ● Simple ● Default ● Custom 	Default	ConsumerType

Name	Description	Default	Type
defaultTaskExecutorType (consumer (advanced))	<p>Specifies what default TaskExecutor type to use in the DefaultMessageListenerContainer, for both consumer endpoints and the ReplyTo consumer of producer endpoints. Possible values: SimpleAsync (uses Spring's SimpleAsyncTaskExecutor) or ThreadPool (uses Spring's ThreadPoolTaskExecutor with optimal values - cached threadpool-like). If not set, it defaults to the previous behaviour, which uses a cached thread pool for consumer endpoints and SimpleAsync for reply consumers. The use of ThreadPool is recommended to reduce thread trash in elastic configurations with dynamically increasing and decreasing concurrent consumers.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • ThreadPool • SimpleAsync 		DefaultTaskExecutorType
eagerLoadingOfProperties (consumer (advanced))	Enables eager loading of JMS properties and payload as soon as a message is loaded which generally is inefficient as the JMS properties may not be required but sometimes can catch early any issues with the underlying JMS provider and the use of JMS properties. See also the option eagerPoisonBody.	false	boolean
eagerPoisonBody (consumer (advanced))	If eagerLoadingOfProperties is enabled and the JMS message payload (JMS body or JMS properties) is poison (cannot be read/mapped), then set this text as the message body instead so the message can be processed (the cause of the poison are already stored as exception on the Exchange). This can be turned off by setting eagerPoisonBody=false. See also the option eagerLoadingOfProperties.	Poison JMS message due to <code>\{exception.message}</code>	String
exposeListenerSession (consumer (advanced))	Specifies whether the listener session should be exposed when consuming messages.	false	boolean
replyToSameDestinationAllowed (consumer (advanced))	Whether a JMS consumer is allowed to send a reply message to the same destination that the consumer is using to consume from. This prevents an endless loop by consuming and sending back the same message to itself.	false	boolean
taskExecutor (consumer (advanced))	Allows you to specify a custom task executor for consuming messages.		TaskExecutor

Name	Description	Default	Type
deliveryDelay (producer)	Sets delivery delay to use for send calls for JMS. This option requires JMS 2.0 compliant broker.	-1	long
deliveryMode (producer)	Specifies the delivery mode to be used. Possible values are those defined by <code>javax.jms.DeliveryMode</code> . <code>NON_PERSISTENT</code> = 1 and <code>PERSISTENT</code> = 2. Enum values: <ul style="list-style-type: none"> • 1 • 2 		Integer
deliveryPersistent (producer)	Specifies whether persistent delivery is used by default.	true	boolean
explicitQoSEnabled (producer)	Set if the <code>deliveryMode</code> , <code>priority</code> or <code>timeToLive</code> qualities of service should be used when sending messages. This option is based on Spring's <code>JmsTemplate</code> . The <code>deliveryMode</code> , <code>priority</code> and <code>timeToLive</code> options are applied to the current endpoint. This contrasts with the <code>preserveMessageQoS</code> option, which operates at message granularity, reading QoS properties exclusively from the Camel In message headers.	false	Boolean
formatDateHeadersToIso8601 (producer)	Sets whether JMS date properties should be formatted according to the ISO 8601 standard.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
preserveMessageQos (producer)	Set to true, if you want to send message using the QoS settings specified on the message, instead of the QoS settings on the JMS endpoint. The following three headers are considered JMSPriority, JMSDeliveryMode, and JMSExpiration. You can provide all or only some of them. If not provided, Camel will fall back to use the values from the endpoint instead. So, when using this option, the headers override the values from the endpoint. The explicitQosEnabled option, by contrast, will only use options set on the endpoint, and not values from the message header.	false	boolean
priority (producer)	<p>Values greater than 1 specify the message priority when sending (where 1 is the lowest priority and 9 is the highest). The explicitQosEnabled option must also be enabled in order for this option to have any effect.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● 1 ● 2 ● 3 ● 4 ● 5 ● 6 ● 7 ● 8 ● 9 	4	int
replyToConcurrentConsumers (producer)	Specifies the default number of concurrent consumers when doing request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.	1	int
replyToMaxConcurrentConsumers (producer)	Specifies the maximum number of concurrent consumers when using request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.		int

Name	Description	Default	Type
replyToOnTimeoutMaxConcurrentConsumers (producer)	Specifies the maximum number of concurrent consumers for continue routing when timeout occurred when using request/reply over JMS.	1	int
replyToOverride (producer)	Provides an explicit ReplyTo destination in the JMS message, which overrides the setting of replyTo. It is useful if you want to forward the message to a remote Queue and receive the reply message from the ReplyTo destination.		String
replyToType (producer)	<p>Allows for explicitly specifying which kind of strategy to use for replyTo queues when doing request/reply over JMS. Possible values are: Temporary, Shared, or Exclusive. By default Camel will use temporary queues. However if replyTo has been configured, then Shared is used by default. This option allows you to use exclusive queues instead of shared ones. See Camel JMS documentation for more details, and especially the notes about the implications if running in a clustered environment, and the fact that Shared reply queues has lower performance than its alternatives Temporary and Exclusive.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● Temporary ● Shared ● Exclusive 		ReplyToType
requestTimeout (producer)	The timeout for waiting for a reply when using the InOut Exchange Pattern (in milliseconds). The default is 20 seconds. You can include the header CamelJmsRequestTimeout to override this endpoint configured timeout value, and thus have per message individual timeout values. See also the requestTimeoutCheckerInterval option.	20000	long
timeToLive (producer)	When sending messages, specifies the time-to-live of the message (in milliseconds).	-1	long

Name	Description	Default	Type
allowAdditionalHeaders (producer (advanced))	This option is used to allow additional headers which may have values that are invalid according to JMS specification. For example some message systems such as WMQ do this with header names using prefix JMS_IBM_MQMD_ containing values with byte array or other invalid types. You can specify multiple header names separated by comma, and use as suffix for wildcard matching.		String
allowNullBody (producer (advanced))	Whether to allow sending messages with no body. If this option is false and the message body is null, then an JMSEException is thrown.	true	boolean
alwaysCopyMessage (producer (advanced))	If true, Camel will always make a JMS message copy of the message when it is passed to the producer for sending. Copying the message is needed in some situations, such as when a <code>replyToDestinationSelectorName</code> is set (incidentally, Camel will set the <code>alwaysCopyMessage</code> option to true, if a <code>replyToDestinationSelectorName</code> is set).	false	boolean
correlationProperty (producer (advanced))	When using InOut exchange pattern use this JMS property instead of <code>JMSCorrelationID</code> JMS property to correlate messages. If set messages will be correlated solely on the value of this property <code>JMSCorrelationID</code> property will be ignored and not set by Camel.		String
disableTimeToLive (producer (advanced))	Use this option to force disabling time to live. For example when you do request/reply over JMS, then Camel will by default use the <code>requestTimeout</code> value as time to live on the message being sent. The problem is that the sender and receiver systems have to have their clocks synchronized, so they are in sync. This is not always so easy to archive. So you can use <code>disableTimeToLive=true</code> to not set a time to live value on the sent message. Then the message will not expire on the receiver system. See below in section About time to live for more details.	false	boolean
forceSendOriginalMessage (producer (advanced))	When using <code>mapJmsMessage=false</code> Camel will create a new JMS message to send to a new JMS destination if you touch the headers (get or set) during the route. Set this option to true to force Camel to send the original JMS message that was received.	false	boolean

Name	Description	Default	Type
includeSentJMSMessageID (producer (advanced))	Only applicable when sending to JMS destination using InOnly (eg fire and forget). Enabling this option will enrich the Camel Exchange with the actual JMSMessageID that was used by the JMS client when the message was sent to the JMS destination.	false	boolean
replyToCacheLevelName (producer (advanced))	<p>Sets the cache level by name for the reply consumer when doing request/reply over JMS. This option only applies when using fixed reply queues (not temporary). Camel will by default use: CACHE_CONSUMER for exclusive or shared w/ replyToSelectorName. And CACHE_SESSION for shared without replyToSelectorName. Some JMS brokers such as IBM WebSphere may require to set the replyToCacheLevelName=CACHE_NONE to work. Note: If using temporary queues then CACHE_NONE is not allowed, and you must use a higher value such as CACHE_CONSUMER or CACHE_SESSION.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● CACHE_AUTO ● CACHE_CONNECTION ● CACHE_CONSUMER ● CACHE_NONE ● CACHE_SESSION 		String
replyToDestinationSelectorName (producer (advanced))	Sets the JMS Selector using the fixed name to be used so you can filter out your own replies from the others when using a shared queue (that is, if you are not using a temporary reply queue).		String
streamMessageTypeEnabled (producer (advanced))	Sets whether StreamMessage type is enabled or not. Message payloads of streaming kind such as files, InputStream, etc will either be sent as BytesMessage or StreamMessage. This option controls which kind will be used. By default BytesMessage is used which enforces the entire message payload to be read into memory. By enabling this option the message payload is read into memory in chunks and each chunk is then written to the StreamMessage until no more data.	false	boolean

Name	Description	Default	Type
allowAutoWiredConnectionFactory (advanced)	Whether to auto-discover <code>ConnectionFactory</code> from the registry, if no connection factory has been configured. If only one instance of <code>ConnectionFactory</code> is found then it will be used. This is enabled by default.	true	boolean
allowAutoWiredDestinationResolver (advanced)	Whether to auto-discover <code>DestinationResolver</code> from the registry, if no destination resolver has been configured. If only one instance of <code>DestinationResolver</code> is found then it will be used. This is enabled by default.	true	boolean
allowSerializedHeaders (advanced)	Controls whether or not to include serialized headers. Applies only when <code>transferExchange</code> is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean
artemisStreamingEnabled (advanced)	Whether optimizing for Apache Artemis streaming mode. This can reduce memory overhead when using Artemis with JMS <code>StreamMessage</code> types. This option must only be enabled if Apache Artemis is being used.	false	boolean
asyncStartListener (advanced)	Whether to startup the <code>JmsConsumer</code> message listener asynchronously, when starting a route. For example if a <code>JmsConsumer</code> cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the <code>JmsConsumer</code> connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	boolean
asyncStopListener (advanced)	Whether to stop the <code>JmsConsumer</code> message listener asynchronously, when stopping a route.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
configuration (advanced)	To use a shared JMS configuration.		JmsConfiguration
destinationResolver (advanced)	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).		DestinationResolver
errorHandler (advanced)	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom errorHandler.		ErrorHandler
exceptionListener (advanced)	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.		ExceptionListener
idleConsumerLimit (advanced)	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	int
idleTaskExecutionLimit (advanced)	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	int
includeAllJMSXProperties (advanced)	Whether to include all JMSXxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as <code>JMSXAppID</code> , and <code>JMSXUserID</code> etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	boolean

Name	Description	Default	Type
jmsKeyFormatStrategy (advanced)	<p>Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the <code>#</code> notation.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • default • passthrough 		JmsKeyFormatStrategy
mapJmsMessage (advanced)	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a String etc.	true	boolean
maxMessagesPerTask (advanced)	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	int
messageConverter (advanced)	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .		MessageConverter
messageCreatedStrategy (advanced)	To use the given <code>MessageCreatedStrategy</code> which are invoked when Camel creates new instances of <code>javax.jms.Message</code> objects when Camel is sending a JMS message.		MessageCreatedStrategy
messageIdEnabled (advanced)	When sending, specifies whether message IDs should be added. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	boolean

Name	Description	Default	Type
messageListenerContainerFactory (advanced)	Registry ID of the MessageListenerContainerFactory used to determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use to consume messages. Setting this will automatically set consumerType to Custom.		MessageListenerContainerFactory
messageTimestampEnabled (advanced)	Specifies whether timestamps should be enabled by default on sending messages. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	boolean
pubSubNoLocal (advanced)	Specifies whether to inhibit the delivery of messages published by its own connection.	false	boolean
queueBrowseStrategy (advanced)	To use a custom QueueBrowseStrategy when browsing queues.		QueueBrowseStrategy
receiveTimeout (advanced)	The timeout for receiving messages (in milliseconds).	1000	long
recoveryInterval (advanced)	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000	long
requestTimeoutCheckerInterval (advanced)	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000	long
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean

Name	Description	Default	Type
transferException (advanced)	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a javax.jms.ObjectMessage. If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have transferExchange enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as org.apache.camel.RuntimeCamelException when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!.	false	boolean
transferExchange (advanced)	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payloads is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the receiver to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumers having to use compatible Camel versions!.	false	boolean
useMessageIDAsCorrelationID (advanced)	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	boolean
waitForProvisionCorrelationToBeUpdatedCounter (advanced)	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option useMessageIDAsCorrelationID is enabled.	50	int
waitForProvisionCorrelationToBeUpdatedThreadSleepingTime (advanced)	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100	long

Name	Description	Default	Type
headerFilterStrategy (filter)	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message.		HeaderFilterStrategy
errorHandlerLoggingLevel (logging)	Allows to configure the default errorHandler logging level for logging uncaught exceptions. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LogLevel
errorHandlerLogStackTrace (logging)	Allows to control whether stacktraces should be logged or not, by the default errorHandler.	true	boolean
password (security)	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
username (security)	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
transacted (transaction)	Specifies whether to use transacted mode.	false	boolean

Name	Description	Default	Type
transactedInOut (transaction)	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	boolean
lazyCreateTransactionManager (transaction advanced))	If true, Camel will create a JmsTransactionManager, if there is no transactionManager injected when option transacted=true.	true	boolean
transactionManager (transaction advanced))	The Spring transaction manager to use.		PlatformTransactionManager
transactionName (transaction advanced))	The name of the transaction to use.		String
transactionTimeout (transaction advanced))	The timeout value of the transaction (in seconds), if using transacted mode.	-1	int

26.4. ENDPOINT OPTIONS

The JMS endpoint is configured using URI syntax:

```
jms:destinationType:destinationName
```

with the following path and query parameters:

26.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
destinationType (common)	The kind of destination to use. Enum values: <ul style="list-style-type: none"> • queue • topic • temp-queue • temp-topic 	queue	String
destinationName (common)	Required Name of the queue or topic to use as destination.		String

26.4.2. Query Parameters (95 parameters)

Name	Description	Default	Type
clientId (common)	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.		String
connectionFactory (common)	The connection factory to be use. A connection factory must be configured either on the component or endpoint.		ConnectionFactory
disableReplyTo (common)	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	boolean
durableSubscriptionName (common)	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.		String

Name	Description	Default	Type
jmsMessageType (common)	<p>Allows you to force the use of a specific <code>javax.jms.Message</code> implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● Bytes ● Map ● Object ● Stream ● Text 		JmsMessageType
replyTo (common)	Provides an explicit ReplyTo destination (overrides any incoming value of <code>Message.getJMSReplyTo()</code> in consumer).		String
testConnectionOnStartup (common)	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	boolean
acknowledgmentModeName (consumer)	<p>The JMS acknowledgement name, which is one of: SESSION_TRANSACTED, CLIENT_ACKNOWLEDGE, AUTO_ACKNOWLEDGE, DUPS_OK_ACKNOWLEDGE.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● SESSION_TRANSACTED ● CLIENT_ACKNOWLEDGE ● AUTO_ACKNOWLEDGE ● DUPS_OK_ACKNOWLEDGE 	AUTO_ACKNOWLEDGE	String

Name	Description	Default	Type
artemisConsumerPriority (consumer)	Consumer priorities allow you to ensure that high priority consumers receive messages while they are active. Normally, active consumers connected to a queue receive messages from it in a round-robin fashion. When consumer priorities are in use, messages are delivered round-robin if multiple active consumers exist with the same high priority. Messages will only go to lower priority consumers when the high priority consumers do not have credit available to consume the message, or those high priority consumers have declined to accept the message (for instance because it does not meet the criteria of any selectors associated with the consumer).		int
asyncConsumer (consumer)	Whether the JmsConsumer processes the Exchange asynchronously. If enabled then the JmsConsumer may pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). This means that messages may be processed not 100% strictly in order. If disabled (as default) then the Exchange is fully processed before the JmsConsumer will pickup the next message from the JMS queue. Note if transacted has been enabled, then <code>asyncConsumer=true</code> does not run asynchronously, as transaction must be executed synchronously (Camel 3.0 may support async transactions).	false	boolean
autoStartup (consumer)	Specifies whether the consumer container should auto-startup.	true	boolean
cacheLevel (consumer)	Sets the cache level by ID for the underlying JMS resources. See <code>cacheLevelName</code> option for more details.		int

Name	Description	Default	Type
cacheLevelName (consumer)	<p>Sets the cache level by name for the underlying JMS resources. Possible values are: <code>CACHE_AUTO</code>, <code>CACHE_CONNECTION</code>, <code>CACHE_CONSUMER</code>, <code>CACHE_NONE</code>, and <code>CACHE_SESSION</code>. The default setting is <code>CACHE_AUTO</code>. See the Spring documentation and Transactions Cache Levels for more information.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● <code>CACHE_AUTO</code> ● <code>CACHE_CONNECTION</code> ● <code>CACHE_CONSUMER</code> ● <code>CACHE_NONE</code> ● <code>CACHE_SESSION</code> 	<code>CACHE_AUTO</code>	String
concurrentConsumers (consumer)	Specifies the default number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	1	int
maxConcurrentConsumers (consumer)	Specifies the maximum number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToMaxConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.		int
replyToDeliveryPersistent (consumer)	Specifies whether to use persistent delivery by default for replies.	true	boolean
selector (consumer)	Sets the JMS selector to use.		String

Name	Description	Default	Type
subscriptionDurable (consumer)	Set whether to make the subscription durable. The durable subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a durable subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well.	false	boolean
subscriptionName (consumer)	Set the name of a subscription to create. To be applied in case of a topic (pub-sub domain) with a shared or durable subscription. The subscription name needs to be unique within this client's JMS client id. Default is the class name of the specified message listener. Note: Only 1 concurrent consumer (which is the default of this message listener container) is allowed for each subscription, except for a shared subscription (which requires JMS 2.0).		String
subscriptionShared (consumer)	Set whether to make the subscription shared. The shared subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a shared subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Note that shared subscriptions may also be durable, so this flag can (and often will) be combined with <code>subscriptionDurable</code> as well. Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well. Requires a JMS 2.0 compatible message broker.	false	boolean
acceptMessagesWhileStopping (consumer (advanced))	Specifies whether the consumer accept messages while it is stopping. You may consider enabling this option, if you start and stop JMS routes at runtime, while there are still messages enqueued on the queue. If this option is false, and you stop the JMS route, then messages may be rejected, and the JMS broker would have to attempt redeliveries, which yet again may be rejected, and eventually the message may be moved at a dead letter queue on the JMS broker. To avoid this its recommended to enable this option.	false	boolean

Name	Description	Default	Type
allowReplyManagerQuickStop (consumer (advanced))	Whether the DefaultMessageListenerContainer used in the reply managers for request-reply messaging allow the DefaultMessageListenerContainer.runningAllowed flag to quick stop in case JmsConfiguration#isAcceptMessagesWhileStopping is enabled, and org.apache.camel.CamelContext is currently being stopped. This quick stop ability is enabled by default in the regular JMS consumers but to enable for reply managers you must enable this flag.	false	boolean
consumerType (consumer (advanced))	The consumer type to use, which can be one of: Simple, Default, or Custom. The consumer type determines which Spring JMS listener to use. Default will use org.springframework.jms.listener.DefaultMessageListenerContainer, Simple will use org.springframework.jms.listener.SimpleMessageListenerContainer. When Custom is specified, the MessageListenerContainerFactory defined by the messageListenerContainerFactory option will determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use. Enum values: <ul style="list-style-type: none"> ● Simple ● Default ● Custom 	Default	ConsumerType

Name	Description	Default	Type
defaultTaskExecutorType (consumer (advanced))	<p>Specifies what default TaskExecutor type to use in the DefaultMessageListenerContainer, for both consumer endpoints and the ReplyTo consumer of producer endpoints. Possible values: SimpleAsync (uses Spring's SimpleAsyncTaskExecutor) or ThreadPool (uses Spring's ThreadPoolTaskExecutor with optimal values - cached threadpool-like). If not set, it defaults to the previous behaviour, which uses a cached thread pool for consumer endpoints and SimpleAsync for reply consumers. The use of ThreadPool is recommended to reduce thread trash in elastic configurations with dynamically increasing and decreasing concurrent consumers.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • ThreadPool • SimpleAsync 		DefaultTaskExecutorType
eagerLoadingOfProperties (consumer (advanced))	Enables eager loading of JMS properties and payload as soon as a message is loaded which generally is inefficient as the JMS properties may not be required but sometimes can catch early any issues with the underlying JMS provider and the use of JMS properties. See also the option eagerPoisonBody.	false	boolean
eagerPoisonBody (consumer (advanced))	If eagerLoadingOfProperties is enabled and the JMS message payload (JMS body or JMS properties) is poison (cannot be read/mapped), then set this text as the message body instead so the message can be processed (the cause of the poison are already stored as exception on the Exchange). This can be turned off by setting eagerPoisonBody=false. See also the option eagerLoadingOfProperties.	Poison JMS message due to \backslash {exception.message}	String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	<p>Sets the exchange pattern when the consumer creates an exchange.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 		ExchangePattern
exposeListenerSession (consumer (advanced))	Specifies whether the listener session should be exposed when consuming messages.	false	boolean
replyToSameDestinationAllowed (consumer (advanced))	Whether a JMS consumer is allowed to send a reply message to the same destination that the consumer is using to consume from. This prevents an endless loop by consuming and sending back the same message to itself.	false	boolean
taskExecutor (consumer (advanced))	Allows you to specify a custom task executor for consuming messages.		TaskExecutor
deliveryDelay (producer)	Sets delivery delay to use for send calls for JMS. This option requires JMS 2.0 compliant broker.	-1	long
deliveryMode (producer)	<p>Specifies the delivery mode to be used. Possible values are those defined by <code>javax.jms.DeliveryMode</code>. <code>NON_PERSISTENT = 1</code> and <code>PERSISTENT = 2</code>.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • 1 • 2 		Integer
deliveryPersistent (producer)	Specifies whether persistent delivery is used by default.	true	boolean

Name	Description	Default	Type
explicitQosEnabled (producer)	Set if the deliveryMode, priority or timeToLive qualities of service should be used when sending messages. This option is based on Spring's JmsTemplate. The deliveryMode, priority and timeToLive options are applied to the current endpoint. This contrasts with the preserveMessageQos option, which operates at message granularity, reading QoS properties exclusively from the Camel In message headers.	false	Boolean
formatDateHeadersToIso8601 (producer)	Sets whether JMS date properties should be formatted according to the ISO 8601 standard.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
preserveMessageQos (producer)	Set to true, if you want to send message using the QoS settings specified on the message, instead of the QoS settings on the JMS endpoint. The following three headers are considered JMSPriority, JMSDeliveryMode, and JMSExpiration. You can provide all or only some of them. If not provided, Camel will fall back to use the values from the endpoint instead. So, when using this option, the headers override the values from the endpoint. The explicitQosEnabled option, by contrast, will only use options set on the endpoint, and not values from the message header.	false	boolean

Name	Description	Default	Type
priority (producer)	<p>Values greater than 1 specify the message priority when sending (where 1 is the lowest priority and 9 is the highest). The <code>explicitQosEnabled</code> option must also be enabled in order for this option to have any effect.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • 1 • 2 • 3 • 4 • 5 • 6 • 7 • 8 • 9 	4	int
replyToConcurrentConsumers (producer)	Specifies the default number of concurrent consumers when doing request/reply over JMS. See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads.	1	int
replyToMaxConcurrentConsumers (producer)	Specifies the maximum number of concurrent consumers when using request/reply over JMS. See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads.		int
replyToOnTimeoutMaxConcurrentConsumers (producer)	Specifies the maximum number of concurrent consumers for continue routing when timeout occurred when using request/reply over JMS.	1	int
replyToOverride (producer)	Provides an explicit <code>ReplyTo</code> destination in the JMS message, which overrides the setting of <code>replyTo</code> . It is useful if you want to forward the message to a remote Queue and receive the reply message from the <code>ReplyTo</code> destination.		String

Name	Description	Default	Type
replyToType (producer)	<p>Allows for explicitly specifying which kind of strategy to use for replyTo queues when doing request/reply over JMS. Possible values are: Temporary, Shared, or Exclusive. By default Camel will use temporary queues. However if replyTo has been configured, then Shared is used by default. This option allows you to use exclusive queues instead of shared ones. See Camel JMS documentation for more details, and especially the notes about the implications if running in a clustered environment, and the fact that Shared reply queues has lower performance than its alternatives Temporary and Exclusive.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • Temporary • Shared • Exclusive 		ReplyToType
requestTimeout (producer)	The timeout for waiting for a reply when using the InOut Exchange Pattern (in milliseconds). The default is 20 seconds. You can include the header CamelJmsRequestTimeout to override this endpoint configured timeout value, and thus have per message individual timeout values. See also the requestTimeoutCheckerInterval option.	20000	long
timeToLive (producer)	When sending messages, specifies the time-to-live of the message (in milliseconds).	-1	long
allowAdditionalHeaders (producer (advanced))	This option is used to allow additional headers which may have values that are invalid according to JMS specification. For example some message systems such as WMQ do this with header names using prefix JMS_IBM_MQMD_ containing values with byte array or other invalid types. You can specify multiple header names separated by comma, and use as suffix for wildcard matching.		String
allowNullBody (producer (advanced))	Whether to allow sending messages with no body. If this option is false and the message body is null, then an JMSException is thrown.	true	boolean

Name	Description	Default	Type
alwaysCopyMessage (producer (advanced))	If true, Camel will always make a JMS message copy of the message when it is passed to the producer for sending. Copying the message is needed in some situations, such as when a <code>replyToDestinationSelectorName</code> is set (incidentally, Camel will set the <code>alwaysCopyMessage</code> option to true, if a <code>replyToDestinationSelectorName</code> is set).	false	boolean
correlationProperty (producer (advanced))	When using InOut exchange pattern use this JMS property instead of <code>JMSCorrelationID</code> JMS property to correlate messages. If set messages will be correlated solely on the value of this property <code>JMSCorrelationID</code> property will be ignored and not set by Camel.		String
disableTimeToLive (producer (advanced))	Use this option to force disabling time to live. For example when you do request/reply over JMS, then Camel will by default use the <code>requestTimeout</code> value as time to live on the message being sent. The problem is that the sender and receiver systems have to have their clocks synchronized, so they are in sync. This is not always so easy to archive. So you can use <code>disableTimeToLive=true</code> to not set a time to live value on the sent message. Then the message will not expire on the receiver system. See below in section About time to live for more details.	false	boolean
forceSendOriginalMessage (producer (advanced))	When using <code>mapJmsMessage=false</code> Camel will create a new JMS message to send to a new JMS destination if you touch the headers (get or set) during the route. Set this option to true to force Camel to send the original JMS message that was received.	false	boolean
includeSentJMSMessageID (producer (advanced))	Only applicable when sending to JMS destination using InOnly (eg fire and forget). Enabling this option will enrich the Camel Exchange with the actual <code>JMSMessageID</code> that was used by the JMS client when the message was sent to the JMS destination.	false	boolean

Name	Description	Default	Type
replyToCacheLevelName (producer (advanced))	<p>Sets the cache level by name for the reply consumer when doing request/reply over JMS. This option only applies when using fixed reply queues (not temporary). Camel will by default use: CACHE_CONSUMER for exclusive or shared w/ replyToSelectorName. And CACHE_SESSION for shared without replyToSelectorName. Some JMS brokers such as IBM WebSphere may require to set the replyToCacheLevelName=CACHE_NONE to work. Note: If using temporary queues then CACHE_NONE is not allowed, and you must use a higher value such as CACHE_CONSUMER or CACHE_SESSION.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● CACHE_AUTO ● CACHE_CONNECTION ● CACHE_CONSUMER ● CACHE_NONE ● CACHE_SESSION 		String
replyToDestinationSelectorName (producer (advanced))	Sets the JMS Selector using the fixed name to be used so you can filter out your own replies from the others when using a shared queue (that is, if you are not using a temporary reply queue).		String
streamMessageTypeEnabled (producer (advanced))	Sets whether StreamMessage type is enabled or not. Message payloads of streaming kind such as files, InputStream, etc will either be sent as BytesMessage or StreamMessage. This option controls which kind will be used. By default BytesMessage is used which enforces the entire message payload to be read into memory. By enabling this option the message payload is read into memory in chunks and each chunk is then written to the StreamMessage until no more data.	false	boolean
allowSerializedHeaders (advanced)	Controls whether or not to include serialized headers. Applies only when transferExchange is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean
artemisStreamingEnabled (advanced)	Whether optimizing for Apache Artemis streaming mode. This can reduce memory overhead when using Artemis with JMS StreamMessage types. This option must only be enabled if Apache Artemis is being used.	false	boolean

Name	Description	Default	Type
asyncStartListener (advanced)	Whether to startup the JmsConsumer message listener asynchronously, when starting a route. For example if a JmsConsumer cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the JmsConsumer connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	boolean
asyncStopListener (advanced)	Whether to stop the JmsConsumer message listener asynchronously, when stopping a route.	false	boolean
destinationResolver (advanced)	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).		DestinationResolver
errorHandler (advanced)	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom errorHandler.		ErrorHandler
exceptionListener (advanced)	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.		ExceptionListener
headerFilterStrategy (advanced)	To use a custom HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
idleConsumerLimit (advanced)	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	int

Name	Description	Default	Type
idleTaskExecutionLimit (advanced)	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	int
includeAllJMSXProperties (advanced)	Whether to include all JMSXxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as JMSXAppID, and JMSXUserID etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	boolean
jmsKeyFormatStrategy (advanced)	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the # notation. Enum values: <ul style="list-style-type: none"> ● default ● passthrough 		JmsKeyFormatStrategy
mapJmsMessage (advanced)	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a String etc.	true	boolean
maxMessagesPerTask (advanced)	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	int
messageConverter (advanced)	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .		MessageConverter

Name	Description	Default	Type
messageCreatedStrategy (advanced)	To use the given MessageCreatedStrategy which are invoked when Camel creates new instances of javax.jms.Message objects when Camel is sending a JMS message.		MessageCreatedStrategy
messageIdEnabled (advanced)	When sending, specifies whether message IDs should be added. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	boolean
messageListenerContainerFactory (advanced)	Registry ID of the MessageListenerContainerFactory used to determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use to consume messages. Setting this will automatically set consumerType to Custom.		MessageListenerContainerFactory
messageTimestampEnabled (advanced)	Specifies whether timestamps should be enabled by default on sending messages. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	boolean
pubSubNoLocal (advanced)	Specifies whether to inhibit the delivery of messages published by its own connection.	false	boolean
receiveTimeout (advanced)	The timeout for receiving messages (in milliseconds).	1000	long
recoveryInterval (advanced)	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000	long
requestTimeoutCheckerInterval (advanced)	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000	long
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean

Name	Description	Default	Type
transferException (advanced)	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a javax.jms.ObjectMessage. If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have transferExchange enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as org.apache.camel.RuntimeCamelException when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!.	false	boolean
transferExchange (advanced)	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payloads is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the receiver to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumers having to use compatible Camel versions!.	false	boolean
useMessageIDAsCorrelationID (advanced)	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	boolean
waitForProvisionCorrelationToBeUpdatedCounter (advanced)	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option useMessageIDAsCorrelationID is enabled.	50	int
waitForProvisionCorrelationToBeUpdatedThreadSleepingTime (advanced)	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100	long

Name	Description	Default	Type
errorHandlerLoggingLevel (logging)	<p>Allows to configure the default errorHandler logging level for logging uncaught exceptions.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LogLevel
errorHandlerLogStackTrace (logging)	Allows to control whether stacktraces should be logged or not, by the default errorHandler.	true	boolean
password (security)	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
username (security)	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
transacted (transaction)	Specifies whether to use transacted mode.	false	boolean

Name	Description	Default	Type
transactedInOut (transaction)	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	boolean
lazyCreateTransactionManager (transaction advanced))	If true, Camel will create a JmsTransactionManager, if there is no transactionManager injected when option transacted=true.	true	boolean
transactionManager (transaction advanced))	The Spring transaction manager to use.		PlatformTransactionManager
transactionName (transaction advanced))	The name of the transaction to use.		String
transactionTimeout (transaction advanced))	The timeout value of the transaction (in seconds), if using transacted mode.	-1	int

26.5. SAMPLES

JMS is used in many examples for other components as well. But we provide a few samples below to get started.

26.5.1. Receiving from JMS

In the following sample we configure a route that receives JMS messages and routes the message to a POJO:

```
from("jms:queue:foo").
  to("bean:myBusinessLogic");
```

You can of course use any of the EIP patterns so the route can be context based. For example, here's how to filter an order topic for the big spenders:

```
from("jms:topic:OrdersTopic").
  filter().method("myBean", "isGoldCustomer").
  to("jms:queue:BigSpendersQueue");
```

26.5.2. Sending to JMS

In the sample below we poll a file folder and send the file content to a JMS topic. As we want the content of the file as a **TextMessage** instead of a **BytesMessage**, we need to convert the body to a **String**:

```
from("file://orders").
  convertBodyTo(String.class).
  to("jms:topic:OrdersTopic");
```

26.5.3. Using Annotations

Camel also has annotations so you can use [POJO Consuming](#) and POJO Producing.

26.5.4. Spring DSL sample

The preceding examples use the Java DSL. Camel also supports Spring XML DSL. Here is the big spender sample using Spring DSL:

```
<route>
  <from uri="jms:topic:OrdersTopic"/>
  <filter>
    <method ref="myBean" method="isGoldCustomer"/>
    <to uri="jms:queue:BigSpendersQueue"/>
  </filter>
</route>
```

26.5.5. Other samples

JMS appears in many of the examples for other components and EIP patterns, as well in this Camel documentation. So feel free to browse the documentation.

26.5.6. Using JMS as a Dead Letter Queue storing Exchange

Normally, when using [JMS](#) as the transport, it only transfers the body and headers as the payload. If you want to use [JMS](#) with a [Dead Letter Channel](#), using a JMS queue as the Dead Letter Queue, then normally the caused Exception is not stored in the JMS message. You can, however, use the **transferExchange** option on the JMS dead letter queue to instruct Camel to store the entire Exchange in the queue as a **javax.jms.ObjectMessage** that holds a

org.apache.camel.support.DefaultExchangeHolder. This allows you to consume from the Dead Letter Queue and retrieve the caused exception from the Exchange property with the key **Exchange.EXCEPTION_CAUGHT**. The demo below illustrates this:

```
// setup error handler to use JMS as queue and store the entire Exchange
errorHandler(deadLetterChannel("jms:queue:dead?transferExchange=true"));
```

Then you can consume from the JMS queue and analyze the problem:

```
from("jms:queue:dead").to("bean:myErrorAnalyzer");

// and in our bean
String body = exchange.getIn().getBody();
Exception cause = exchange.getProperty(Exchange.EXCEPTION_CAUGHT, Exception.class);
// the cause message is
String problem = cause.getMessage();
```

26.5.7. Using JMS as a Dead Letter Channel storing error only

You can use JMS to store the cause error message or to store a custom body, which you can initialize yourself. The following example uses the Message Translator EIP to do a transformation on the failed exchange before it is moved to the [JMS](#) dead letter queue:

```
// we sent it to a seda dead queue first
errorHandler(deadLetterChannel("seda:dead"));

// and on the seda dead queue we can do the custom transformation before its sent to the JMS queue
from("seda:dead").transform(exceptionMessage()).to("jms:queue:dead");
```

Here we only store the original cause error message in the transform. You can, however, use any Expression to send whatever you like. For example, you can invoke a method on a Bean or use a custom processor.

26.6. MESSAGE MAPPING BETWEEN JMS AND CAMEL

Camel automatically maps messages between **javax.jms.Message** and **org.apache.camel.Message**.

When sending a JMS message, Camel converts the message body to the following JMS message types:

Body Type	JMS Message	Comment
String	javax.jms.TextMessage	
org.w3c.dom.Node	javax.jms.TextMessage	The DOM will be converted to String .
Map	javax.jms.MapMessage	
java.io.Serializable	javax.jms.ObjectMessage	
byte[]	javax.jms.BytesMessage	

Body Type	JMS Message	Comment
<code>java.io.File</code>	<code>javax.jms.BytesMessage</code>	
<code>java.io.Reader</code>	<code>javax.jms.BytesMessage</code>	
<code>java.io.InputStream</code>	<code>javax.jms.BytesMessage</code>	
<code>java.nio.ByteBuffer</code>	<code>javax.jms.BytesMessage</code>	

When receiving a JMS message, Camel converts the JMS message to the following body type:

JMS Message	Body Type
<code>javax.jms.TextMessage</code>	<code>String</code>
<code>javax.jms.BytesMessage</code>	<code>byte[]</code>
<code>javax.jms.MapMessage</code>	<code>Map<String, Object></code>
<code>javax.jms.ObjectMessage</code>	<code>Object</code>

26.6.1. Disabling auto-mapping of JMS messages

You can use the `mapJmsMessage` option to disable the auto-mapping above. If disabled, Camel will not try to map the received JMS message, but instead uses it directly as the payload. This allows you to avoid the overhead of mapping and let Camel just pass through the JMS message. For instance, it even allows you to route `javax.jms.ObjectMessage` JMS messages with classes you do **not** have on the classpath.

26.6.2. Using a custom MessageConverter

You can use the `messageConverter` option to do the mapping yourself in a Spring `org.springframework.jms.support.converter.MessageConverter` class.

For example, in the route below we use a custom message converter when sending a message to the JMS order queue:

```
from("file://inbox/order").to("jms:queue:order?messageConverter=#myMessageConverter");
```

You can also use a custom message converter when consuming from a JMS destination.

26.6.3. Controlling the mapping strategy selected

You can use the `jmsMessageType` option on the endpoint URL to force a specific message type for all messages.

In the route below, we poll files from a folder and send them as `javax.jms.TextMessage` as we have forced the JMS producer endpoint to use text messages:

```
from("file://inbox/order").to("jms:queue:order?jmsMessageType=Text");
```

You can also specify the message type to use for each message by setting the header with the key **CamelJmsMessageType**. For example:

```
from("file://inbox/order").setHeader("CamelJmsMessageType",
  JmsMessageType.Text).to("jms:queue:order");
```

The possible values are defined in the **enum** class, **org.apache.camel.jms.JmsMessageType**.

26.7. MESSAGE FORMAT WHEN SENDING

The exchange that is sent over the JMS wire must conform to the [JMS Message spec](#).

For the **exchange.in.header** the following rules apply for the header **keys**:

- Keys starting with **JMS** or **JMSX** are reserved.
- **exchange.in.headers** keys must be literals and all be valid Java identifiers (do not use dots in the key name).
- Camel replaces dots & hyphens and the reverse when consuming JMS messages:
 - `.` is replaced by ``DOT`` and the reverse replacement when Camel consumes the message.
 - `-` is replaced by ``HYPHEN`` and the reverse replacement when Camel consumes the message.
- See also the option **jmsKeyFormatStrategy**, which allows use of your own custom strategy for formatting keys.

For the **exchange.in.header**, the following rules apply for the header **values**:

- The values must be primitives or their counter objects (such as **Integer**, **Long**, **Character**). The types, **String**, **CharSequence**, **Date**, **BigDecimal** and **BigInteger** are all converted to their **toString()** representation. All other types are dropped.

Camel will log with category **org.apache.camel.component.jms.JmsBinding** at **DEBUG** level if it drops a given header value. For example:

```
2008-07-09 06:43:04,046 [main          ] DEBUG JmsBinding
- Ignoring non primitive header: order of class:
org.apache.camel.component.jms.issues.DummyOrder with value: DummyOrder{orderId=333,
itemId=4444, quantity=2}
```

26.8. MESSAGE FORMAT WHEN RECEIVING

Camel adds the following properties to the **Exchange** when it receives a message:

Property	Type	Description
org.apache.camel.jms.replyDestination	javax.jms.Destination	The reply destination.

Camel adds the following JMS properties to the In message headers when it receives a JMS message:

Header	Type	Description
JMSCorrelationID	String	The JMS correlation ID.
JMSDeliveryMode	int	The JMS delivery mode.
JMSDestination	javax.jms.Destination	The JMS destination.
JMSExpiration	long	The JMS expiration.
JMSMessageID	String	The JMS unique message ID.
JMSPriority	int	The JMS priority (with 0 as the lowest priority and 9 as the highest).
JMSRedelivered	boolean	Is the JMS message redelivered.
JMSReplyTo	javax.jms.Destination	The JMS reply-to destination.
JMSTimestamp	long	The JMS timestamp.
JMSType	String	The JMS type.
JMSXGroupID	String	The JMS group ID.

As all the above information is standard JMS you can check the [JMS documentation](#) for further details.

26.9. ABOUT USING CAMEL TO SEND AND RECEIVE MESSAGES AND JMSREPLYTO

The JMS component is complex and you have to pay close attention to how it works in some cases. So this is a short summary of some of the areas/pitfalls to look for.

When Camel sends a message using its **JMSProducer**, it checks the following conditions:

- The message exchange pattern,
- Whether a **JMSReplyTo** was set in the endpoint or in the message headers,
- Whether any of the following options have been set on the JMS endpoint: **disableReplyTo**, **preserveMessageQos**, **explicitQosEnabled**.

All this can be a tad complex to understand and configure to support your use case.

26.9.1. JmsProducer

The **JmsProducer** behaves as follows, depending on configuration:

Exchange Pattern	Other options	Description
<i>InOut</i>	-	Camel will expect a reply, set a temporary JMSReplyTo , and after sending the message, it will start to listen for the reply message on the temporary queue.
<i>InOut</i>	JMSReplyTo is set	Camel will expect a reply and, after sending the message, it will start to listen for the reply message on the specified JMSReplyTo queue.
<i>InOnly</i>	-	Camel will send the message and not expect a reply.
<i>InOnly</i>	JMSReplyTo is set	By default, Camel discards the JMSReplyTo destination and clears the JMSReplyTo header before sending the message. Camel then sends the message and does not expect a reply. Camel logs this in the log at WARN level (changed to DEBUG level from Camel 2.6 onwards. You can use preserveMessageQuo=true to instruct Camel to keep the JMSReplyTo . In all situations the JmsProducer does not expect any reply and thus continue after sending the message.

26.9.2. JmsConsumer

The **JmsConsumer** behaves as follows, depending on configuration:

Exchange Pattern	Other options	Description
<i>InOut</i>	-	Camel will send the reply back to the JMSReplyTo queue.
<i>InOnly</i>	-	Camel will not send a reply back, as the pattern is <i>InOnly</i> .
-	disableReplyTo=true	This option suppresses replies.

So pay attention to the message exchange pattern set on your exchanges.

If you send a message to a JMS destination in the middle of your route you can specify the exchange pattern to use, see more at Request Reply.

This is useful if you want to send an **InOnly** message to a JMS topic:

```
from("activemq:queue:in")
    .to("bean:validateOrder")
    .to(ExchangePattern.InOnly, "activemq:topic:order")
    .to("bean:handleOrder");
```


26.10. REUSE ENDPOINT AND SEND TO DIFFERENT DESTINATIONS COMPUTED AT RUNTIME

If you need to send messages to a lot of different JMS destinations, it makes sense to reuse a JMS endpoint and specify the real destination in a message header. This allows Camel to reuse the same endpoint, but send to different destinations. This greatly reduces the number of endpoints created and economizes on memory and thread resources.

You can specify the destination in the following headers:

Header	Type	Description
CamelJmsDestination	javax.jms.Destination	A destination object.
CamelJmsDestinationName	String	The destination name.

For example, the following route shows how you can compute a destination at run time and use it to override the destination appearing in the JMS URL:

```
from("file://inbox")
  .to("bean:computeDestination")
  .to("activemq:queue:dummy");
```

The queue name, **dummy**, is just a placeholder. It must be provided as part of the JMS endpoint URL, but it will be ignored in this example.

In the **computeDestination** bean, specify the real destination by setting the **CamelJmsDestinationName** header as follows:

```
public void setJmsHeader(Exchange exchange) {
    String id = ....
    exchange.getIn().setHeader("CamelJmsDestinationName", "order:" + id);
}
```

Then Camel will read this header and use it as the destination instead of the one configured on the endpoint. So, in this example Camel sends the message to **activemq:queue:order:2**, assuming the **id** value was 2.

If both the **CamelJmsDestination** and the **CamelJmsDestinationName** headers are set, **CamelJmsDestination** takes priority. Keep in mind that the JMS producer removes both **CamelJmsDestination** and **CamelJmsDestinationName** headers from the exchange and do not propagate them to the created JMS message in order to avoid the accidental loops in the routes (in scenarios when the message will be forwarded to the another JMS endpoint).

26.11. CONFIGURING DIFFERENT JMS PROVIDERS

You can configure your JMS provider in Spring XML as follows:

Basically, you can configure as many JMS component instances as you wish and give them a **unique name using the id attribute**. The preceding example configures an **activemq** component. You could do the same to configure MQSeries, TibCo, BEA, Sonic and so on.

Once you have a named JMS component, you can then refer to endpoints within that component using URIs. For example for the component name, **activemq**, you can then refer to destinations using the URI format, **activemq:[queue:]topic:destinationName**. You can use the same approach for all other JMS providers.

This works by the SpringCamelContext lazily fetching components from the spring context for the scheme name you use for Endpoint URIs and having the Component resolve the endpoint URIs.

26.11.1. Using JNDI to find the ConnectionFactory

If you are using a J2EE container, you might need to look up JNDI to find the JMS **ConnectionFactory** rather than use the usual **<bean>** mechanism in Spring. You can do this using Spring's factory bean or the new Spring XML namespace. For example:

```
<bean id="weblogic" class="org.apache.camel.component.jms.JmsComponent">
  <property name="connectionFactory" ref="myConnectionFactory"/>
</bean>

<jee:jndi-lookup id="myConnectionFactory" jndi-name="jms/connectionFactory"/>
```

See [The jee schema](#) in the Spring reference documentation for more details about JNDI lookup.

26.12. CONCURRENT CONSUMING

A common requirement with JMS is to consume messages concurrently in multiple threads in order to make an application more responsive. You can set the **concurrentConsumers** option to specify the number of threads servicing the JMS endpoint, as follows:

```
from("jms:SomeQueue?concurrentConsumers=20").
  bean(MyClass.class);
```

You can configure this option in one of the following ways:

- On the **JmsComponent**,
- On the endpoint URI or,
- By invoking **setConcurrentConsumers()** directly on the **JmsEndpoint**.

26.12.1. Concurrent Consuming with async consumer

Notice that each concurrent consumer will only pickup the next available message from the JMS broker, when the current message has been fully processed. You can set the option **asyncConsumer=true** to let the consumer pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). See more details in the table on top of the page about the **asyncConsumer** option.

```
from("jms:SomeQueue?concurrentConsumers=20&asyncConsumer=true").
  bean(MyClass.class);
```

26.13. REQUEST-REPLY OVER JMS

Camel supports Request Reply over JMS. In essence the MEP of the Exchange should be **InOut** when you send a message to a JMS queue.

Camel offers a number of options to configure request/reply over JMS that influence performance and clustered environments. The table below summaries the options.

Option	Performance	Cluster	Description
Temporary	Fast	Yes	A temporary queue is used as reply queue, and automatic created by Camel. To use this do not specify a replyTo queue name. And you can optionally configure replyToType=Temporary to make it stand out that temporary queues are in use.
Shared	Slow	Yes	A shared persistent queue is used as reply queue. The queue must be created beforehand, although some brokers can create them on the fly such as Apache ActiveMQ. To use this you must specify the replyTo queue name. And you can optionally configure replyToType=Shared to make it stand out that shared queues are in use. A shared queue can be used in a clustered environment with multiple nodes running this Camel application at the same time. All using the same shared reply queue. This is possible because JMS Message selectors are used to correlate expected reply messages; this impacts performance though. JMS Message selectors is slower, and therefore not as fast as Temporary or Exclusive queues. See further below how to tweak this for better performance.

Option	Performance	Cluster	Description
Exclusive	Fast	No (*Yes)	An exclusive persistent queue is used as reply queue. The queue must be created beforehand, although some brokers can create them on the fly such as Apache ActiveMQ. To use this you must specify the replyTo queue name. And you must configure replyToType=Exclusive to instruct Camel to use exclusive queues, as Shared is used by default, if a replyTo queue name was configured. When using exclusive reply queues, then JMS Message selectors are not in use, and therefore other applications must not use this queue as well. An exclusive queue cannot be used in a clustered environment with multiple nodes running this Camel application at the same time; as we do not have control if the reply queue comes back to the same node that sent the request message; that is why shared queues use JMS Message selectors to make sure of this. Though if you configure each Exclusive reply queue with a unique name per node, then you can run this in a clustered environment. As then the reply message will be sent back to that queue for the given node, that awaits the reply message.
concurrentConsumers	Fast	Yes	Allows to process reply messages concurrently using concurrent message listeners in use. You can specify a range using the concurrentConsumers and maxConcurrentConsumers options. Notice: That using Shared reply queues may not work as well with concurrent listeners, so use this option with care.
maxConcurrentConsumers	Fast	Yes	Allows to process reply messages concurrently using concurrent message listeners in use. You can specify a range using the concurrentConsumers and maxConcurrentConsumers options. Notice: That using Shared reply queues may not work as well with concurrent listeners, so use this option with care.

The **JmsProducer** detects the **InOut** and provides a **JMSReplyTo** header with the reply destination to be used. By default Camel uses a temporary queue, but you can use the **replyTo** option on the endpoint to specify a fixed reply queue (see more below about fixed reply queue).

Camel will automatically setup a consumer which listen on the reply queue, so you should **not** do anything.

This consumer is a Spring **DefaultMessageListenerContainer** which listen for replies. However it's fixed to 1 concurrent consumer.

That means replies will be processed in sequence as there are only 1 thread to process the replies. You can configure the listener to use concurrent threads using the **concurrentConsumers** and **maxConcurrentConsumers** options. This allows you to easier configure this in Camel as shown below:

■

```

from(xxx)
.inOut().to("activemq:queue:foo?concurrentConsumers=5")
.to(yyy)
.to(zzz);

```

In this route we instruct Camel to route replies asynchronously using a thread pool with 5 threads.

26.13.1. Request-reply over JMS and using a shared fixed reply queue

If you use a fixed reply queue when doing Request Reply over JMS as shown in the example below, then pay attention.

```

from(xxx)
.inOut().to("activemq:queue:foo?replyTo=bar")
.to(yyy)

```

In this example the fixed reply queue named "bar" is used. By default Camel assumes the queue is shared when using fixed reply queues, and therefore it uses a **JMSSelector** to only pickup the expected reply messages (eg based on the **JMSCorrelationID**). See next section for exclusive fixed reply queues. That means its not as fast as temporary queues. You can speedup how often Camel will pull for reply messages using the **receiveTimeout** option. By default its 1000 millis. So to make it faster you can set it to 250 millis to pull 4 times per second as shown:

```

from(xxx)
.inOut().to("activemq:queue:foo?replyTo=bar&receiveTimeout=250")
.to(yyy)

```

Notice this will cause the Camel to send pull requests to the message broker more frequent, and thus require more network traffic.

It is generally recommended to use temporary queues if possible.

26.13.2. Request-reply over JMS and using an exclusive fixed reply queue

In the previous example, Camel would anticipate the fixed reply queue named "bar" was shared, and thus it uses a **JMSSelector** to only consume reply messages which it expects. However there is a drawback doing this as the JMS selector is slower. Also the consumer on the reply queue is slower to update with new JMS selector ids. In fact it only updates when the **receiveTimeout** option times out, which by default is 1 second. So in theory the reply messages could take up till about 1 sec to be detected. On the other hand if the fixed reply queue is exclusive to the Camel reply consumer, then we can avoid using the JMS selectors, and thus be more performant. In fact as fast as using temporary queues. There is the **ReplyToType** option which you can configure to **Exclusive** to tell Camel that the reply queue is exclusive as shown in the example below:

```

from(xxx)
.inOut().to("activemq:queue:foo?replyTo=bar&replyToType=Exclusive")
.to(yyy)

```

Mind that the queue must be exclusive to each and every endpoint. So if you have two routes, then they each need an unique reply queue as shown in the next example:

```

from(xxx)
.inOut().to("activemq:queue:foo?replyTo=bar&replyToType=Exclusive")
.to(yyy)

```

```
from(aaa)
.inOut().to("activemq:queue:order?replyTo=order.reply&replyToType=Exclusive")
.to(bbb)
```

The same applies if you run in a clustered environment. Then each node in the cluster must use an unique reply queue name. As otherwise each node in the cluster may pickup messages which was intended as a reply on another node. For clustered environments its recommended to use shared reply queues instead.

26.14. SYNCHRONIZING CLOCKS BETWEEN SENDERS AND RECEIVERS

When doing messaging between systems, its desirable that the systems have synchronized clocks. For example when sending a [JMS](#) message, then you can set a time to live value on the message. Then the receiver can inspect this value, and determine if the message is already expired, and thus drop the message instead of consume and process it. However this requires that both sender and receiver have synchronized clocks. If you are using [ActiveMQ](#) then you can use the [timestamp plugin](#) to synchronize clocks.

26.15. ABOUT TIME TO LIVE

Read first above about synchronized clocks.

When you do request/reply (InOut) over [JMS](#) with Camel then Camel uses a timeout on the sender side, which is default 20 seconds from the **requestTimeout** option. You can control this by setting a higher/lower value. However the time to live value is still set on the message being send. So that requires the clocks to be synchronized between the systems. If they are not, then you may want to disable the time to live value being set. This is now possible using the **disableTimeToLive** option from **Camel 2.8** onwards. So if you set this option to **disableTimeToLive=true**, then Camel does **not** set any time to live value when sending [JMS](#) messages. **But** the request timeout is still active. So for example if you do request/reply over [JMS](#) and have disabled time to live, then Camel will still use a timeout by 20 seconds (the **requestTimeout** option). That option can of course also be configured. So the two options **requestTimeout** and **disableTimeToLive** gives you fine grained control when doing request/reply.

You can provide a header in the message to override and use as the request timeout value instead of the endpoint configured value. For example:

```
from("direct:someWhere")
.to("jms:queue:foo?replyTo=bar&requestTimeout=30s")
.to("bean:processReply");
```

In the route above we have a endpoint configured **requestTimeout** of 30 seconds. So Camel will wait up till 30 seconds for that reply message to come back on the bar queue. If no reply message is received then a **org.apache.camel.ExchangeTimedOutException** is set on the Exchange and Camel continues routing the message, which would then fail due the exception, and Camel's error handler reacts.

If you want to use a per message timeout value, you can set the header with key **org.apache.camel.component.jms.JmsConstants#JMS_REQUEST_TIMEOUT** which has constant value **"CamelJmsRequestTimeout"** with a timeout value as long type.

For example we can use a bean to compute the timeout value per individual message, such as calling the **"whatIsTheTimeout"** method on the service bean as shown below:

```
from("direct:someWhere")
```

```
.setHeader("CamelJmsRequestTimeout", method(ServiceBean.class, "whatIsTheTimeout"))
.to("jms:queue:foo?replyTo=bar&requestTimeout=30s")
.to("bean:processReply");
```

When you do fire and forget (InOut) over [JMS](#) with Camel then Camel by default does **not** set any time to live value on the message. You can configure a value by using the **timeToLive** option. For example to indicate a 5 sec., you set **timeToLive=5000**. The option **disableTimeToLive** can be used to force disabling the time to live, also for InOnly messaging. The **requestTimeout** option is not being used for InOnly messaging.

26.16. ENABLING TRANSACTED CONSUMPTION

A common requirement is to consume from a queue in a transaction and then process the message using the Camel route. To do this, just ensure that you set the following properties on the component/endpoint:

- **transacted** = true
- **transactionManager** = a *Transaction Manager* - typically the **JmsTransactionManager**

See the Transactional Client EIP pattern for further details.

Transactions and [Request Reply] over JMS

When using Request Reply over JMS you cannot use a single transaction; JMS will not send any messages until a commit is performed, so the server side won't receive anything at all until the transaction commits. Therefore to use [Request Reply](#) you must commit a transaction after sending the request and then use a separate transaction for receiving the response.

To address this issue the JMS component uses different properties to specify transaction use for oneway messaging and request reply messaging:

The **transacted** property applies **only** to the InOnly message Exchange Pattern (MEP).

You can leverage the [DMLC transacted session API](#) using the following properties on component/endpoint:

- **transacted** = true
- **lazyCreateTransactionManager** = false

The benefit of doing so is that the cacheLevel setting will be honored when using local transactions without a configured TransactionManager. When a TransactionManager is configured, no caching happens at DMLC level and it is necessary to rely on a pooled connection factory. For more details about this kind of setup, see [here](#) and [here](#).

26.17. USING JMSREPLYTO FOR LATE REPLIES

When using Camel as a JMS listener, it sets an Exchange property with the value of the ReplyTo **javax.jms.Destination** object, having the key **ReplyTo**. You can obtain this **Destination** as follows:

```
Destination replyDestination =
exchange.getIn().getHeader(JmsConstants.JMS_REPLY_DESTINATION, Destination.class);
```

And then later use it to send a reply using regular JMS or Camel.

```
// we need to pass in the JMS component, and in this sample we use ActiveMQ
JmsEndpoint endpoint = JmsEndpoint.newInstance(replyDestination, activeMQComponent);
// now we have the endpoint we can use regular Camel API to send a message to it
template.sendBody(endpoint, "Here is the late reply.");
```

A different solution to sending a reply is to provide the **replyDestination** object in the same Exchange property when sending. Camel will then pick up this property and use it for the real destination. The endpoint URI must include a dummy destination, however. For example:

```
// we pretend to send it to some non existing dummy queue
template.send("activemq:queue:dummy, new Processor() {
    public void process(Exchange exchange) throws Exception {
        // and here we override the destination with the ReplyTo destination object so the message is sent
        // to there instead of dummy
        exchange.getIn().setHeader(JmsConstants.JMS_DESTINATION, replyDestination);
        exchange.getIn().setBody("Here is the late reply.");
    }
}
```

26.18. USING A REQUEST TIMEOUT

In the sample below we send a Request Reply style message Exchange (we use the **requestBody** method = **InOut**) to the slow queue for further processing in Camel and we wait for a return reply:

26.19. SENDING AN INONLY MESSAGE AND KEEPING THE JMSREPLYTO HEADER

When sending to a [JMS](#) destination using **camel-jms** the producer will use the MEP to detect if its *InOnly* or *InOut* messaging. However there can be times where you want to send an *InOnly* message but keeping the **JMSReplyTo** header. To do so you have to instruct Camel to keep it, otherwise the **JMSReplyTo** header will be dropped.

For example to send an *InOnly* message to the foo queue, but with a **JMSReplyTo** with bar queue you can do as follows:

```
template.send("activemq:queue:foo?preserveMessageQos=true", new Processor() {
    public void process(Exchange exchange) throws Exception {
        exchange.getIn().setBody("World");
        exchange.getIn().setHeader("JMSReplyTo", "bar");
    }
});
```

Notice we use **preserveMessageQos=true** to instruct Camel to keep the **JMSReplyTo** header.

26.20. SETTING JMS PROVIDER OPTIONS ON THE DESTINATION

Some JMS providers, like IBM's WebSphere MQ need options to be set on the JMS destination. For example, you may need to specify the **targetClient** option. Since **targetClient** is a WebSphere MQ option and not a Camel URI option, you need to set that on the JMS destination name like so:


```
// ...
.setHeader("CamelJmsDestinationName", constant("queue:///MY_QUEUE?targetClient=1"))
.to("wmq:queue:MY_QUEUE?useMessageIDAsCorrelationID=true");
```

Some versions of WMQ won't accept this option on the destination name and you will get an exception like:

```
com.ibm.msg.client.jms.DetailedJMSEException: JMSSC0005: The specified
value 'MY_QUEUE?targetClient=1' is not allowed for
'XMSC_DESTINATION_NAME'
```

A workaround is to use a custom DestinationResolver:

```
JmsComponent wmq = new JmsComponent(connectionFactory);

wmq.setDestinationResolver(new DestinationResolver() {
    public Destination resolveDestinationName(Session session, String destinationName, boolean
pubSubDomain) throws JMSEException {
        MQQueueSession wmqSession = (MQQueueSession) session;
        return wmqSession.createQueue("queue://" + destinationName + "?targetClient=1");
    }
});
```

26.21. SPRING BOOT AUTO-CONFIGURATION

When using jms with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jms-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 99 options, which are listed below.

Name	Description	Default	Type
camel.component.jms.accept-messages-while-stopping	Specifies whether the consumer accept messages while it is stopping. You may consider enabling this option, if you start and stop JMS routes at runtime, while there are still messages enqueued on the queue. If this option is false, and you stop the JMS route, then messages may be rejected, and the JMS broker would have to attempt redeliveries, which yet again may be rejected, and eventually the message may be moved at a dead letter queue on the JMS broker. To avoid this its recommended to enable this option.	false	Boolean

Name	Description	Default	Type
<code>camel.component.jms.acknowledgment-mode-name</code>	The JMS acknowledgement name, which is one of: SESSION_TRANSACTED, CLIENT_ACKNOWLEDGE, AUTO_ACKNOWLEDGE, DUPS_OK_ACKNOWLEDGE.	AUTO_ACKNOWLEDGE	String
<code>camel.component.jms.allow-additional-headers</code>	This option is used to allow additional headers which may have values that are invalid according to JMS specification. For example some message systems such as WMQ do this with header names using prefix JMS_IBM_MQMD_ containing values with byte array or other invalid types. You can specify multiple header names separated by comma, and use as suffix for wildcard matching.		String
<code>camel.component.jms.allow-auto-wired-connection-factory</code>	Whether to auto-discover ConnectionFactory from the registry, if no connection factory has been configured. If only one instance of ConnectionFactory is found then it will be used. This is enabled by default.	true	Boolean
<code>camel.component.jms.allow-auto-wired-destination-resolver</code>	Whether to auto-discover DestinationResolver from the registry, if no destination resolver has been configured. If only one instance of DestinationResolver is found then it will be used. This is enabled by default.	true	Boolean
<code>camel.component.jms.allow-null-body</code>	Whether to allow sending messages with no body. If this option is false and the message body is null, then an JMSEException is thrown.	true	Boolean
<code>camel.component.jms.allow-reply-manager-quick-stop</code>	Whether the DefaultMessageListenerContainer used in the reply managers for request-reply messaging allow the DefaultMessageListenerContainer.runningAllowed flag to quick stop in case JmsConfiguration#isAcceptMessagesWhileStopping is enabled, and org.apache.camel.CamelContext is currently being stopped. This quick stop ability is enabled by default in the regular JMS consumers but to enable for reply managers you must enable this flag.	false	Boolean
<code>camel.component.jms.allow-serialized-headers</code>	Controls whether or not to include serialized headers. Applies only when transferExchange is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	Boolean

Name	Description	Default	Type
camel.component.jms.always-copy-message	If true, Camel will always make a JMS message copy of the message when it is passed to the producer for sending. Copying the message is needed in some situations, such as when a <code>replyToDestinationSelectorName</code> is set (incidentally, Camel will set the <code>alwaysCopyMessage</code> option to true, if a <code>replyToDestinationSelectorName</code> is set).	false	Boolean
camel.component.jms.artemis-consumer-priority	Consumer priorities allow you to ensure that high priority consumers receive messages while they are active. Normally, active consumers connected to a queue receive messages from it in a round-robin fashion. When consumer priorities are in use, messages are delivered round-robin if multiple active consumers exist with the same high priority. Messages will only going to lower priority consumers when the high priority consumers do not have credit available to consume the message, or those high priority consumers have declined to accept the message (for instance because it does not meet the criteria of any selectors associated with the consumer).		Integer
camel.component.jms.artemis-streaming-enabled	Whether optimizing for Apache Artemis streaming mode. This can reduce memory overhead when using Artemis with JMS <code>StreamMessage</code> types. This option must only be enabled if Apache Artemis is being used.	false	Boolean
camel.component.jms.async-consumer	Whether the <code>JmsConsumer</code> processes the Exchange asynchronously. If enabled then the <code>JmsConsumer</code> may pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). This means that messages may be processed not 100% strictly in order. If disabled (as default) then the Exchange is fully processed before the <code>JmsConsumer</code> will pickup the next message from the JMS queue. Note if <code>transacted</code> has been enabled, then <code>asyncConsumer=true</code> does not run asynchronously, as transaction must be executed synchronously (Camel 3.0 may support async transactions).	false	Boolean

Name	Description	Default	Type
<code>camel.component.jms.async-start-listener</code>	Whether to startup the JmsConsumer message listener asynchronously, when starting a route. For example if a JmsConsumer cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the JmsConsumer connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	Boolean
<code>camel.component.jms.async-stop-listener</code>	Whether to stop the JmsConsumer message listener asynchronously, when stopping a route.	false	Boolean
<code>camel.component.jms.auto-startup</code>	Specifies whether the consumer container should auto-startup.	true	Boolean
<code>camel.component.jms.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.jms.cache-level</code>	Sets the cache level by ID for the underlying JMS resources. See <code>cacheLevelName</code> option for more details.		Integer
<code>camel.component.jms.cache-level-name</code>	Sets the cache level by name for the underlying JMS resources. Possible values are: <code>CACHE_AUTO</code> , <code>CACHE_CONNECTION</code> , <code>CACHE_CONSUMER</code> , <code>CACHE_NONE</code> , and <code>CACHE_SESSION</code> . The default setting is <code>CACHE_AUTO</code> . See the Spring documentation and Transactions Cache Levels for more information.	<code>CACHE_AUTO</code>	String
<code>camel.component.jms.client-id</code>	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.		String

Name	Description	Default	Type
camel.component.jms.concurrent-consumers	Specifies the default number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	1	Integer
camel.component.jms.configuration	To use a shared JMS configuration. The option is a <code>org.apache.camel.component.jms.JmsConfiguration</code> type.		JmsConfiguration
camel.component.jms.connection-factory	The connection factory to be use. A connection factory must be configured either on the component or endpoint. The option is a <code>javax.jms.ConnectionFactory</code> type.		ConnectionFactory
camel.component.jms.consumer-type	The consumer type to use, which can be one of: Simple, Default, or Custom. The consumer type determines which Spring JMS listener to use. Default will use <code>org.springframework.jms.listener.DefaultMessageListenerContainer</code> , Simple will use <code>org.springframework.jms.listener.SimpleMessageListenerContainer</code> . When Custom is specified, the <code>MessageListenerContainerFactory</code> defined by the <code>messageListenerContainerFactory</code> option will determine what <code>org.springframework.jms.listener.AbstractMessageListenerContainer</code> to use.		ConsumerType
camel.component.jms.correlation-property	When using InOut exchange pattern use this JMS property instead of <code>JMSCorrelationID</code> JMS property to correlate messages. If set messages will be correlated solely on the value of this property <code>JMSCorrelationID</code> property will be ignored and not set by Camel.		String

Name	Description	Default	Type
camel.component.jms.default-task-executor-type	Specifies what default TaskExecutor type to use in the DefaultMessageListenerContainer, for both consumer endpoints and the ReplyTo consumer of producer endpoints. Possible values: SimpleAsync (uses Spring's SimpleAsyncTaskExecutor) or ThreadPool (uses Spring's ThreadPoolTaskExecutor with optimal values - cached threadpool-like). If not set, it defaults to the previous behaviour, which uses a cached thread pool for consumer endpoints and SimpleAsync for reply consumers. The use of ThreadPool is recommended to reduce thread trash in elastic configurations with dynamically increasing and decreasing concurrent consumers.		DefaultTaskExecutorType
camel.component.jms.delivery-delay	Sets delivery delay to use for send calls for JMS. This option requires JMS 2.0 compliant broker.	-1	Long
camel.component.jms.delivery-mode	Specifies the delivery mode to be used. Possible values are those defined by javax.jms.DeliveryMode. NON_PERSISTENT = 1 and PERSISTENT = 2.		Integer
camel.component.jms.delivery-persistent	Specifies whether persistent delivery is used by default.	true	Boolean
camel.component.jms.destination-resolver	A pluggable org.springframework.jms.support.destination.DestinationResolver that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry). The option is a org.springframework.jms.support.destination.DestinationResolver type.		DestinationResolver
camel.component.jms.disable-reply-to	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	Boolean

Name	Description	Default	Type
camel.component.jms.disable-time-to-live	Use this option to force disabling time to live. For example when you do request/reply over JMS, then Camel will by default use the requestTimeout value as time to live on the message being sent. The problem is that the sender and receiver systems have to have their clocks synchronized, so they are in sync. This is not always so easy to archive. So you can use disableTimeToLive=true to not set a time to live value on the sent message. Then the message will not expire on the receiver system. See below in section About time to live for more details.	false	Boolean
camel.component.jms.durable-subscription-name	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.		String
camel.component.jms.eager-loading-of-properties	Enables eager loading of JMS properties and payload as soon as a message is loaded which generally is inefficient as the JMS properties may not be required but sometimes can catch early any issues with the underlying JMS provider and the use of JMS properties. See also the option eagerPoisonBody.	false	Boolean
camel.component.jms.eager-poison-body	If eagerLoadingOfProperties is enabled and the JMS message payload (JMS body or JMS properties) is poison (cannot be read/mapped), then set this text as the message body instead so the message can be processed (the cause of the poison are already stored as exception on the Exchange). This can be turned off by setting eagerPoisonBody=false. See also the option eagerLoadingOfProperties.	Poison JMS message due to \${exception.message}	String
camel.component.jms.enabled	Whether to enable auto configuration of the jms component. This is enabled by default.		Boolean
camel.component.jms.error-handler	Specifies a org.springframework.util.ErrorHandler to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using errorHandlerLoggingLevel and errorHandlerLogStackTrace options. This makes it much easier to configure, than having to code a custom errorHandler. The option is a org.springframework.util.ErrorHandler type.		ErrorHandler

Name	Description	Default	Type
<code>camel.component.jms.error-handler-log-stack-trace</code>	Allows to control whether stacktraces should be logged or not, by the default errorHandler.	true	Boolean
<code>camel.component.jms.error-handler-logging-level</code>	Allows to configure the default errorHandler logging level for logging uncaught exceptions.		LogLevel
<code>camel.component.jms.exception-listener</code>	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions. The option is a <code>javax.jms.ExceptionListener</code> type.		ExceptionListener
<code>camel.component.jms.explicit-qos-enabled</code>	Set if the <code>deliveryMode</code> , <code>priority</code> or <code>timeToLive</code> qualities of service should be used when sending messages. This option is based on Spring's <code>JmsTemplate</code> . The <code>deliveryMode</code> , <code>priority</code> and <code>timeToLive</code> options are applied to the current endpoint. This contrasts with the <code>preserveMessageQos</code> option, which operates at message granularity, reading QoS properties exclusively from the Camel In message headers.	false	Boolean
<code>camel.component.jms.expose-listener-session</code>	Specifies whether the listener session should be exposed when consuming messages.	false	Boolean
<code>camel.component.jms.force-send-original-message</code>	When using <code>mapJmsMessage=false</code> Camel will create a new JMS message to send to a new JMS destination if you touch the headers (<code>get</code> or <code>set</code>) during the route. Set this option to true to force Camel to send the original JMS message that was received.	false	Boolean
<code>camel.component.jms.format-date-headers-to-iso8601</code>	Sets whether JMS date properties should be formatted according to the ISO 8601 standard.	false	Boolean
<code>camel.component.jms.header-filter-strategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		HeaderFilterStrategy
<code>camel.component.jms.idle-consumer-limit</code>	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	Integer

Name	Description	Default	Type
camel.component.jms.idle-task-execution-limit	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	Integer
camel.component.jms.include-all-jms-x-properties	Whether to include all JMSXxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as JMSXAppID, and JMSXUserID etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	Boolean
camel.component.jms.include-sent-jms-message-id	Only applicable when sending to JMS destination using <code>InOnly</code> (eg fire and forget). Enabling this option will enrich the Camel Exchange with the actual <code>JMSMessageID</code> that was used by the JMS client when the message was sent to the JMS destination.	false	Boolean
camel.component.jms.jms-key-format-strategy	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: <code>default</code> and <code>passthrough</code> . The default strategy will safely marshal dots and hyphens (<code>.</code> and <code>-</code>). The <code>passthrough</code> strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the <code>#</code> notation.		<code>JmsKeyFormatStrategy</code>
camel.component.jms.jms-message-type	Allows you to force the use of a specific <code>javax.jms.Message</code> implementation for sending JMS messages. Possible values are: <code>Bytes</code> , <code>Map</code> , <code>Object</code> , <code>Stream</code> , <code>Text</code> . By default, Camel would determine which JMS message type to use from the <code>In</code> body type. This option allows you to specify it.		<code>JmsMessageType</code>
camel.component.jms.lazy-create-transaction-manager	If true, Camel will create a <code>JmsTransactionManager</code> , if there is no <code>transactionManager</code> injected when option <code>transacted=true</code> .	true	Boolean

Name	Description	Default	Type
<code>camel.component.jms.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.jms.map-jms-message</code>	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a <code>String</code> etc.	true	Boolean
<code>camel.component.jms.max-concurrent-consumers</code>	Specifies the maximum number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToMaxConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.		Integer
<code>camel.component.jms.max-messages-per-task</code>	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	Integer
<code>camel.component.jms.message-converter</code>	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> . The option is a <code>org.springframework.jms.support.converter.MessageConverter</code> type.		<code>MessageConverter</code>
<code>camel.component.jms.message-created-strategy</code>	To use the given <code>MessageCreatedStrategy</code> which are invoked when Camel creates new instances of <code>javax.jms.Message</code> objects when Camel is sending a JMS message. The option is a <code>org.apache.camel.component.jms.MessageCreatedStrategy</code> type.		<code>MessageCreatedStrategy</code>

Name	Description	Default	Type
camel.component.jms.message-id-enabled	When sending, specifies whether message IDs should be added. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	Boolean
camel.component.jms.message-listener-container-factory	Registry ID of the MessageListenerContainerFactory used to determine what <code>org.springframework.jms.listener.AbstractMessageListenerContainer</code> to use to consume messages. Setting this will automatically set <code>consumerType</code> to <code>Custom</code> . The option is a <code>org.apache.camel.component.jms.MessageListenerContainerFactory</code> type.		MessageListenerContainerFactory
camel.component.jms.message-timestamp-enabled	Specifies whether timestamps should be enabled by default on sending messages. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	Boolean
camel.component.jms.password	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
camel.component.jms.preserve-message-qos	Set to true, if you want to send message using the QoS settings specified on the message, instead of the QoS settings on the JMS endpoint. The following three headers are considered <code>JMSPriority</code> , <code>JMSDeliveryMode</code> , and <code>JMSExpiration</code> . You can provide all or only some of them. If not provided, Camel will fall back to use the values from the endpoint instead. So, when using this option, the headers override the values from the endpoint. The <code>explicitQosEnabled</code> option, by contrast, will only use options set on the endpoint, and not values from the message header.	false	Boolean
camel.component.jms.priority	Values greater than 1 specify the message priority when sending (where 1 is the lowest priority and 9 is the highest). The <code>explicitQosEnabled</code> option must also be enabled in order for this option to have any effect.	4	Integer

Name	Description	Default	Type
<code>camel.component.jms.pub-sub-no-local</code>	Specifies whether to inhibit the delivery of messages published by its own connection.	false	Boolean
<code>camel.component.jms.queue-browse-strategy</code>	To use a custom <code>QueueBrowseStrategy</code> when browsing queues. The option is a <code>org.apache.camel.component.jms.QueueBrowseStrategy</code> type.		<code>QueueBrowseStrategy</code>
<code>camel.component.jms.receive-timeout</code>	The timeout for receiving messages (in milliseconds). The option is a long type.	1000	Long
<code>camel.component.jms.recovery-interval</code>	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds. The option is a long type.	5000	Long
<code>camel.component.jms.reply-to</code>	Provides an explicit <code>ReplyTo</code> destination (overrides any incoming value of <code>Message.getJMSReplyTo()</code> in consumer).		String
<code>camel.component.jms.reply-to-cache-level-name</code>	Sets the cache level by name for the reply consumer when doing request/reply over JMS. This option only applies when using fixed reply queues (not temporary). Camel will by default use: <code>CACHE_CONSUMER</code> for exclusive or shared w/ <code>replyToSelectorName</code> . And <code>CACHE_SESSION</code> for shared without <code>replyToSelectorName</code> . Some JMS brokers such as IBM WebSphere may require to set the <code>replyToCacheLevelName=CACHE_NONE</code> to work. Note: If using temporary queues then <code>CACHE_NONE</code> is not allowed, and you must use a higher value such as <code>CACHE_CONSUMER</code> or <code>CACHE_SESSION</code> .		String
<code>camel.component.jms.reply-to-concurrent-consumers</code>	Specifies the default number of concurrent consumers when doing request/reply over JMS. See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads.	1	Integer
<code>camel.component.jms.reply-to-delivery-persistent</code>	Specifies whether to use persistent delivery by default for replies.	true	Boolean

Name	Description	Default	Type
<code>camel.component.jms.reply-to-destination-selector-name</code>	Sets the JMS Selector using the fixed name to be used so you can filter out your own replies from the others when using a shared queue (that is, if you are not using a temporary reply queue).		String
<code>camel.component.jms.reply-to-max-concurrent-consumers</code>	Specifies the maximum number of concurrent consumers when using request/reply over JMS. See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads.		Integer
<code>camel.component.jms.reply-to-on-timeout-max-concurrent-consumers</code>	Specifies the maximum number of concurrent consumers for continue routing when timeout occurred when using request/reply over JMS.	1	Integer
<code>camel.component.jms.reply-to-override</code>	Provides an explicit <code>ReplyTo</code> destination in the JMS message, which overrides the setting of <code>replyTo</code> . It is useful if you want to forward the message to a remote Queue and receive the reply message from the <code>ReplyTo</code> destination.		String
<code>camel.component.jms.reply-to-same-destination-allowed</code>	Whether a JMS consumer is allowed to send a reply message to the same destination that the consumer is using to consume from. This prevents an endless loop by consuming and sending back the same message to itself.	false	Boolean
<code>camel.component.jms.reply-to-type</code>	Allows for explicitly specifying which kind of strategy to use for <code>replyTo</code> queues when doing request/reply over JMS. Possible values are: <code>Temporary</code> , <code>Shared</code> , or <code>Exclusive</code> . By default Camel will use temporary queues. However if <code>replyTo</code> has been configured, then <code>Shared</code> is used by default. This option allows you to use exclusive queues instead of shared ones. See Camel JMS documentation for more details, and especially the notes about the implications if running in a clustered environment, and the fact that <code>Shared</code> reply queues has lower performance than its alternatives <code>Temporary</code> and <code>Exclusive</code> .		<code>ReplyToType</code>
<code>camel.component.jms.request-timeout</code>	The timeout for waiting for a reply when using the InOut Exchange Pattern (in milliseconds). The default is 20 seconds. You can include the header <code>CamelJmsRequestTimeout</code> to override this endpoint configured timeout value, and thus have per message individual timeout values. See also the <code>requestTimeoutCheckerInterval</code> option. The option is a long type.	20000	Long

Name	Description	Default	Type
camel.component.jms.request-timeout-checker-interval	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout. The option is a long type.	1000	Long
camel.component.jms.selector	Sets the JMS selector to use.		String
camel.component.jms.stream-message-type-enabled	Sets whether StreamMessage type is enabled or not. Message payloads of streaming kind such as files, InputStream, etc will either be sent as BytesMessage or StreamMessage. This option controls which kind will be used. By default BytesMessage is used which enforces the entire message payload to be read into memory. By enabling this option the message payload is read into memory in chunks and each chunk is then written to the StreamMessage until no more data.	false	Boolean
camel.component.jms.subscription-durable	Set whether to make the subscription durable. The durable subscription name to be used can be specified through the subscriptionName property. Default is false. Set this to true to register a durable subscription, typically in combination with a subscriptionName value (unless your message listener class name is good enough as subscription name). Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the pubSubDomain flag as well.	false	Boolean
camel.component.jms.subscription-name	Set the name of a subscription to create. To be applied in case of a topic (pub-sub domain) with a shared or durable subscription. The subscription name needs to be unique within this client's JMS client id. Default is the class name of the specified message listener. Note: Only 1 concurrent consumer (which is the default of this message listener container) is allowed for each subscription, except for a shared subscription (which requires JMS 2.0).		String

Name	Description	Default	Type
<code>camel.component.jms.subscription-shared</code>	Set whether to make the subscription shared. The shared subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a shared subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Note that shared subscriptions may also be durable, so this flag can (and often will) be combined with <code>subscriptionDurable</code> as well. Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well. Requires a JMS 2.0 compatible message broker.	false	Boolean
<code>camel.component.jms.synchronous</code>	Sets whether synchronous processing should be strictly used.	false	Boolean
<code>camel.component.jms.task-executor</code>	Allows you to specify a custom task executor for consuming messages. The option is a <code>org.springframework.core.task.TaskExecutor</code> type.		TaskExecutor
<code>camel.component.jms.test-connection-on-startup</code>	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	Boolean
<code>camel.component.jms.time-to-live</code>	When sending messages, specifies the time-to-live of the message (in milliseconds).	-1	Long
<code>camel.component.jms.transacted</code>	Specifies whether to use transacted mode.	false	Boolean

Name	Description	Default	Type
camel.component.jms.transacted-in-out	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	Boolean
camel.component.jms.transaction-manager	The Spring transaction manager to use. The option is a <code>org.springframework.transaction.PlatformTransactionManager</code> type.		PlatformTransactionManager
camel.component.jms.transaction-name	The name of the transaction to use.		String
camel.component.jms.transaction-timeout	The timeout value of the transaction (in seconds), if using transacted mode.	-1	Integer

Name	Description	Default	Type
camel.component.jms.transfer-exception	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a javax.jms.ObjectMessage. If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have transferExchange enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as org.apache.camel.RuntimeCamelException when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!.	false	Boolean
camel.component.jms.transfer-exchange	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payloads is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the receiver to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumers having to use compatible Camel versions!.	false	Boolean
camel.component.jms.use-message-i-d-as-correlation-i-d	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	Boolean
camel.component.jms.username	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.		String
camel.component.jms.wait-for-provision-correlation-to-be-updated-counter	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option useMessageIDAsCorrelationID is enabled.	50	Integer

Name	Description	Default	Type
<code>camel.component.jms.wait-for-provision-correlation-to-be-updated-thread-sleeping-time</code>	Interval in millis to sleep each time while waiting for provisional correlation id to be updated. The option is a long type.	100	Long

CHAPTER 27. KAFKA

Both producer and consumer are supported

The Kafka component is used for communicating with [Apache Kafka](#) message broker.

Maven users will need to add the following dependency to their **pom.xml** for this component.

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-kafka</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

27.1. URI FORMAT

```
kafka:topic[?options]
```

27.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

27.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

27.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

27.3. COMPONENT OPTIONS

The Kafka component supports 104 options, which are listed below.

Name	Description	Default	Type
additionalProperties (common)	Sets additional properties for either kafka consumer or kafka producer in case they can't be set directly on the camel configurations (e.g: new Kafka properties that are not reflected yet in Camel configurations), the properties have to be prefixed with <code>additionalProperties.</code> . E.g: <code>additionalProperties.transactional.id=12345&additionalProperties.schema.registry.url=http://localhost:8811/avro.</code>		Map
brokers (common)	URL of the Kafka brokers to use. The format is <code>host1:port1,host2:port2</code> , and the list can be a subset of brokers or a VIP pointing to a subset of brokers. This option is known as <code>bootstrap.servers</code> in the Kafka documentation.		String
clientId (common)	The client id is a user-specified string sent in each request to help trace calls. It should logically identify the application making the request.		String
configuration (common)	Allows to pre-configure the Kafka component with common options that the endpoints will reuse.		KafkaConfiguration
headerFilterStrategy (common)	To use a custom <code>HeaderFilterStrategy</code> to filter header to and from Camel message.		HeaderFilterStrategy
reconnectBackoffMaxMs (common)	The maximum amount of time in milliseconds to wait when reconnecting to a broker that has repeatedly failed to connect. If provided, the backoff per host will increase exponentially for each consecutive connection failure, up to this maximum. After calculating the backoff increase, 20% random jitter is added to avoid connection storms.	1000	Integer
shutdownTimeout (common)	Timeout in milliseconds to wait gracefully for the consumer or producer to shutdown and terminate its worker threads.	30000	int

Name	Description	Default	Type
allowManualCommit (consumer)	Whether to allow doing manual commits via <code>KafkaManualCommit</code> . If this option is enabled then an instance of <code>KafkaManualCommit</code> is stored on the Exchange message header, which allows end users to access this API and perform manual offset commits via the Kafka consumer.	false	boolean
autoCommitEnable (consumer)	If true, periodically commit to ZooKeeper the offset of messages already fetched by the consumer. This committed offset will be used when the process fails as the position from which the new consumer will begin.	true	Boolean
autoCommitIntervalMs (consumer)	The frequency in ms that the consumer offsets are committed to zookeeper.	5000	Integer
autoCommitOnStop (consumer)	Whether to perform an explicit auto commit when the consumer stops to ensure the broker has a commit from the last consumed message. This requires the option <code>autoCommitEnable</code> is turned on. The possible values are: <code>sync</code> , <code>async</code> , or <code>none</code> . And <code>sync</code> is the default value. Enum values: <ul style="list-style-type: none"> ● <code>sync</code> ● <code>async</code> ● <code>none</code> 	sync	String
autoOffsetReset (consumer)	What to do when there is no initial offset in ZooKeeper or if an offset is out of range: <code>earliest</code> : automatically reset the offset to the earliest offset <code>latest</code> : automatically reset the offset to the latest offset <code>fail</code> : throw exception to the consumer. Enum values: <ul style="list-style-type: none"> ● <code>latest</code> ● <code>earliest</code> ● <code>none</code> 	latest	String

Name	Description	Default	Type
breakOnFirstError (consumer)	This options controls what happens when a consumer is processing an exchange and it fails. If the option is false then the consumer continues to the next message and processes it. If the option is true then the consumer breaks out, and will seek back to offset of the message that caused a failure, and then re-attempt to process this message. However this can lead to endless processing of the same message if its bound to fail every time, eg a poison message. Therefore its recommended to deal with that for example by using Camel's error handler.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
checkCrcs (consumer)	Automatically check the CRC32 of the records consumed. This ensures no on-the-wire or on-disk corruption to the messages occurred. This check adds some overhead, so it may be disabled in cases seeking extreme performance.	true	Boolean
commitTimeoutMs (consumer)	The maximum time, in milliseconds, that the code will wait for a synchronous commit to complete.	5000	Long
consumerRequestTimeoutMs (consumer)	The configuration controls the maximum amount of time the client will wait for the response of a request. If the response is not received before the timeout elapses the client will resend the request if necessary or fail the request if retries are exhausted.	40000	Integer
consumersCount (consumer)	The number of consumers that connect to kafka server. Each consumer is run on a separate thread, that retrieves and process the incoming data.	1	int

Name	Description	Default	Type
fetchMaxBytes (consumer)	The maximum amount of data the server should return for a fetch request. This is not an absolute maximum, if the first message in the first non-empty partition of the fetch is larger than this value, the message will still be returned to ensure that the consumer can make progress. The maximum message size accepted by the broker is defined via <code>message.max.bytes</code> (broker config) or <code>max.message.bytes</code> (topic config). Note that the consumer performs multiple fetches in parallel.	52428800	Integer
fetchMinBytes (consumer)	The minimum amount of data the server should return for a fetch request. If insufficient data is available the request will wait for that much data to accumulate before answering the request.	1	Integer
fetchWaitMaxMs (consumer)	The maximum amount of time the server will block before answering the fetch request if there isn't sufficient data to immediately satisfy <code>fetch.min.bytes</code> .	500	Integer
groupId (consumer)	A string that uniquely identifies the group of consumer processes to which this consumer belongs. By setting the same group id multiple processes indicate that they are all part of the same consumer group. This option is required for consumers.		String
groupInstanceId (consumer)	A unique identifier of the consumer instance provided by the end user. Only non-empty strings are permitted. If set, the consumer is treated as a static member, which means that only one instance with this ID is allowed in the consumer group at any time. This can be used in combination with a larger session timeout to avoid group rebalances caused by transient unavailability (e.g. process restarts). If not set, the consumer will join the group as a dynamic member, which is the traditional behavior.		String
headerDeserializer (consumer)	To use a custom <code>KafkaHeaderDeserializer</code> to deserialize kafka headers values.		<code>KafkaHeaderDeserializer</code>

Name	Description	Default	Type
heartbeatIntervalMs (consumer)	The expected time between heartbeats to the consumer coordinator when using Kafka's group management facilities. Heartbeats are used to ensure that the consumer's session stays active and to facilitate rebalancing when new consumers join or leave the group. The value must be set lower than <code>session.timeout.ms</code> , but typically should be set no higher than 1/3 of that value. It can be adjusted even lower to control the expected time for normal rebalances.	3000	Integer
keyDeserializer (consumer)	Deserializer class for key that implements the <code>Deserializer</code> interface.	<code>org.apache.kafka.common.serialization.StringDeserializer</code>	String
maxPartitionFetchBytes (consumer)	The maximum amount of data per-partition the server will return. The maximum total memory used for a request will be <code>#partitions max.partition.fetch.bytes</code> . This size must be at least as large as the maximum message size the server allows or else it is possible for the producer to send messages larger than the consumer can fetch. If that happens, the consumer can get stuck trying to fetch a large message on a certain partition.	1048576	Integer
maxPollIntervalMs (consumer)	The maximum delay between invocations of <code>poll()</code> when using consumer group management. This places an upper bound on the amount of time that the consumer can be idle before fetching more records. If <code>poll()</code> is not called before expiration of this timeout, then the consumer is considered failed and the group will rebalance in order to reassign the partitions to another member.		Long
maxPollRecords (consumer)	The maximum number of records returned in a single call to <code>poll()</code> .	500	Integer
offsetRepository (consumer)	The offset repository to use in order to locally store the offset of each partition of the topic. Defining one will disable the <code>autocommit</code> .		<code>StateRepository</code>

Name	Description	Default	Type
partitionAssignor (consumer)	The class name of the partition assignment strategy that the client will use to distribute partition ownership amongst consumer instances when group management is used.	org.apache.kafka.clients.consumer.RangeAssignor	String
pollOnError (consumer)	<p>What to do if kafka threw an exception while polling for new messages. Will by default use the value from the component configuration unless an explicit value has been configured on the endpoint level. DISCARD will discard the message and continue to poll next message. ERROR_HANDLER will use Camel's error handler to process the exception, and afterwards continue to poll next message. RECONNECT will reconnect the consumer and try poll the message again. RETRY will let the consumer retry polling the same message again. STOP will stop the consumer (have to be manually started/restarted if the consumer should be able to consume messages again).</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● DISCARD ● ERROR_HANDLER ● RECONNECT ● RETRY ● STOP 	ERROR_HANDLER	PollOnError
pollTimeoutMs (consumer)	The timeout used when polling the KafkaConsumer.	5000	Long
resumeStrategy (consumer)	This option allows the user to set a custom resume strategy. The resume strategy is executed when partitions are assigned (i.e.: when connecting or reconnecting). It allows implementations to customize how to resume operations and serve as more flexible alternative to the seekTo and the offsetRepository mechanisms. See the KafkaConsumerResumeStrategy for implementation details. This option does not affect the auto commit setting. It is likely that implementations using this setting will also want to evaluate using the manual commit option along with this.		KafkaConsumerResumeStrategy

Name	Description	Default	Type
seekTo (consumer)	Set if KafkaConsumer will read from beginning or end on startup: beginning : read from beginning end : read from end This is replacing the earlier property seekToBeginning. Enum values: <ul style="list-style-type: none"> ● beginning ● end 		String
sessionTimeoutMs (consumer)	The timeout used to detect failures when using Kafka's group management facilities.	10000	Integer
specificAvroReader (consumer)	This enables the use of a specific Avro reader for use with the Confluent Platform schema registry and the io.confluent.kafka.serializers.KafkaAvroDeserializer. This option is only available in the Confluent Platform (not standard Apache Kafka).	false	boolean
topicsPattern (consumer)	Whether the topic is a pattern (regular expression). This can be used to subscribe to dynamic number of topics matching the pattern.	false	boolean
valueDeserializer (consumer)	Deserializer class for value that implements the Deserializer interface.	org.apache.kafka.common.serialization.StringDeserializer	String
kafkaManualCommitFactory (consumer (advanced))	Autowired Factory to use for creating KafkaManualCommit instances. This allows to plugin a custom factory to create custom KafkaManualCommit instances in case special logic is needed when doing manual commits that deviates from the default implementation that comes out of the box.		KafkaManualCommitFactory
pollExceptionStrategy (consumer (advanced))	Autowired To use a custom strategy with the consumer to control how to handle exceptions thrown from the Kafka broker while pooling messages.		PollExceptionStrategy

Name	Description	Default	Type
bufferMemorySize (producer)	The total bytes of memory the producer can use to buffer records waiting to be sent to the server. If records are sent faster than they can be delivered to the server the producer will either block or throw an exception based on the preference specified by <code>block.on.buffer.full</code> . This setting should correspond roughly to the total memory the producer will use, but is not a hard bound since not all memory the producer uses is used for buffering. Some additional memory will be used for compression (if compression is enabled) as well as for maintaining in-flight requests.	33554432	Integer
compressionCodec (producer)	This parameter allows you to specify the compression codec for all data generated by this producer. Valid values are <code>none</code> , <code>gzip</code> and <code>snappy</code> . Enum values: <ul style="list-style-type: none"> • <code>none</code> • <code>gzip</code> • <code>snappy</code> • <code>lz4</code> 	<code>none</code>	String
connectionMaxIdleMs (producer)	Close idle connections after the number of milliseconds specified by this config.	540000	Integer
deliveryTimeoutMs (producer)	An upper bound on the time to report success or failure after a call to <code>send()</code> returns. This limits the total time that a record will be delayed prior to sending, the time to await acknowledgement from the broker (if expected), and the time allowed for retrievable send failures.	120000	Integer
enableIdempotence (producer)	If set to <code>'true'</code> the producer will ensure that exactly one copy of each message is written in the stream. If <code>'false'</code> , producer retries may write duplicates of the retried message in the stream. If set to <code>true</code> this option will require <code>max.in.flight.requests.per.connection</code> to be set to 1 and retries cannot be zero and additionally <code>acks</code> must be set to <code>'all'</code> .	<code>false</code>	boolean
headerSerializer (producer)	To use a custom <code>KafkaHeaderSerializer</code> to serialize kafka headers values.		<code>KafkaHeaderSerializer</code>

Name	Description	Default	Type
key (producer)	The record key (or null if no key is specified). If this option has been configured then it take precedence over header <code>KafkaConstants#KEY</code> .		String
keySerializer (producer)	The serializer class for keys (defaults to the same as for messages if nothing is given).	<code>org.apache.kafka.common.serialization.StringSerializer</code>	String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
lingerMs (producer)	The producer groups together any records that arrive in between request transmissions into a single batched request. Normally this occurs only under load when records arrive faster than they can be sent out. However in some circumstances the client may want to reduce the number of requests even under moderate load. This setting accomplishes this by adding a small amount of artificial delay that is, rather than immediately sending out a record the producer will wait for up to the given delay to allow other records to be sent so that the sends can be batched together. This can be thought of as analogous to Nagle's algorithm in TCP. This setting gives the upper bound on the delay for batching: once we get batch.size worth of records for a partition it will be sent immediately regardless of this setting, however if we have fewer than this many bytes accumulated for this partition we will 'linger' for the specified time waiting for more records to show up. This setting defaults to 0 (i.e. no delay). Setting <code>linger.ms=5</code> , for example, would have the effect of reducing the number of requests sent but would add up to 5ms of latency to records sent in the absence of load.	0	Integer

Name	Description	Default	Type
maxBlockMs (producer)	The configuration controls how long sending to kafka will block. These methods can be blocked for multiple reasons. For e.g: buffer full, metadata unavailable. This configuration imposes maximum limit on the total time spent in fetching metadata, serialization of key and value, partitioning and allocation of buffer memory when doing a send(). In case of partitionsFor(), this configuration imposes a maximum time threshold on waiting for metadata.	60000	Integer
maxInFlightRequest (producer)	The maximum number of unacknowledged requests the client will send on a single connection before blocking. Note that if this setting is set to be greater than 1 and there are failed sends, there is a risk of message re-ordering due to retries (i.e., if retries are enabled).	5	Integer
maxRequestSize (producer)	The maximum size of a request. This is also effectively a cap on the maximum record size. Note that the server has its own cap on record size which may be different from this. This setting will limit the number of record batches the producer will send in a single request to avoid sending huge requests.	1048576	Integer
metadataMaxAgeMs (producer)	The period of time in milliseconds after which we force a refresh of metadata even if we haven't seen any partition leadership changes to proactively discover any new brokers or partitions.	300000	Integer
metricReporters (producer)	A list of classes to use as metrics reporters. Implementing the MetricReporter interface allows plugging in classes that will be notified of new metric creation. The JmxReporter is always included to register JMX statistics.		String
metricsSampleWindowMs (producer)	The number of samples maintained to compute metrics.	30000	Integer
noOfMetricsSample (producer)	The number of samples maintained to compute metrics.	2	Integer

Name	Description	Default	Type
partitioner (producer)	The partitioner class for partitioning messages amongst sub-topics. The default partitioner is based on the hash of the key.	org.apache.kafka.clients.producer.internals.DefaultPartitioner	String
partitionKey (producer)	The partition to which the record will be sent (or null if no partition was specified). If this option has been configured then it take precedence over header <code>KafkaConstants#PARTITION_KEY</code> .		Integer
producerBatchSize (producer)	The producer will attempt to batch records together into fewer requests whenever multiple records are being sent to the same partition. This helps performance on both the client and the server. This configuration controls the default batch size in bytes. No attempt will be made to batch records larger than this size. Requests sent to brokers will contain multiple batches, one for each partition with data available to be sent. A small batch size will make batching less common and may reduce throughput (a batch size of zero will disable batching entirely). A very large batch size may use memory a bit more wastefully as we will always allocate a buffer of the specified batch size in anticipation of additional records.	16384	Integer
queueBufferingMaxMessages (producer)	The maximum number of unsent messages that can be queued up the producer when using async mode before either the producer must be blocked or data must be dropped.	10000	Integer
receiveBufferBytes (producer)	The size of the TCP receive buffer (<code>SO_RCVBUF</code>) to use when reading data.	65536	Integer
reconnectBackoffMs (producer)	The amount of time to wait before attempting to reconnect to a given host. This avoids repeatedly connecting to a host in a tight loop. This backoff applies to all requests sent by the consumer to the broker.	50	Integer

Name	Description	Default	Type
recordMetadata (producer)	Whether the producer should store the RecordMetadata results from sending to Kafka. The results are stored in a List containing the RecordMetadata metadata's. The list is stored on a header with the key <code>KafkaConstants#KAFKA_RECORDMETA</code> .	true	boolean
requestRequiredAcks (producer)	<p>The number of acknowledgments the producer requires the leader to have received before considering a request complete. This controls the durability of records that are sent. The following settings are common: <code>acks=0</code> If set to zero then the producer will not wait for any acknowledgment from the server at all. The record will be immediately added to the socket buffer and considered sent. No guarantee can be made that the server has received the record in this case, and the retries configuration will not take effect (as the client won't generally know of any failures). The offset given back for each record will always be set to -1. <code>acks=1</code> This will mean the leader will write the record to its local log but will respond without awaiting full acknowledgement from all followers. In this case should the leader fail immediately after acknowledging the record but before the followers have replicated it then the record will be lost. <code>acks=all</code> This means the leader will wait for the full set of in-sync replicas to acknowledge the record. This guarantees that the record will not be lost as long as at least one in-sync replica remains alive. This is the strongest available guarantee.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● -1 ● 0 ● 1 ● all 	1	String
requestTimeoutMs (producer)	The amount of time the broker will wait trying to meet the <code>request.required.acks</code> requirement before sending back an error to the client.	30000	Integer

Name	Description	Default	Type
retries (producer)	Setting a value greater than zero will cause the client to resend any record whose send fails with a potentially transient error. Note that this retry is no different than if the client resent the record upon receiving the error. Allowing retries will potentially change the ordering of records because if two records are sent to a single partition, and the first fails and is retried but the second succeeds, then the second record may appear first.	0	Integer
retryBackoffMs (producer)	Before each retry, the producer refreshes the metadata of relevant topics to see if a new leader has been elected. Since leader election takes a bit of time, this property specifies the amount of time that the producer waits before refreshing the metadata.	100	Integer
sendBufferBytes (producer)	Socket write buffer size.	131072	Integer
valueSerializer (producer)	The serializer class for messages.	org.apache.kafka.common.serialization.StringSerializer	String
workerPool (producer)	To use a custom worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing. If using this option then you must handle the lifecycle of the thread pool to shut the pool down when no longer needed.		ExecutorService
workerPoolCoreSize (producer)	Number of core threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	10	Integer
workerPoolMaxSize (producer)	Maximum number of threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	20	Integer

Name	Description	Default	Type
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
kafkaClientFactory (advanced)	Autowired Factory to use for creating <code>org.apache.kafka.clients.consumer.KafkaConsumer</code> and <code>org.apache.kafka.clients.producer.KafkaProducer</code> instances. This allows to configure a custom factory to create instances with logic that extends the vanilla Kafka clients.		KafkaClientFactory
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
schemaRegistryURL (confluent)	URL of the Confluent Platform schema registry servers to use. The format is <code>host1:port1,host2:port2</code> . This is known as <code>schema.registry.url</code> in the Confluent Platform documentation. This option is only available in the Confluent Platform (not standard Apache Kafka).		String
interceptorClasses (monitoring)	Sets interceptors for producer or consumers. Producer interceptors have to be classes implementing <code>org.apache.kafka.clients.producer.ProducerInterceptor</code> or Consumer interceptors have to be classes implementing <code>org.apache.kafka.clients.consumer.ConsumerInterceptor</code> . Note that if you use Producer interceptor on a consumer it will throw a class cast exception in runtime.		String
kerberosBeforeReloginMinTime (security)	Login thread sleep time between refresh attempts.	60000	Integer
kerberosInitCmd (security)	Kerberos kinit command path. Default is <code>/usr/bin/kinit</code> .	<code>/usr/bin/kinit</code>	String

Name	Description	Default	Type
kerberosPrincipalToLocalRules (security)	A list of rules for mapping from principal names to short names (typically operating system usernames). The rules are evaluated in order and the first rule that matches a principal name is used to map it to a short name. Any later rules in the list are ignored. By default, principal names of the form {username}/{hostname}{REALM} are mapped to {username}. For more details on the format please see the security authorization and acls documentation.. Multiple values can be separated by comma.	DEFAULT	String
kerberosRenewJitter (security)	Percentage of random jitter added to the renewal time.	0.05	Double
kerberosRenewWindowFactor (security)	Login thread will sleep until the specified window factor of time from last refresh to ticket's expiry has been reached, at which time it will try to renew the ticket.	0.8	Double
saslJaasConfig (security)	Expose the kafka sasl.jaas.config parameter Example: org.apache.kafka.common.security.plain.PlainLoginModule required username=USERNAME password=PASSWORD;.		String
saslKerberosServiceName (security)	The Kerberos principal name that Kafka runs as. This can be defined either in Kafka's JAAS config or in Kafka's config.		String
saslMechanism (security)	The Simple Authentication and Security Layer (SASL) Mechanism used. For the valid values see .	GSSAPI	String
securityProtocol (security)	Protocol used to communicate with brokers. SASL_PLAINTEXT, PLAINTEXT and SSL are supported.	PLAINTEXT	String
sslCipherSuites (security)	A list of cipher suites. This is a named combination of authentication, encryption, MAC and key exchange algorithm used to negotiate the security settings for a network connection using TLS or SSL network protocol. By default all the available cipher suites are supported.		String

Name	Description	Default	Type
sslContextParameters (security)	SSL configuration using a Camel SSLContextParameters object. If configured it's applied before the other SSL endpoint parameters. NOTE: Kafka only supports loading keystore from file locations, so prefix the location with file: in the KeyStoreParameters.resource option.		SSLContextParameters
sslEnabledProtocols (security)	The list of protocols enabled for SSL connections. TLSv1.2, TLSv1.1 and TLSv1 are enabled by default.		String
sslEndpointAlgorithm (security)	The endpoint identification algorithm to validate server hostname using server certificate.	https	String
sslKeymanagerAlgorithm (security)	The algorithm used by key manager factory for SSL connections. Default value is the key manager factory algorithm configured for the Java Virtual Machine.	SunX509	String
sslKeyPassword (security)	The password of the private key in the key store file. This is optional for client.		String
sslKeystoreLocation (security)	The location of the key store file. This is optional for client and can be used for two-way authentication for client.		String
sslKeystorePassword (security)	The store password for the key store file. This is optional for client and only needed if ssl.keystore.location is configured.		String
sslKeystoreType (security)	The file format of the key store file. This is optional for client. Default value is JKS.	JKS	String
sslProtocol (security)	The SSL protocol used to generate the SSLContext. Default setting is TLS, which is fine for most cases. Allowed values in recent JVMs are TLS, TLSv1.1 and TLSv1.2. SSL, SSLv2 and SSLv3 may be supported in older JVMs, but their usage is discouraged due to known security vulnerabilities.		String
sslProvider (security)	The name of the security provider used for SSL connections. Default value is the default security provider of the JVM.		String

Name	Description	Default	Type
sslTrustmanagerAlgorithm (security)	The algorithm used by trust manager factory for SSL connections. Default value is the trust manager factory algorithm configured for the Java Virtual Machine.	PKIX	String
sslTruststoreLocation (security)	The location of the trust store file.		String
sslTruststorePassword (security)	The password for the trust store file.		String
sslTruststoreType (security)	The file format of the trust store file. Default value is JKS.	JKS	String
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean

27.4. ENDPOINT OPTIONS

The Kafka endpoint is configured using URI syntax:

```
kafka:topic
```

with the following path and query parameters:

27.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
topic (common)	Required Name of the topic to use. On the consumer you can use comma to separate multiple topics. A producer can only send a message to a single topic.		String

27.4.2. Query Parameters (102 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
additionalProperties (common)	Sets additional properties for either kafka consumer or kafka producer in case they can't be set directly on the camel configurations (e.g: new Kafka properties that are not reflected yet in Camel configurations), the properties have to be prefixed with <code>additionalProperties</code> .. E.g: <code>additionalProperties.transactional.id=12345&additionalProperties.schema.registry.url=http://localhost:8811/avro</code> .		Map
brokers (common)	URL of the Kafka brokers to use. The format is <code>host1:port1,host2:port2</code> , and the list can be a subset of brokers or a VIP pointing to a subset of brokers. This option is known as <code>bootstrap.servers</code> in the Kafka documentation.		String
clientId (common)	The client id is a user-specified string sent in each request to help trace calls. It should logically identify the application making the request.		String
headerFilterStrategy (common)	To use a custom <code>HeaderFilterStrategy</code> to filter header to and from Camel message.		<code>HeaderFilterStrategy</code>
reconnectBackoffMaxMs (common)	The maximum amount of time in milliseconds to wait when reconnecting to a broker that has repeatedly failed to connect. If provided, the backoff per host will increase exponentially for each consecutive connection failure, up to this maximum. After calculating the backoff increase, 20% random jitter is added to avoid connection storms.	1000	Integer
shutdownTimeout (common)	Timeout in milliseconds to wait gracefully for the consumer or producer to shutdown and terminate its worker threads.	30000	int
allowManualCommit (consumer)	Whether to allow doing manual commits via <code>KafkaManualCommit</code> . If this option is enabled then an instance of <code>KafkaManualCommit</code> is stored on the Exchange message header, which allows end users to access this API and perform manual offset commits via the Kafka consumer.	false	boolean
autoCommitEnable (consumer)	If true, periodically commit to ZooKeeper the offset of messages already fetched by the consumer. This committed offset will be used when the process fails as the position from which the new consumer will begin.	true	Boolean

Name	Description	Default	Type
autoCommitIntervalMs (consumer)	The frequency in ms that the consumer offsets are committed to zookeeper.	5000	Integer
autoCommitOnStop (consumer)	Whether to perform an explicit auto commit when the consumer stops to ensure the broker has a commit from the last consumed message. This requires the option <code>autoCommitEnable</code> is turned on. The possible values are: <code>sync</code> , <code>async</code> , or <code>none</code> . And <code>sync</code> is the default value. Enum values: <ul style="list-style-type: none"> • <code>sync</code> • <code>async</code> • <code>none</code> 	<code>sync</code>	String
autoOffsetReset (consumer)	What to do when there is no initial offset in ZooKeeper or if an offset is out of range: <code>earliest</code> : automatically reset the offset to the earliest offset <code>latest</code> : automatically reset the offset to the latest offset <code>fail</code> : throw exception to the consumer. Enum values: <ul style="list-style-type: none"> • <code>latest</code> • <code>earliest</code> • <code>none</code> 	<code>latest</code>	String
breakOnFirstError (consumer)	This options controls what happens when a consumer is processing an exchange and it fails. If the option is <code>false</code> then the consumer continues to the next message and processes it. If the option is <code>true</code> then the consumer breaks out, and will seek back to offset of the message that caused a failure, and then re-attempt to process this message. However this can lead to endless processing of the same message if its bound to fail every time, eg a poison message. Therefore its recommended to deal with that for example by using Camel's error handler.	<code>false</code>	boolean

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
checkCrcs (consumer)	Automatically check the CRC32 of the records consumed. This ensures no on-the-wire or on-disk corruption to the messages occurred. This check adds some overhead, so it may be disabled in cases seeking extreme performance.	true	Boolean
commitTimeoutMs (consumer)	The maximum time, in milliseconds, that the code will wait for a synchronous commit to complete.	5000	Long
consumerRequestTimeoutMs (consumer)	The configuration controls the maximum amount of time the client will wait for the response of a request. If the response is not received before the timeout elapses the client will resend the request if necessary or fail the request if retries are exhausted.	40000	Integer
consumersCount (consumer)	The number of consumers that connect to kafka server. Each consumer is run on a separate thread, that retrieves and process the incoming data.	1	int
fetchMaxBytes (consumer)	The maximum amount of data the server should return for a fetch request This is not an absolute maximum, if the first message in the first non-empty partition of the fetch is larger than this value, the message will still be returned to ensure that the consumer can make progress. The maximum message size accepted by the broker is defined via <code>message.max.bytes</code> (broker config) or <code>max.message.bytes</code> (topic config). Note that the consumer performs multiple fetches in parallel.	52428800	Integer
fetchMinBytes (consumer)	The minimum amount of data the server should return for a fetch request. If insufficient data is available the request will wait for that much data to accumulate before answering the request.	1	Integer
fetchWaitMaxMs (consumer)	The maximum amount of time the server will block before answering the fetch request if there isn't sufficient data to immediately satisfy <code>fetch.min.bytes</code> .	500	Integer

Name	Description	Default	Type
groupId (consumer)	A string that uniquely identifies the group of consumer processes to which this consumer belongs. By setting the same group id multiple processes indicate that they are all part of the same consumer group. This option is required for consumers.		String
groupId (consumer)	A unique identifier of the consumer instance provided by the end user. Only non-empty strings are permitted. If set, the consumer is treated as a static member, which means that only one instance with this ID is allowed in the consumer group at any time. This can be used in combination with a larger session timeout to avoid group rebalances caused by transient unavailability (e.g. process restarts). If not set, the consumer will join the group as a dynamic member, which is the traditional behavior.		String
headerDeserializer (consumer)	To use a custom <code>KafkaHeaderDeserializer</code> to deserialize kafka headers values.		<code>KafkaHeaderDeserializer</code>
heartbeatIntervalMs (consumer)	The expected time between heartbeats to the consumer coordinator when using Kafka's group management facilities. Heartbeats are used to ensure that the consumer's session stays active and to facilitate rebalancing when new consumers join or leave the group. The value must be set lower than <code>session.timeout.ms</code> , but typically should be set no higher than 1/3 of that value. It can be adjusted even lower to control the expected time for normal rebalances.	3000	Integer
keyDeserializer (consumer)	Deserializer class for key that implements the <code>Deserializer</code> interface.	<code>org.apache.kafka.common.serialization.StringDeserializer</code>	String

Name	Description	Default	Type
maxPartitionFetchBytes (consumer)	The maximum amount of data per-partition the server will return. The maximum total memory used for a request will be <code>#partitions max.partition.fetch.bytes</code> . This size must be at least as large as the maximum message size the server allows or else it is possible for the producer to send messages larger than the consumer can fetch. If that happens, the consumer can get stuck trying to fetch a large message on a certain partition.	1048576	Integer
maxPollIntervalMs (consumer)	The maximum delay between invocations of <code>poll()</code> when using consumer group management. This places an upper bound on the amount of time that the consumer can be idle before fetching more records. If <code>poll()</code> is not called before expiration of this timeout, then the consumer is considered failed and the group will rebalance in order to reassign the partitions to another member.		Long
maxPollRecords (consumer)	The maximum number of records returned in a single call to <code>poll()</code> .	500	Integer
offsetRepository (consumer)	The offset repository to use in order to locally store the offset of each partition of the topic. Defining one will disable the autocommit.		StateRepository
partitionAssignor (consumer)	The class name of the partition assignment strategy that the client will use to distribute partition ownership amongst consumer instances when group management is used.	<code>org.apache.kafka.clients.consumer.RangeAssignor</code>	String

Name	Description	Default	Type
pollOnError (consumer)	<p>What to do if kafka threw an exception while polling for new messages. Will by default use the value from the component configuration unless an explicit value has been configured on the endpoint level. DISCARD will discard the message and continue to poll next message. ERROR_HANDLER will use Camel's error handler to process the exception, and afterwards continue to poll next message. RECONNECT will reconnect the consumer and try poll the message again. RETRY will let the consumer retry polling the same message again. STOP will stop the consumer (have to be manually started/restarted if the consumer should be able to consume messages again).</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● DISCARD ● ERROR_HANDLER ● RECONNECT ● RETRY ● STOP 	ERROR_HANDLER	PollOnError
pollTimeoutMs (consumer)	The timeout used when polling the KafkaConsumer.	5000	Long
resumeStrategy (consumer)	<p>This option allows the user to set a custom resume strategy. The resume strategy is executed when partitions are assigned (i.e.: when connecting or reconnecting). It allows implementations to customize how to resume operations and serve as more flexible alternative to the seekTo and the offsetRepository mechanisms. See the KafkaConsumerResumeStrategy for implementation details. This option does not affect the auto commit setting. It is likely that implementations using this setting will also want to evaluate using the manual commit option along with this.</p>		KafkaConsumerResumeStrategy

Name	Description	Default	Type
seekTo (consumer)	Set if KafkaConsumer will read from beginning or end on startup: beginning : read from beginning end : read from end This is replacing the earlier property seekToBeginning. Enum values: <ul style="list-style-type: none"> • beginning • end 		String
sessionTimeoutMs (consumer)	The timeout used to detect failures when using Kafka's group management facilities.	10000	Integer
specificAvroReader (consumer)	This enables the use of a specific Avro reader for use with the Confluent Platform schema registry and the io.confluent.kafka.serializers.KafkaAvroDeserializer. This option is only available in the Confluent Platform (not standard Apache Kafka).	false	boolean
topicsPattern (consumer)	Whether the topic is a pattern (regular expression). This can be used to subscribe to dynamic number of topics matching the pattern.	false	boolean
valueDeserializer (consumer)	Deserializer class for value that implements the Deserializer interface.	org.apache.kafka.common.serialization.StringDeserializer	String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	<p>Sets the exchange pattern when the consumer creates an exchange.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
kafkaManualCommitFactory (consumer (advanced))	<p>Factory to use for creating KafkaManualCommit instances. This allows to plugin a custom factory to create custom KafkaManualCommit instances in case special logic is needed when doing manual commits that deviates from the default implementation that comes out of the box.</p>		KafkaManualCommitFactory
bufferMemorySize (producer)	<p>The total bytes of memory the producer can use to buffer records waiting to be sent to the server. If records are sent faster than they can be delivered to the server the producer will either block or throw an exception based on the preference specified by <code>block.on.buffer.full</code>. This setting should correspond roughly to the total memory the producer will use, but is not a hard bound since not all memory the producer uses is used for buffering. Some additional memory will be used for compression (if compression is enabled) as well as for maintaining in-flight requests.</p>	33554432	Integer
compressionCodec (producer)	<p>This parameter allows you to specify the compression codec for all data generated by this producer. Valid values are none, gzip and snappy.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● none ● gzip ● snappy ● lz4 	none	String
connectionMaxIdleMs (producer)	<p>Close idle connections after the number of milliseconds specified by this config.</p>	540000	Integer

Name	Description	Default	Type
deliveryTimeoutMs (producer)	An upper bound on the time to report success or failure after a call to <code>send()</code> returns. This limits the total time that a record will be delayed prior to sending, the time to await acknowledgement from the broker (if expected), and the time allowed for retrievable send failures.	120000	Integer
enableIdempotence (producer)	If set to 'true' the producer will ensure that exactly one copy of each message is written in the stream. If 'false', producer retries may write duplicates of the retried message in the stream. If set to true this option will require <code>max.in.flight.requests.per.connection</code> to be set to 1 and retries cannot be zero and additionally <code>acks</code> must be set to 'all'.	false	boolean
headerSerializer (producer)	To use a custom <code>KafkaHeaderSerializer</code> to serialize kafka headers values.		<code>KafkaHeaderSerializer</code>
key (producer)	The record key (or null if no key is specified). If this option has been configured then it take precedence over header <code>KafkaConstants#KEY</code> .		String
keySerializer (producer)	The serializer class for keys (defaults to the same as for messages if nothing is given).	<code>org.apache.kafka.common.serialization.StringSerializer</code>	String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
lingerMs (producer)	The producer groups together any records that arrive in between request transmissions into a single batched request. Normally this occurs only under load when records arrive faster than they can be sent out. However in some circumstances the client may want to reduce the number of requests even under moderate load. This setting accomplishes this by adding a small amount of artificial delay that is, rather than immediately sending out a record the producer will wait for up to the given delay to allow other records to be sent so that the sends can be batched together. This can be thought of as analogous to Nagle's algorithm in TCP. This setting gives the upper bound on the delay for batching: once we get batch.size worth of records for a partition it will be sent immediately regardless of this setting, however if we have fewer than this many bytes accumulated for this partition we will 'linger' for the specified time waiting for more records to show up. This setting defaults to 0 (i.e. no delay). Setting <code>linger.ms=5</code> , for example, would have the effect of reducing the number of requests sent but would add up to 5ms of latency to records sent in the absence of load.	0	Integer
maxBlockMs (producer)	The configuration controls how long sending to kafka will block. These methods can be blocked for multiple reasons. For e.g: buffer full, metadata unavailable. This configuration imposes maximum limit on the total time spent in fetching metadata, serialization of key and value, partitioning and allocation of buffer memory when doing a <code>send()</code> . In case of <code>partitionsFor()</code> , this configuration imposes a maximum time threshold on waiting for metadata.	60000	Integer
maxInFlightRequest (producer)	The maximum number of unacknowledged requests the client will send on a single connection before blocking. Note that if this setting is set to be greater than 1 and there are failed sends, there is a risk of message re-ordering due to retries (i.e., if retries are enabled).	5	Integer
maxRequestSize (producer)	The maximum size of a request. This is also effectively a cap on the maximum record size. Note that the server has its own cap on record size which may be different from this. This setting will limit the number of record batches the producer will send in a single request to avoid sending huge requests.	1048576	Integer

Name	Description	Default	Type
metadataMaxAgeMs (producer)	The period of time in milliseconds after which we force a refresh of metadata even if we haven't seen any partition leadership changes to proactively discover any new brokers or partitions.	300000	Integer
metricReporters (producer)	A list of classes to use as metrics reporters. Implementing the MetricReporter interface allows plugging in classes that will be notified of new metric creation. The JmxReporter is always included to register JMX statistics.		String
metricsSampleWindowMs (producer)	The number of samples maintained to compute metrics.	30000	Integer
noOfMetricsSample (producer)	The number of samples maintained to compute metrics.	2	Integer
partitioner (producer)	The partitioner class for partitioning messages amongst sub-topics. The default partitioner is based on the hash of the key.	org.apache.kafka.clients.producer.internals.DefaultPartitioner	String
partitionKey (producer)	The partition to which the record will be sent (or null if no partition was specified). If this option has been configured then it take precedence over header <code>KafkaConstants#PARTITION_KEY</code> .		Integer
producerBatchSize (producer)	The producer will attempt to batch records together into fewer requests whenever multiple records are being sent to the same partition. This helps performance on both the client and the server. This configuration controls the default batch size in bytes. No attempt will be made to batch records larger than this size. Requests sent to brokers will contain multiple batches, one for each partition with data available to be sent. A small batch size will make batching less common and may reduce throughput (a batch size of zero will disable batching entirely). A very large batch size may use memory a bit more wastefully as we will always allocate a buffer of the specified batch size in anticipation of additional records.	16384	Integer

Name	Description	Default	Type
queueBufferingMaxMessages (producer)	The maximum number of unsent messages that can be queued up the producer when using async mode before either the producer must be blocked or data must be dropped.	10000	Integer
receiveBufferBytes (producer)	The size of the TCP receive buffer (SO_RCVBUF) to use when reading data.	65536	Integer
reconnectBackoffMs (producer)	The amount of time to wait before attempting to reconnect to a given host. This avoids repeatedly connecting to a host in a tight loop. This backoff applies to all requests sent by the consumer to the broker.	50	Integer
recordMetadata (producer)	Whether the producer should store the RecordMetadata results from sending to Kafka. The results are stored in a List containing the RecordMetadata metadata's. The list is stored on a header with the key <code>KafkaConstants#KAFKA_RECORDMETA</code> .	true	boolean

Name	Description	Default	Type
requestRequiredAcks (producer)	<p>The number of acknowledgments the producer requires the leader to have received before considering a request complete. This controls the durability of records that are sent. The following settings are common: <code>acks=0</code> If set to zero then the producer will not wait for any acknowledgment from the server at all. The record will be immediately added to the socket buffer and considered sent. No guarantee can be made that the server has received the record in this case, and the <code>retries</code> configuration will not take effect (as the client won't generally know of any failures). The offset given back for each record will always be set to <code>-1</code>. <code>acks=1</code> This will mean the leader will write the record to its local log but will respond without awaiting full acknowledgement from all followers. In this case should the leader fail immediately after acknowledging the record but before the followers have replicated it then the record will be lost. <code>acks=all</code> This means the leader will wait for the full set of in-sync replicas to acknowledge the record. This guarantees that the record will not be lost as long as at least one in-sync replica remains alive. This is the strongest available guarantee.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● <code>-1</code> ● <code>0</code> ● <code>1</code> ● <code>all</code> 	1	String
requestTimeoutMs (producer)	The amount of time the broker will wait trying to meet the <code>request.required.acks</code> requirement before sending back an error to the client.	30000	Integer
retries (producer)	Setting a value greater than zero will cause the client to resend any record whose <code>send</code> fails with a potentially transient error. Note that this retry is no different than if the client resent the record upon receiving the error. Allowing retries will potentially change the ordering of records because if two records are sent to a single partition, and the first fails and is retried but the second succeeds, then the second record may appear first.	0	Integer

Name	Description	Default	Type
retryBackoffMs (producer)	Before each retry, the producer refreshes the metadata of relevant topics to see if a new leader has been elected. Since leader election takes a bit of time, this property specifies the amount of time that the producer waits before refreshing the metadata.	100	Integer
sendBufferBytes (producer)	Socket write buffer size.	131072	Integer
valueSerializer (producer)	The serializer class for messages.	org.apache.kafka.common.serialization.StringSerializer	String
workerPool (producer)	To use a custom worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing. If using this option then you must handle the lifecycle of the thread pool to shut the pool down when no longer needed.		ExecutorService
workerPoolCoreSize (producer)	Number of core threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	10	Integer
workerPoolMaxSize (producer)	Maximum number of threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	20	Integer
kafkaClientFactory (advanced)	Factory to use for creating org.apache.kafka.clients.consumer.KafkaConsumer and org.apache.kafka.clients.producer.KafkaProducer instances. This allows to configure a custom factory to create instances with logic that extends the vanilla Kafka clients.		KafkaClientFactory
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean

Name	Description	Default	Type
schemaRegistryURL (confluent)	URL of the Confluent Platform schema registry servers to use. The format is host1:port1,host2:port2. This is known as schema.registry.url in the Confluent Platform documentation. This option is only available in the Confluent Platform (not standard Apache Kafka).		String
interceptorClasses (monitoring)	Sets interceptors for producer or consumers. Producer interceptors have to be classes implementing <code>org.apache.kafka.clients.producer.ProducerInterceptor</code> or Consumer interceptors have to be classes implementing <code>org.apache.kafka.clients.consumer.ConsumerInterceptor</code> Note that if you use Producer interceptor on a consumer it will throw a class cast exception in runtime.		String
kerberosBeforeRefreshMinTime (security)	Login thread sleep time between refresh attempts.	60000	Integer
kerberosInitCmd (security)	Kerberos kinit command path. Default is <code>/usr/bin/kinit</code> .	<code>/usr/bin/kinit</code>	String
kerberosPrincipalToLocalRules (security)	A list of rules for mapping from principal names to short names (typically operating system usernames). The rules are evaluated in order and the first rule that matches a principal name is used to map it to a short name. Any later rules in the list are ignored. By default, principal names of the form <code>{username}/{hostname}{REALM}</code> are mapped to <code>{username}</code> . For more details on the format please see the security authorization and acls documentation.. Multiple values can be separated by comma.	DEFAULT	String
kerberosRenewJitter (security)	Percentage of random jitter added to the renewal time.	0.05	Double
kerberosRenewWindowFactor (security)	Login thread will sleep until the specified window factor of time from last refresh to ticket's expiry has been reached, at which time it will try to renew the ticket.	0.8	Double

Name	Description	Default	Type
saslJaasConfig (security)	Expose the kafka sasl.jaas.config parameter Example: org.apache.kafka.common.security.plain.PlainLoginModule required username=USERNAME password=PASSWORD;.		String
saslKerberosServiceName (security)	The Kerberos principal name that Kafka runs as. This can be defined either in Kafka's JAAS config or in Kafka's config.		String
saslMechanism (security)	The Simple Authentication and Security Layer (SASL) Mechanism used. For the valid values see .	GSSAPI	String
securityProtocol (security)	Protocol used to communicate with brokers. SASL_PLAINTEXT, PLAINTEXT and SSL are supported.	PLAINTEXT	String
sslCipherSuites (security)	A list of cipher suites. This is a named combination of authentication, encryption, MAC and key exchange algorithm used to negotiate the security settings for a network connection using TLS or SSL network protocol. By default all the available cipher suites are supported.		String
sslContextParameters (security)	SSL configuration using a Camel SSLContextParameters object. If configured it's applied before the other SSL endpoint parameters. NOTE: Kafka only supports loading keystore from file locations, so prefix the location with file: in the KeyStoreParameters.resource option.		SSLContextParameters
sslEnabledProtocols (security)	The list of protocols enabled for SSL connections. TLSv1.2, TLSv1.1 and TLSv1 are enabled by default.		String
sslEndpointAlgorithm (security)	The endpoint identification algorithm to validate server hostname using server certificate.	https	String
sslKeymanagerAlgorithm (security)	The algorithm used by key manager factory for SSL connections. Default value is the key manager factory algorithm configured for the Java Virtual Machine.	SunX509	String
sslKeyPassword (security)	The password of the private key in the key store file. This is optional for client.		String
sslKeystoreLocation (security)	The location of the key store file. This is optional for client and can be used for two-way authentication for client.		String

Name	Description	Default	Type
sslKeystorePassword (security)	The store password for the key store file. This is optional for client and only needed if <code>ssl.keystore.location</code> is configured.		String
sslKeystoreType (security)	The file format of the key store file. This is optional for client. Default value is JKS.	JKS	String
sslProtocol (security)	The SSL protocol used to generate the <code>SSLContext</code> . Default setting is TLS, which is fine for most cases. Allowed values in recent JVMs are TLS, TLSv1.1 and TLSv1.2. SSL, SSLv2 and SSLv3 may be supported in older JVMs, but their usage is discouraged due to known security vulnerabilities.		String
sslProvider (security)	The name of the security provider used for SSL connections. Default value is the default security provider of the JVM.		String
sslTrustmanagerAlgorithm (security)	The algorithm used by trust manager factory for SSL connections. Default value is the trust manager factory algorithm configured for the Java Virtual Machine.	PKIX	String
sslTruststoreLocation (security)	The location of the trust store file.		String
sslTruststorePassword (security)	The password for the trust store file.		String
sslTruststoreType (security)	The file format of the trust store file. Default value is JKS.	JKS	String

For more information about Producer/Consumer configuration see:

- <http://kafka.apache.org/documentation.html#newconsumerconfigs>
- <http://kafka.apache.org/documentation.html#producerconfigs>

27.5. MESSAGE HEADERS

27.5.1. Consumer headers

The following headers are available when consuming messages from Kafka.

Header constant	Header value	Type	Description
KafkaConstants.TOPIC	"kafka.TOPIC"	String	The topic from where the message originated
KafkaConstants.PARTITION	"kafka.PARTITION"	Integer	The partition where the message was stored
KafkaConstants.OFFSET	"kafka.OFFSET"	Long	The offset of the message
KafkaConstants.KEY	"kafka.KEY"	Object	The key of the message if configured
KafkaConstants.HEADERS	"kafka.HEADERS"	org.apache.kafka.common.header.Headers	The record headers
KafkaConstants.LAST_RECORD_BEFORE_COMMIT	"kafka.LAST_RECORD_BEFORE_COMMIT"	Boolean	Whether or not it's the last record before commit (only available if autoCommitEnable endpoint parameter is false)
KafkaConstants.LAST_POLL_RECORD	"kafka.LAST_POLL_RECORD"	Boolean	Indicates the last record within the current poll request (only available if autoCommitEnable endpoint parameter is false or allowManualCommit is true)
KafkaConstants.MANUAL_COMMIT	"CamelKafkaManualCommit"	KafkaManualCommit	Can be used for forcing manual offset commit when using Kafka consumer.

27.5.2. Producer headers

Before sending a message to Kafka you can configure the following headers.

Header constant	Header value	Type	Description
KafkaConstants.KEY	"kafka.KEY"	Object	Required The key of the message in order to ensure that all related message goes in the same partition

Header constant	Header value	Type	Description
KafkaConstants.OVERRIDE_TOPIC	"kafka.OVERRIDE_TOPIC"	String	The topic to which send the message (override and takes precedence), and the header is not preserved.
KafkaConstants.OVERRIDE_TIMESTAMP	"kafka.OVERRIDE_TIMESTAMP"	Long	The ProducerRecord also has an associated timestamp. If the user did provide a timestamp, the producer will stamp the record with the provided timestamp and the header is not preserved.
KafkaConstants.PARTITION_KEY	"kafka.PARTITION_KEY"	Integer	Explicitly specify the partition

If you want to send a message to a dynamic topic then use **KafkaConstants.OVERRIDE_TOPIC** as its used as a one-time header that are not send along the message, as its removed in the producer.

After the message is sent to Kafka, the following headers are available

Header constant	Header value	Type	Description
KafkaConstants.KAFKA_RECORDMETA	"org.apache.kafka.clients.producer.RecordMetadata"	List<RecordMetadata>	The metadata (only configured if recordMetadata endpoint parameter is true)

27.6. CONSUMER ERROR HANDLING

While kafka consumer is polling messages from the kafka broker, then errors can happen. This section describes what happens and what you can configure.

The consumer may throw exception when invoking the Kafka **poll** API. For example if the message cannot be de-serialized due invalid data, and many other kind of errors. Those errors are in the form of **KafkaException** which are either *retryable* or not. The exceptions which can be retried (**RetriableException**) will be retried again (with a poll timeout in between). All other kind of exceptions are handled according to the *pollOnError* configuration. This configuration has the following values:

- DISCARD will discard the message and continue to poll next message.
- ERROR_HANDLER will use Camel's error handler to process the exception, and afterwards continue to poll next message.
- RECONNECT will re-connect the consumer and try poll the message again.
- RETRY will let the consumer retry polling the same message again

- STOP will stop the consumer (have to be manually started/restarted if the consumer should be able to consume messages again).

The default is **ERROR_HANDLER** which will let Camel's error handler (if any configured) process the caused exception. And then afterwards continue to poll the next message. This behavior is similar to the `bridgeErrorHandler` option that Camel components have.

For advanced control then a custom implementation of **org.apache.camel.component.kafka.PollExceptionHandler** can be configured on the component level, which allows to control which exceptions causes which of the strategies above.

27.7. SAMPLES

27.7.1. Consuming messages from Kafka

Here is the minimal route you need in order to read messages from Kafka.

```
from("kafka:test?brokers=localhost:9092")
  .log("Message received from Kafka : ${body}")
  .log("  on the topic ${headers[kafka.TOPIC]}")
  .log("  on the partition ${headers[kafka.PARTITION]}")
  .log("  with the offset ${headers[kafka.OFFSET]}")
  .log("  with the key ${headers[kafka.KEY]}")
```

If you need to consume messages from multiple topics you can use a comma separated list of topic names.

```
from("kafka:test,test1,test2?brokers=localhost:9092")
  .log("Message received from Kafka : ${body}")
  .log("  on the topic ${headers[kafka.TOPIC]}")
  .log("  on the partition ${headers[kafka.PARTITION]}")
  .log("  with the offset ${headers[kafka.OFFSET]}")
  .log("  with the key ${headers[kafka.KEY]}")
```

It's also possible to subscribe to multiple topics giving a pattern as the topic name and using the **topicsPattern** option.

```
from("kafka:test*?brokers=localhost:9092&topicsPattern=true")
  .log("Message received from Kafka : ${body}")
  .log("  on the topic ${headers[kafka.TOPIC]}")
  .log("  on the partition ${headers[kafka.PARTITION]}")
  .log("  with the offset ${headers[kafka.OFFSET]}")
  .log("  with the key ${headers[kafka.KEY]}")
```

When consuming messages from Kafka you can use your own offset management and not delegate this management to Kafka. In order to keep the offsets the component needs a **StateRepository** implementation such as **FileStateRepository**. This bean should be available in the registry. Here how to use it :

```
// Create the repository in which the Kafka offsets will be persisted
FileStateRepository repository = FileStateRepository.fileStateRepository(new
File("/path/to/repo.dat"));

// Bind this repository into the Camel registry
```



```

Registry registry = createCamelRegistry();
registry.bind("offsetRepo", repository);

// Configure the camel context
DefaultCamelContext camelContext = new DefaultCamelContext(registry);
camelContext.addRoutes(new RouteBuilder() {
    @Override
    public void configure() throws Exception {
        from("kafka:" + TOPIC + "?brokers=localhost:{{kafkaPort}}" +
            // Setup the topic and broker address
            "&groupId=A" +
            // The consumer processor group ID
            "&autoOffsetReset=earliest" +
            // Ask to start from the beginning if we have unknown offset
            "&offsetRepository=#offsetRepo")
            // Keep the offsets in the previously configured repository
            .to("mock:result");
    }
});

```

27.7.2. Producing messages to Kafka

Here is the minimal route you need in order to write messages to Kafka.

```

from("direct:start")
    .setBody(constant("Message from Camel")) // Message to send
    .setHeader(KafkaConstants.KEY, constant("Camel")) // Key of the message
    .to("kafka:test?brokers=localhost:9092");

```

27.8. SSL CONFIGURATION

You have 2 different ways to configure the SSL communication on the Kafka component.

The first way is through the many SSL endpoint parameters

```

from("kafka:" + TOPIC + "?brokers=localhost:{{kafkaPort}}" +
    "&groupId=A" +
    "&sslKeystoreLocation=/path/to/keystore.jks" +
    "&sslKeystorePassword=changeit" +
    "&sslKeyPassword=changeit" +
    "&securityProtocol=SSL")
    .to("mock:result");

```

The second way is to use the **sslContextParameters** endpoint parameter.

```

// Configure the SSLContextParameters object
KeyStoreParameters ksp = new KeyStoreParameters();
ksp.setResource("/path/to/keystore.jks");
ksp.setPassword("changeit");
KeyManagersParameters kmp = new KeyManagersParameters();
kmp.setKeyStore(ksp);
kmp.setKeyPassword("changeit");
SSLContextParameters scp = new SSLContextParameters();
scp.setKeyManagers(kmp);

```

```

// Bind this SSLContextParameters into the Camel registry
Registry registry = createCamelRegistry();
registry.bind("ssl", scp);

// Configure the camel context
DefaultCamelContext camelContext = new DefaultCamelContext(registry);
camelContext.addRoutes(new RouteBuilder() {
    @Override
    public void configure() throws Exception {
        from("kafka:" + TOPIC + "?brokers=localhost:{{kafkaPort}}" +
            // Setup the topic and broker address
            "&groupId=A" +
            // The consumer processor group ID
            "&sslContextParameters=#ssl" +
            // The security protocol
            "&securityProtocol=SSL)
            // Reference the SSL configuration
            .to("mock:result");
    }
});

```

27.9. USING THE KAFKA IDEMPOTENT REPOSITORY

The **camel-kafka** library provides a Kafka topic-based idempotent repository.

This repository stores broadcasts all changes to idempotent state (add/remove) in a Kafka topic, and populates a local in-memory cache for each repository's process instance through event sourcing. The topic used must be unique per idempotent repository instance.

The mechanism does not have any requirements about the number of topic partitions; as the repository consumes from all partitions at the same time. It also does not have any requirements about the replication factor of the topic.

Each repository instance that uses the topic (e.g. typically on different machines running in parallel) controls its own consumer group, so in a cluster of 10 Camel processes using the same topic each will control its own offset.

On startup, the instance subscribes to the topic and rewinds the offset to the beginning, rebuilding the cache to the latest state. The cache will not be considered warmed up until one poll of **pollDurationMs** in length returns 0 records. Startup will not be completed until either the cache has warmed up, or 30 seconds go by; if the latter happens the idempotent repository may be in an inconsistent state until its consumer catches up to the end of the topic.

Be mindful of the format of the header used for the uniqueness check. By default, it uses Strings as the data types. When using primitive numeric formats, the header must be deserialized accordingly. Check the samples below for examples.

A **KafkaldempotentRepository** has the following properties:

Property	Description
topic	The name of the Kafka topic to use to broadcast changes. (required)

Property	Description
bootstrapServers	The bootstrap.servers property on the internal Kafka producer and consumer. Use this as shorthand if not setting consumerConfig and producerConfig . If used, this component will apply sensible default configurations for the producer and consumer.
producerConfig	Sets the properties that will be used by the Kafka producer that broadcasts changes. Overrides bootstrapServers , so must define the Kafka bootstrap.servers property itself
consumerConfig	Sets the properties that will be used by the Kafka consumer that populates the cache from the topic. Overrides bootstrapServers , so must define the Kafka bootstrap.servers property itself
maxCacheSize	How many of the most recently used keys should be stored in memory (default 1000).
pollDurationMs	The poll duration of the Kafka consumer. The local caches are updated immediately. This value will affect how far behind other peers that update their caches from the topic are relative to the idempotent consumer instance that sent the cache action message. The default value of this is 100 ms. If setting this value explicitly, be aware that there is a tradeoff between the remote cache liveness and the volume of network traffic between this repository's consumer and the Kafka brokers. The cache warmup process also depends on there being one poll that fetches nothing - this indicates that the stream has been consumed up to the current point. If the poll duration is excessively long for the rate at which messages are sent on the topic, there exists a possibility that the cache cannot be warmed up and will operate in an inconsistent state relative to its peers until it catches up.

The repository can be instantiated by defining the **topic** and **bootstrapServers**, or the **producerConfig** and **consumerConfig** property sets can be explicitly defined to enable features such as SSL/SASL. To use, this repository must be placed in the Camel registry, either manually or by registration as a bean in Spring/Blueprint, as it is **CamelContext** aware.

Sample usage is as follows:

```
KafkaldempotentRepository kafkaldempotentRepository = new
KafkaldempotentRepository("idempotent-db-inserts", "localhost:9091");

SimpleRegistry registry = new SimpleRegistry();
registry.put("insertDbIdemRepo", kafkaldempotentRepository); // must be registered in the registry, to
enable access to the CamelContext
CamelContext context = new CamelContext(registry);

// later in RouteBuilder...
from("direct:performInsert")
    .idempotentConsumer(header("id")).messageIdRepositoryRef("insertDbIdemRepo")
        // once-only insert into database
    .end()
```

In XML:

```

<!-- simple -->
<bean id="insertDbldemRepo"
  class="org.apache.camel.processor.idempotent.kafka.KafkaIdempotentRepository">
  <property name="topic" value="idempotent-db-inserts"/>
  <property name="bootstrapServers" value="localhost:9091"/>
</bean>

<!-- complex -->
<bean id="insertDbldemRepo"
  class="org.apache.camel.processor.idempotent.kafka.KafkaIdempotentRepository">
  <property name="topic" value="idempotent-db-inserts"/>
  <property name="maxCacheSize" value="10000"/>
  <property name="consumerConfig">
    <props>
      <prop key="bootstrap.servers">localhost:9091</prop>
    </props>
  </property>
  <property name="producerConfig">
    <props>
      <prop key="bootstrap.servers">localhost:9091</prop>
    </props>
  </property>
</bean>

```

There are 3 alternatives to choose from when using idempotency with numeric identifiers. The first one is to use the static method `numericHeader` method from `org.apache.camel.component.kafka.serde.KafkaSerdeHelper` to perform the conversion for you:

```

from("direct:performInsert")
  .idempotentConsumer(numericHeader("id").messageIdRepositoryRef("insertDbldemRepo")
    // once-only insert into database
  )
  .end()

```

Alternatively, it is possible use a custom serializer configured via the route URL to perform the conversion:

```

public class CustomHeaderDeserializer extends DefaultKafkaHeaderDeserializer {
  private static final Logger LOG = LoggerFactory.getLogger(CustomHeaderDeserializer.class);

  @Override
  public Object deserialize(String key, byte[] value) {
    if (key.equals("id")) {
      BigInteger bi = new BigInteger(value);

      return String.valueOf(bi.longValue());
    } else {
      return super.deserialize(key, value);
    }
  }
}

```

Lastly, it is also possible to do so in a processor:

```

from(from).routeId("foo")
    .process(exchange -> {
        byte[] id = exchange.getIn().getHeader("id", byte[].class);

        BigInteger bi = new BigInteger(id);
        exchange.getIn().setHeader("id", String.valueOf(bi.longValue()));
    })
    .idempotentConsumer(header("id"))
    .messageIdRepositoryRef("kafkaIdempotentRepository")
    .to(to);

```

27.10. USING MANUAL COMMIT WITH KAFKA CONSUMER

By default the Kafka consumer will use auto commit, where the offset will be committed automatically in the background using a given interval.

In case you want to force manual commits, you can use **KafkaManualCommit** API from the Camel Exchange, stored on the message header. This requires to turn on manual commits by either setting the option **allowManualCommit** to **true** on the **KafkaComponent** or on the endpoint, for example:

```

KafkaComponent kafka = new KafkaComponent();
kafka.setAllowManualCommit(true);
...
camelContext.addComponent("kafka", kafka);

```

You can then use the **KafkaManualCommit** from Java code such as a Camel **Processor**:

```

public void process(Exchange exchange) {
    KafkaManualCommit manual =
        exchange.getIn().getHeader(KafkaConstants.MANUAL_COMMIT, KafkaManualCommit.class);
    manual.commit();
}

```

This will force a synchronous commit which will block until the commit is acknowledge on Kafka, or if it fails an exception is thrown. You can use an asynchronous commit as well by configuring the **KafkaManualCommitFactory** with the `DefaultKafkaManualAsyncCommitFactory` implementation.

The commit will then be done in the next consumer loop using the kafka asynchronous commit api. Be aware that records from a partition must be processed and committed by a unique thread. If not, this could lead with non consistent behaviors. This is mostly useful with aggregation's completion timeout strategies.

If you want to use a custom implementation of **KafkaManualCommit** then you can configure a custom **KafkaManualCommitFactory** on the **KafkaComponent** that creates instances of your custom implementation.

27.11. KAFKA HEADERS PROPAGATION

When consuming messages from Kafka, headers will be propagated to camel exchange headers automatically. Producing flow backed by same behaviour - camel headers of particular exchange will be propagated to kafka message headers.

Since kafka headers allows only **byte[]** values, in order camel exchange header to be propagated its value should be serialized to **bytes[]**, otherwise header will be skipped. Following header value types are

supported: **String, Integer, Long, Double, Boolean, byte[]**. Note: all headers propagated **from kafka to camel exchange** will contain **byte[]** value by default. In order to override default functionality uri parameters can be set: **headerDeserializer** for **from** route and **headerSerializer** for **to** route. Example:

```
from("kafka:my_topic?headerDeserializer=#myDeserializer")
...
.to("kafka:my_topic?headerSerializer=#mySerializer")
```

By default all headers are being filtered by **KafkaHeaderFilterStrategy**. Strategy filters out headers which start with **Camel** or **org.apache.camel** prefixes. Default strategy can be overridden by using **headerFilterStrategy** uri parameter in both **to** and **from** routes:

```
from("kafka:my_topic?headerFilterStrategy=#myStrategy")
...
.to("kafka:my_topic?headerFilterStrategy=#myStrategy")
```

myStrategy object should be subclass of **HeaderFilterStrategy** and must be placed in the Camel registry, either manually or by registration as a bean in Spring/Blueprint, as it is **CamelContext** aware.

27.12. SPRING BOOT AUTO-CONFIGURATION

When using kafka with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-kafka-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 105 options, which are listed below.

Name	Description	Default	Type
camel.component.kafka.additional-properties	Sets additional properties for either kafka consumer or kafka producer in case they can't be set directly on the camel configurations (e.g: new Kafka properties that are not reflected yet in Camel configurations), the properties have to be prefixed with <code>additionalProperties</code> . E.g: <code>additionalProperties.transactional.id=12345&additionalProperties.schema.registry.url=http://localhost:8811/avro</code> .		Map
camel.component.kafka.allow-manual-commit	Whether to allow doing manual commits via <code>KafkaManualCommit</code> . If this option is enabled then an instance of <code>KafkaManualCommit</code> is stored on the Exchange message header, which allows end users to access this API and perform manual offset commits via the Kafka consumer.	false	Boolean

Name	Description	Default	Type
<code>camel.component.kafka.auto-commit-enable</code>	If true, periodically commit to ZooKeeper the offset of messages already fetched by the consumer. This committed offset will be used when the process fails as the position from which the new consumer will begin.	true	Boolean
<code>camel.component.kafka.auto-commit-interval-ms</code>	The frequency in ms that the consumer offsets are committed to zookeeper.	5000	Integer
<code>camel.component.kafka.auto-commit-on-stop</code>	Whether to perform an explicit auto commit when the consumer stops to ensure the broker has a commit from the last consumed message. This requires the option <code>autoCommitEnable</code> is turned on. The possible values are: <code>sync</code> , <code>async</code> , or <code>none</code> . And <code>sync</code> is the default value.	sync	String
<code>camel.component.kafka.auto-offset-reset</code>	What to do when there is no initial offset in ZooKeeper or if an offset is out of range: <code>earliest</code> : automatically reset the offset to the earliest offset <code>latest</code> : automatically reset the offset to the latest offset <code>fail</code> : throw exception to the consumer.	latest	String
<code>camel.component.kafka.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.kafka.break-on-first-error</code>	This options controls what happens when a consumer is processing an exchange and it fails. If the option is false then the consumer continues to the next message and processes it. If the option is true then the consumer breaks out, and will seek back to offset of the message that caused a failure, and then re-attempt to process this message. However this can lead to endless processing of the same message if its bound to fail every time, eg a poison message. Therefore its recommended to deal with that for example by using Camel's error handler.	false	Boolean

Name	Description	Default	Type
camel.component.kafka.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.kafka.brokers	URL of the Kafka brokers to use. The format is <code>host1:port1,host2:port2</code> , and the list can be a subset of brokers or a VIP pointing to a subset of brokers. This option is known as <code>bootstrap.servers</code> in the Kafka documentation.		String
camel.component.kafka.buffer-memory-size	The total bytes of memory the producer can use to buffer records waiting to be sent to the server. If records are sent faster than they can be delivered to the server the producer will either block or throw an exception based on the preference specified by <code>block.on.buffer.full</code> . This setting should correspond roughly to the total memory the producer will use, but is not a hard bound since not all memory the producer uses is used for buffering. Some additional memory will be used for compression (if compression is enabled) as well as for maintaining in-flight requests.	33554432	Integer
camel.component.kafka.check-crcs	Automatically check the CRC32 of the records consumed. This ensures no on-the-wire or on-disk corruption to the messages occurred. This check adds some overhead, so it may be disabled in cases seeking extreme performance.	true	Boolean
camel.component.kafka.client-id	The client id is a user-specified string sent in each request to help trace calls. It should logically identify the application making the request.		String
camel.component.kafka.commit-timeout-ms	The maximum time, in milliseconds, that the code will wait for a synchronous commit to complete. The option is a <code>java.lang.Long</code> type.	5000	Long
camel.component.kafka.compression-codec	This parameter allows you to specify the compression codec for all data generated by this producer. Valid values are <code>none</code> , <code>gzip</code> and <code>snappy</code> .	none	String

Name	Description	Default	Type
camel.component.kafka.configuration	Allows to pre-configure the Kafka component with common options that the endpoints will reuse. The option is a <code>org.apache.camel.component.kafka.KafkaConfiguration</code> type.		KafkaConfiguration
camel.component.kafka.connection-max-idle-ms	Close idle connections after the number of milliseconds specified by this config.	540000	Integer
camel.component.kafka.consumer-request-timeout-ms	The configuration controls the maximum amount of time the client will wait for the response of a request. If the response is not received before the timeout elapses the client will resend the request if necessary or fail the request if retries are exhausted.	40000	Integer
camel.component.kafka.consumers-count	The number of consumers that connect to kafka server. Each consumer is run on a separate thread, that retrieves and process the incoming data.	1	Integer
camel.component.kafka.delivery-timeout-ms	An upper bound on the time to report success or failure after a call to <code>send()</code> returns. This limits the total time that a record will be delayed prior to sending, the time to await acknowledgement from the broker (if expected), and the time allowed for retrievable send failures.	120000	Integer
camel.component.kafka.enable-idempotence	If set to 'true' the producer will ensure that exactly one copy of each message is written in the stream. If 'false', producer retries may write duplicates of the retried message in the stream. If set to true this option will require <code>max.in.flight.requests.per.connection</code> to be set to 1 and retries cannot be zero and additionally acks must be set to 'all'.	false	Boolean
camel.component.kafka.enabled	Whether to enable auto configuration of the kafka component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component.kafka.fetch-max-bytes	The maximum amount of data the server should return for a fetch request. This is not an absolute maximum, if the first message in the first non-empty partition of the fetch is larger than this value, the message will still be returned to ensure that the consumer can make progress. The maximum message size accepted by the broker is defined via <code>message.max.bytes</code> (broker config) or <code>max.message.bytes</code> (topic config). Note that the consumer performs multiple fetches in parallel.	52428800	Integer
camel.component.kafka.fetch-min-bytes	The minimum amount of data the server should return for a fetch request. If insufficient data is available the request will wait for that much data to accumulate before answering the request.	1	Integer
camel.component.kafka.fetch-wait-max-ms	The maximum amount of time the server will block before answering the fetch request if there isn't sufficient data to immediately satisfy <code>fetch.min.bytes</code> .	500	Integer
camel.component.kafka.group-id	A string that uniquely identifies the group of consumer processes to which this consumer belongs. By setting the same group id multiple processes indicate that they are all part of the same consumer group. This option is required for consumers.		String
camel.component.kafka.group-instance-id	A unique identifier of the consumer instance provided by the end user. Only non-empty strings are permitted. If set, the consumer is treated as a static member, which means that only one instance with this ID is allowed in the consumer group at any time. This can be used in combination with a larger session timeout to avoid group rebalances caused by transient unavailability (e.g. process restarts). If not set, the consumer will join the group as a dynamic member, which is the traditional behavior.		String
camel.component.kafka.header-deserializer	To use a custom <code>KafkaHeaderDeserializer</code> to deserialize kafka headers values. The option is a <code>org.apache.camel.component.kafka.serde.KafkaHeaderDeserializer</code> type.		<code>KafkaHeaderDeserializer</code>
camel.component.kafka.header-filter-strategy	To use a custom <code>HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		<code>HeaderFilterStrategy</code>

Name	Description	Default	Type
<code>camel.component.kafka.header-serializer</code>	To use a custom <code>KafkaHeaderSerializer</code> to serialize kafka headers values. The option is a <code>org.apache.camel.component.kafka.serde.KafkaHeaderSerializer</code> type.		<code>KafkaHeaderSerializer</code>
<code>camel.component.kafka.heartbeat-interval-ms</code>	The expected time between heartbeats to the consumer coordinator when using Kafka's group management facilities. Heartbeats are used to ensure that the consumer's session stays active and to facilitate rebalancing when new consumers join or leave the group. The value must be set lower than <code>session.timeout.ms</code> , but typically should be set no higher than 1/3 of that value. It can be adjusted even lower to control the expected time for normal rebalances.	3000	Integer
<code>camel.component.kafka.interceptor-classes</code>	Sets interceptors for producer or consumers. Producer interceptors have to be classes implementing <code>org.apache.kafka.clients.producer.ProducerInterceptor</code> or Consumer interceptors have to be classes implementing <code>org.apache.kafka.clients.consumer.ConsumerInterceptor</code> Note that if you use Producer interceptor on a consumer it will throw a class cast exception in runtime.		String
<code>camel.component.kafka.kafka-client-factory</code>	Factory to use for creating <code>org.apache.kafka.clients.consumer.KafkaConsumer</code> and <code>org.apache.kafka.clients.producer.KafkaProducer</code> instances. This allows to configure a custom factory to create instances with logic that extends the vanilla Kafka clients. The option is a <code>org.apache.camel.component.kafka.KafkaClientFactory</code> type.		<code>KafkaClientFactory</code>
<code>camel.component.kafka.kafka-manual-commit-factory</code>	Factory to use for creating <code>KafkaManualCommit</code> instances. This allows to plugin a custom factory to create custom <code>KafkaManualCommit</code> instances in case special logic is needed when doing manual commits that deviates from the default implementation that comes out of the box. The option is a <code>org.apache.camel.component.kafka.KafkaManualCommitFactory</code> type.		<code>KafkaManualCommitFactory</code>
<code>camel.component.kafka.kerberos-before-relogin-min-time</code>	Login thread sleep time between refresh attempts.	60000	Integer

Name	Description	Default	Type
<code>camel.component.kafka.kerberos-init-cmd</code>	Kerberos kinit command path. Default is <code>/usr/bin/kinit</code> .	<code>/usr/bin/kinit</code>	String
<code>camel.component.kafka.kerberos-principal-to-local-rules</code>	A list of rules for mapping from principal names to short names (typically operating system usernames). The rules are evaluated in order and the first rule that matches a principal name is used to map it to a short name. Any later rules in the list are ignored. By default, principal names of the form <code>{username}/{hostname}{REALM}</code> are mapped to <code>{username}</code> . For more details on the format please see the security authorization and acls documentation.. Multiple values can be separated by comma.	DEFAULT	String
<code>camel.component.kafka.kerberos-renew-jitter</code>	Percentage of random jitter added to the renewal time.		Double
<code>camel.component.kafka.kerberos-renew-window-factor</code>	Login thread will sleep until the specified window factor of time from last refresh to ticket's expiry has been reached, at which time it will try to renew the ticket.		Double
<code>camel.component.kafka.key</code>	The record key (or null if no key is specified). If this option has been configured then it take precedence over header <code>KafkaConstants#KEY</code> .		String
<code>camel.component.kafka.key-deserializer</code>	Deserializer class for key that implements the Deserializer interface.	<code>org.apache.kafka.common.serialization.StringDeserializer</code>	String
<code>camel.component.kafka.key-serializer</code>	The serializer class for keys (defaults to the same as for messages if nothing is given).	<code>org.apache.kafka.common.serialization.StringSerializer</code>	String

Name	Description	Default	Type
camel.component.kafka.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.kafka.linger-ms	The producer groups together any records that arrive in between request transmissions into a single batched request. Normally this occurs only under load when records arrive faster than they can be sent out. However in some circumstances the client may want to reduce the number of requests even under moderate load. This setting accomplishes this by adding a small amount of artificial delay that is, rather than immediately sending out a record the producer will wait for up to the given delay to allow other records to be sent so that the sends can be batched together. This can be thought of as analogous to Nagle's algorithm in TCP. This setting gives the upper bound on the delay for batching: once we get batch.size worth of records for a partition it will be sent immediately regardless of this setting, however if we have fewer than this many bytes accumulated for this partition we will 'linger' for the specified time waiting for more records to show up. This setting defaults to 0 (i.e. no delay). Setting linger.ms=5, for example, would have the effect of reducing the number of requests sent but would add up to 5ms of latency to records sent in the absence of load.	0	Integer
camel.component.kafka.max-block-ms	The configuration controls how long sending to kafka will block. These methods can be blocked for multiple reasons. For e.g: buffer full, metadata unavailable. This configuration imposes maximum limit on the total time spent in fetching metadata, serialization of key and value, partitioning and allocation of buffer memory when doing a send(). In case of partitionsFor(), this configuration imposes a maximum time threshold on waiting for metadata.	60000	Integer

Name	Description	Default	Type
camel.component.kafka.max-in-flight-request	The maximum number of unacknowledged requests the client will send on a single connection before blocking. Note that if this setting is set to be greater than 1 and there are failed sends, there is a risk of message re-ordering due to retries (i.e., if retries are enabled).	5	Integer
camel.component.kafka.max-partition-fetch-bytes	The maximum amount of data per-partition the server will return. The maximum total memory used for a request will be #partitions max.partition.fetch.bytes. This size must be at least as large as the maximum message size the server allows or else it is possible for the producer to send messages larger than the consumer can fetch. If that happens, the consumer can get stuck trying to fetch a large message on a certain partition.	1048576	Integer
camel.component.kafka.max-poll-interval-ms	The maximum delay between invocations of poll() when using consumer group management. This places an upper bound on the amount of time that the consumer can be idle before fetching more records. If poll() is not called before expiration of this timeout, then the consumer is considered failed and the group will rebalance in order to reassign the partitions to another member. The option is a java.lang.Long type.		Long
camel.component.kafka.max-poll-records	The maximum number of records returned in a single call to poll().	500	Integer
camel.component.kafka.max-request-size	The maximum size of a request. This is also effectively a cap on the maximum record size. Note that the server has its own cap on record size which may be different from this. This setting will limit the number of record batches the producer will send in a single request to avoid sending huge requests.	1048576	Integer
camel.component.kafka.metadata-max-age-ms	The period of time in milliseconds after which we force a refresh of metadata even if we haven't seen any partition leadership changes to proactively discover any new brokers or partitions.	300000	Integer
camel.component.kafka.metric-reporters	A list of classes to use as metrics reporters. Implementing the MetricReporter interface allows plugging in classes that will be notified of new metric creation. The JmxReporter is always included to register JMX statistics.		String

Name	Description	Default	Type
<code>camel.component.kafka.metrics-sample-window-ms</code>	The number of samples maintained to compute metrics.	30000	Integer
<code>camel.component.kafka.no-of-metrics-sample</code>	The number of samples maintained to compute metrics.	2	Integer
<code>camel.component.kafka.offset-repository</code>	The offset repository to use in order to locally store the offset of each partition of the topic. Defining one will disable the autocommit. The option is a <code>org.apache.camel.spi.StateRepository<java.lang.String, java.lang.String></code> type.		StateRepository
<code>camel.component.kafka.partition-assignor</code>	The class name of the partition assignment strategy that the client will use to distribute partition ownership amongst consumer instances when group management is used.	<code>org.apache.kafka.clients.consumer.RangeAssignor</code>	String
<code>camel.component.kafka.partition-key</code>	The partition to which the record will be sent (or null if no partition was specified). If this option has been configured then it take precedence over header <code>KafkaConstants#PARTITION_KEY</code> .		Integer
<code>camel.component.kafka.partitioner</code>	The partitioner class for partitioning messages amongst sub-topics. The default partitioner is based on the hash of the key.	<code>org.apache.kafka.clients.producer.internals.DefaultPartitioner</code>	String
<code>camel.component.kafka.poll-exception-strategy</code>	To use a custom strategy with the consumer to control how to handle exceptions thrown from the Kafka broker while pooling messages. The option is a <code>org.apache.camel.component.kafka.PollExceptionStrategy</code> type.		PollExceptionStrategy

Name	Description	Default	Type
camel.component.kafka.poll-on-error	What to do if kafka threw an exception while polling for new messages. Will by default use the value from the component configuration unless an explicit value has been configured on the endpoint level. DISCARD will discard the message and continue to poll next message. ERROR_HANDLER will use Camel's error handler to process the exception, and afterwards continue to poll next message. RECONNECT will reconnect the consumer and try poll the message again. RETRY will let the consumer retry polling the same message again. STOP will stop the consumer (have to be manually started/restarted if the consumer should be able to consume messages again).		PollOnError
camel.component.kafka.poll-timeout-ms	The timeout used when polling the KafkaConsumer. The option is a java.lang.Long type.	5000	Long
camel.component.kafka.producer-batch-size	The producer will attempt to batch records together into fewer requests whenever multiple records are being sent to the same partition. This helps performance on both the client and the server. This configuration controls the default batch size in bytes. No attempt will be made to batch records larger than this size. Requests sent to brokers will contain multiple batches, one for each partition with data available to be sent. A small batch size will make batching less common and may reduce throughput (a batch size of zero will disable batching entirely). A very large batch size may use memory a bit more wastefully as we will always allocate a buffer of the specified batch size in anticipation of additional records.	16384	Integer
camel.component.kafka.queue-buffering-max-messages	The maximum number of unsent messages that can be queued up the producer when using async mode before either the producer must be blocked or data must be dropped.	10000	Integer
camel.component.kafka.receive-buffer-bytes	The size of the TCP receive buffer (SO_RCVBUF) to use when reading data.	65536	Integer

Name	Description	Default	Type
<code>camel.component.kafka.reconnect-backoff-max-ms</code>	The maximum amount of time in milliseconds to wait when reconnecting to a broker that has repeatedly failed to connect. If provided, the backoff per host will increase exponentially for each consecutive connection failure, up to this maximum. After calculating the backoff increase, 20% random jitter is added to avoid connection storms.	1000	Integer
<code>camel.component.kafka.reconnect-backoff-ms</code>	The amount of time to wait before attempting to reconnect to a given host. This avoids repeatedly connecting to a host in a tight loop. This backoff applies to all requests sent by the consumer to the broker.	50	Integer
<code>camel.component.kafka.record-metadata</code>	Whether the producer should store the RecordMetadata results from sending to Kafka. The results are stored in a List containing the RecordMetadata metadata's. The list is stored on a header with the key <code>KafkaConstants#KAFKA_RECORDMETA</code> .	true	Boolean
<code>camel.component.kafka.request-required-acks</code>	The number of acknowledgments the producer requires the leader to have received before considering a request complete. This controls the durability of records that are sent. The following settings are common: <code>acks=0</code> If set to zero then the producer will not wait for any acknowledgment from the server at all. The record will be immediately added to the socket buffer and considered sent. No guarantee can be made that the server has received the record in this case, and the retries configuration will not take effect (as the client won't generally know of any failures). The offset given back for each record will always be set to -1. <code>acks=1</code> This will mean the leader will write the record to its local log but will respond without awaiting full acknowledgement from all followers. In this case should the leader fail immediately after acknowledging the record but before the followers have replicated it then the record will be lost. <code>acks=all</code> This means the leader will wait for the full set of in-sync replicas to acknowledge the record. This guarantees that the record will not be lost as long as at least one in-sync replica remains alive. This is the strongest available guarantee.	1	String
<code>camel.component.kafka.request-timeout-ms</code>	The amount of time the broker will wait trying to meet the <code>request.required.acks</code> requirement before sending back an error to the client.	30000	Integer

Name	Description	Default	Type
camel.component.kafka.resume-strategy	This option allows the user to set a custom resume strategy. The resume strategy is executed when partitions are assigned (i.e.: when connecting or reconnecting). It allows implementations to customize how to resume operations and serve as more flexible alternative to the seekTo and the offsetRepository mechanisms. See the KafkaConsumerResumeStrategy for implementation details. This option does not affect the auto commit setting. It is likely that implementations using this setting will also want to evaluate using the manual commit option along with this. The option is a org.apache.camel.component.kafka.consumer.support.KafkaConsumerResumeStrategy type.		KafkaConsumerResumeStrategy
camel.component.kafka.retries	Setting a value greater than zero will cause the client to resend any record whose send fails with a potentially transient error. Note that this retry is no different than if the client resent the record upon receiving the error. Allowing retries will potentially change the ordering of records because if two records are sent to a single partition, and the first fails and is retried but the second succeeds, then the second record may appear first.	0	Integer
camel.component.kafka.retry-backoff-ms	Before each retry, the producer refreshes the metadata of relevant topics to see if a new leader has been elected. Since leader election takes a bit of time, this property specifies the amount of time that the producer waits before refreshing the metadata.	100	Integer
camel.component.kafka.sasl-jaas-config	Expose the kafka sasl.jaas.config parameter Example: org.apache.kafka.common.security.plain.PlainLoginModule required username=USERNAME password=PASSWORD;.		String
camel.component.kafka.sasl-kerberos-service-name	The Kerberos principal name that Kafka runs as. This can be defined either in Kafka's JAAS config or in Kafka's config.		String
camel.component.kafka.sasl-mechanism	The Simple Authentication and Security Layer (SASL) Mechanism used. For the valid values see .	GSSAPI	String

Name	Description	Default	Type
camel.component.kafka.schema-registry-u-r-l	URL of the Confluent Platform schema registry servers to use. The format is host1:port1,host2:port2. This is known as schema.registry.url in the Confluent Platform documentation. This option is only available in the Confluent Platform (not standard Apache Kafka).		String
camel.component.kafka.security-protocol	Protocol used to communicate with brokers. SASL_PLAINTEXT, PLAINTEXT and SSL are supported.	PLAINTEXT	String
camel.component.kafka.seek-to	Set if KafkaConsumer will read from beginning or end on startup: beginning : read from beginning end : read from end This is replacing the earlier property seekToBeginning.		String
camel.component.kafka.send-buffer-bytes	Socket write buffer size.	131072	Integer
camel.component.kafka.session-timeout-ms	The timeout used to detect failures when using Kafka's group management facilities.	10000	Integer
camel.component.kafka.shutdown-timeout	Timeout in milliseconds to wait gracefully for the consumer or producer to shutdown and terminate its worker threads.	30000	Integer
camel.component.kafka.specific-avro-reader	This enables the use of a specific Avro reader for use with the Confluent Platform schema registry and the io.confluent.kafka.serializers.KafkaAvroDeserializer. This option is only available in the Confluent Platform (not standard Apache Kafka).	false	Boolean
camel.component.kafka.ssl-cipher-suites	A list of cipher suites. This is a named combination of authentication, encryption, MAC and key exchange algorithm used to negotiate the security settings for a network connection using TLS or SSL network protocol. By default all the available cipher suites are supported.		String

Name	Description	Default	Type
<code>camel.component.kafka.ssl-context-parameters</code>	SSL configuration using a Camel <code>SSLContextParameters</code> object. If configured it's applied before the other SSL endpoint parameters. NOTE: Kafka only supports loading keystore from file locations, so prefix the location with <code>file:</code> in the <code>KeyStoreParameters.resource</code> option. The option is a <code>org.apache.camel.support.jsse.SSLContextParameters</code> type.		<code>SSLContextParameters</code>
<code>camel.component.kafka.ssl-enabled-protocols</code>	The list of protocols enabled for SSL connections. TLSv1.2, TLSv1.1 and TLSv1 are enabled by default.		String
<code>camel.component.kafka.ssl-endpoint-algorithm</code>	The endpoint identification algorithm to validate server hostname using server certificate.	https	String
<code>camel.component.kafka.ssl-key-password</code>	The password of the private key in the key store file. This is optional for client.		String
<code>camel.component.kafka.ssl-keymanager-algorithm</code>	The algorithm used by key manager factory for SSL connections. Default value is the key manager factory algorithm configured for the Java Virtual Machine.	SunX509	String
<code>camel.component.kafka.ssl-keystore-location</code>	The location of the key store file. This is optional for client and can be used for two-way authentication for client.		String
<code>camel.component.kafka.ssl-keystore-password</code>	The store password for the key store file. This is optional for client and only needed if <code>ssl.keystore.location</code> is configured.		String
<code>camel.component.kafka.ssl-keystore-type</code>	The file format of the key store file. This is optional for client. Default value is JKS.	JKS	String
<code>camel.component.kafka.ssl-protocol</code>	The SSL protocol used to generate the <code>SSLContext</code> . Default setting is TLS, which is fine for most cases. Allowed values in recent JVMs are TLS, TLSv1.1 and TLSv1.2. SSL, SSLv2 and SSLv3 may be supported in older JVMs, but their usage is discouraged due to known security vulnerabilities.		String

Name	Description	Default	Type
<code>camel.component.kafka.ssl-provider</code>	The name of the security provider used for SSL connections. Default value is the default security provider of the JVM.		String
<code>camel.component.kafka.ssl-trustmanager-algorithm</code>	The algorithm used by trust manager factory for SSL connections. Default value is the trust manager factory algorithm configured for the Java Virtual Machine.	PKIX	String
<code>camel.component.kafka.ssl-truststore-location</code>	The location of the trust store file.		String
<code>camel.component.kafka.ssl-truststore-password</code>	The password for the trust store file.		String
<code>camel.component.kafka.ssl-truststore-type</code>	The file format of the trust store file. Default value is JKS.	JKS	String
<code>camel.component.kafka.synchronous</code>	Sets whether synchronous processing should be strictly used.	false	Boolean
<code>camel.component.kafka.topic-is-pattern</code>	Whether the topic is a pattern (regular expression). This can be used to subscribe to dynamic number of topics matching the pattern.	false	Boolean
<code>camel.component.kafka.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	false	Boolean
<code>camel.component.kafka.value-deserializer</code>	Deserializer class for value that implements the Deserializer interface.	<code>org.apache.kafka.common.serialization.StringDeserializer</code>	String

Name	Description	Default	Type
camel.component.kafka.value-serializer	The serializer class for messages.	org.apache.kafka.common.serialization.StringSerializer	String
camel.component.kafka.worker-pool	To use a custom worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing. If using this option then you must handle the lifecycle of the thread pool to shut the pool down when no longer needed. The option is a <code>java.util.concurrent.ExecutorService</code> type.		ExecutorService
camel.component.kafka.worker-pool-core-size	Number of core threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	10	Integer
camel.component.kafka.worker-pool-max-size	Maximum number of threads for the worker pool for continue routing Exchange after kafka server has acknowledge the message that was sent to it from KafkaProducer using asynchronous non-blocking processing.	20	Integer

CHAPTER 28. KAMELET

Both producer and consumer are supported

The Kamelet Component provides support for interacting with the [Camel Route Template](#) engine using Endpoint semantic.

28.1. URI FORMAT

```
kamelet:templateId/routeId[?options]
```

28.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

28.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

28.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

28.3. COMPONENT OPTIONS

The Kamelet component supports 9 options, which are listed below.

Name	Description	Default	Type
location (common)	The location(s) of the Kamelets on the file system. Multiple locations can be set separated by comma.	classpath:/kamelets	String
routeProperties (common)	Set route local parameters.		Map
templateProperties (common)	Set template local parameters.		Map
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
block (producer)	If sending a message to a kamelet endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
timeout (producer)	The timeout value to use if block is enabled.	30000	long
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
routeTemplateLoaderListener (advanced)	Autowired To plugin a custom listener for when the Kamelet component is loading Kamelets from external resources.		RouteTemplateLoaderListener

28.4. ENDPOINT OPTIONS

The Kamelet endpoint is configured using URI syntax:

```
kamelet:templateId/routeId
```

with the following path and query parameters:

28.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
templateId (common)	Required The Route Template ID.		String
routeId (common)	The Route ID. Default value notice: The ID will be auto-generated if not provided.		String

28.4.2. Query Parameters (8 parameters)

Name	Description	Default	Type
location (common)	Location of the Kamelet to use which can be specified as a resource from file system, classpath etc. The location cannot use wildcards, and must refer to a file including extension, for example file:/etc/foo-kamelet.xml.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 		ExchangePattern
block (producer)	If sending a message to a direct endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	boolean
failIfNoConsumers (producer)	Whether the producer should fail by throwing an exception, when sending to a kamelet endpoint with no active consumers.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
timeout (producer)	The timeout value to use if block is enabled.	30000	long



NOTE

The **kamelet** endpoint is **lenient**, which means that the endpoint accepts additional parameters that are passed to the engine and consumed upon route materialization.

28.5. DISCOVERY

If a [Route Template](#) is not found, the **kamelet** endpoint tries to load the related **kamelet** definition from the file system (by default **classpath:/kamelets**). The default resolution mechanism expects kamelet files to have the extension **.kamelet.yaml**.

28.6. SAMPLES

Kamelets can be used as if they were standard Camel components. For example, suppose that we have created a Route Template as follows:

```
routeTemplate("setMyBody")
  .templateParameter("bodyValue")
  .from("kamelet:source")
  .setBody().constant("{{bodyValue}}");
```



NOTE

To let the **Kamelet** component wiring the materialized route to the caller processor, we need to be able to identify the input and output endpoint of the route and this is done by using **kamelet:source** to mark the input endpoint and **kamelet:sink** for the output endpoint.

Then the template can be instantiated and invoked as shown below:

```
from("direct:setMyBody")
  .to("kamelet:setMyBody?bodyValue=myKamelet");
```

Behind the scenes, the **Kamelet** component does the following things:

1. It instantiates a route out of the Route Template identified by the given **templateId** path parameter (in this case **setBody**)
2. It will act like the **direct** component and connect the current route to the materialized one.

If you had to do it programmatically, it would have been something like:

```
routeTemplate("setMyBody")
  .templateParameter("bodyValue")
  .from("direct:{{foo}}")
  .setBody().constant("{{bodyValue}}");

TemplatedRouteBuilder.builder(context, "setMyBody")
  .parameter("foo", "bar")
  .parameter("bodyValue", "myKamelet")
  .add();

from("direct:template")
  .to("direct:bar");
```

28.7. SPRING BOOT AUTO-CONFIGURATION

When using kamelet with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

■

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-kamelet-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 10 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.kamelet.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.kamelet.block</code>	If sending a message to a kamelet endpoint which has no active consumer, then we can tell the producer to block and wait for the consumer to become active.	true	Boolean
<code>camel.component.kamelet.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.kamelet.enabled</code>	Whether to enable auto configuration of the kamelet component. This is enabled by default.		Boolean
<code>camel.component.kamelet.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.kamelet.location</code>	The location(s) of the Kamelets on the file system. Multiple locations can be set separated by comma.	classpath:/kamelets	String

Name	Description	Default	Type
<code>camel.component.kamelet.route-properties</code>	Set route local parameters.		Map
<code>camel.component.kamelet.route-template-loader-listener</code>	To plugin a custom listener for when the Kamelet component is loading Kamelets from external resources. The option is a <code>org.apache.camel.spi.RouteTemplateLoaderListener</code> type.		<code>RouteTemplateLoaderListener</code>
<code>camel.component.kamelet.template-properties</code>	Set template local parameters.		Map
<code>camel.component.kamelet.timeout</code>	The timeout value to use if block is enabled.	30000	Long

CHAPTER 29. LANGUAGE

Only producer is supported

The Language component allows you to send Exchange to an endpoint which executes a script by any of the supported Languages in Camel. By having a component to execute language scripts, it allows more dynamic routing capabilities. For example by using the Routing Slip or [Dynamic Router](#) EIPs you can send messages to language endpoints where the script is dynamic defined as well.

This component is provided out of the box in camel-core and hence no additional JARs is needed. You only have to include additional Camel components if the language of choice mandates it, such as using [Groovy](#) or JavaScript languages.

29.1. URI FORMAT

```
language://languageName[:script][?options]
```

You can refer to an external resource for the script using same notation as supported by the other [Languages](#) in Camel.

```
language://languageName:resource:scheme:location][?options]
```

29.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

29.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

29.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

29.3. COMPONENT OPTIONS

The Language component supports 2 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

29.4. ENDPOINT OPTIONS

The Language endpoint is configured using URI syntax:

```
language:languageName:resourceUri
```

with the following path and query parameters:

29.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
languageName (producer)	<p>Required Sets the name of the language to use.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • bean • constant • exchangeProperty • file • groovy • header • javascript • jsonpath • mvel • ognl • ref • simple • spel • sql • terser • tokenize • xpath • xquery • xtokenize 		String
resourceUri (producer)	Path to the resource, or a reference to lookup a bean in the Registry to use as the resource.		String

29.4.2. Query Parameters (7 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
allowContextMapAll (producer)	Sets whether the context map should allow access to all details. By default only the message body and headers can be accessed. This option can be enabled for full access to the current Exchange and CamelContext. Doing so impose a potential security risk as this opens access to the full power of CamelContext API.	false	boolean
binary (producer)	Whether the script is binary content or text content. By default the script is read as text content (eg java.lang.String).	false	boolean
cacheScript (producer)	Whether to cache the compiled script and reuse. Notice reusing the script can cause side effects from processing one Camel org.apache.camel.Exchange to the next org.apache.camel.Exchange.	false	boolean
contentCache (producer)	Sets whether to use resource content cache or not.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
script (producer)	Sets the script to execute.		String
transform (producer)	Whether or not the result of the script should be used as message body. This options is default true.	true	boolean

29.5. MESSAGE HEADERS

The following message headers can be used to affect the behavior of the component

Header	Description
CamelLanguageScript	The script to execute provided in the header. Takes precedence over script configured on the endpoint.

29.6. EXAMPLES

For example you can use the [Simple](#) language to Message Translator a message.

You can also provide the script as a header as shown below. Here we use XPath language to extract the text from the <foo> tag.

```
Object out = producer.requestBodyAndHeader("language:xpath", "<foo>Hello World</foo>",
Exchange.LANGUAGE_SCRIPT, "/foo/text()");
assertEquals("Hello World", out);
```

29.7. LOADING SCRIPTS FROM RESOURCES

You can specify a resource uri for a script to load in either the endpoint uri, or in the **Exchange.LANGUAGE_SCRIPT** header. The uri must start with one of the following schemes: file:, classpath:, or http:

By default the script is loaded once and cached. However you can disable the **contentCache** option and have the script loaded on each evaluation. For example if the file myscript.txt is changed on disk, then the updated script is used:

You can refer to the resource similar to the other [Languages](#) in Camel by prefixing with "resource:" as shown below.

29.8. SPRING BOOT AUTO-CONFIGURATION

When using language with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-language-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.language.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.language.enabled</code>	Whether to enable auto configuration of the language component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component.language.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 30. LOG

Only producer is supported

The Log component logs message exchanges to the underlying logging mechanism.

Camel uses [SLF4J](#) which allows you to configure logging via, among others:

- Log4j
- Logback
- Java Util Logging

30.1. URI FORMAT

```
log:loggingCategory[?options]
```

Where **loggingCategory** is the name of the logging category to use. You can append query options to the URI in the following format,

?option=value&option=value&...



NOTE

Using Logger instance from the Registry

If there's single instance of **org.slf4j.Logger** found in the Registry, the **loggingCategory** is no longer used to create logger instance. The registered instance is used instead. Also it is possible to reference particular **Logger** instance using **?logger=#myLogger** URI parameter. Eventually, if there's no registered and URI **logger** parameter, the logger instance is created using **loggingCategory**.

For example, a log endpoint typically specifies the logging level using the **level** option, as follows:

```
log:org.apache.camel.example?level=DEBUG
```

The default logger logs every exchange (*regular logging*). But Camel also ships with the **Throughput** logger, which is used whenever the **groupSize** option is specified.



NOTE

Also a log in the DSL

There is also a **log** directly in the DSL, but it has a different purpose. Its meant for lightweight and human logs. See more details at [LogEIP](#).

30.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

30.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

30.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

30.3. COMPONENT OPTIONS

The Log component supports 3 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
exchangeFormatter (advanced)	Autowired Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to DefaultExchangeFormatter.		ExchangeFormatter

30.4. ENDPOINT OPTIONS

The Log endpoint is configured using URI syntax:

```
log:loggerName
```

with the following path and query parameters:

30.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
loggerName (producer)	Required Name of the logging category to use.		String

30.4.2. Query Parameters (27 parameters)

Name	Description	Default	Type
groupActiveOnly (producer)	If true, will hide stats when no new messages have been received for a time interval, if false, show stats regardless of message traffic.	true	Boolean
groupDelay (producer)	Set the initial delay for stats (in millis).		Long
groupInterval (producer)	If specified will group message stats by this time interval (in millis).		Long
groupSize (producer)	An integer that specifies a group size for throughput logging.		Integer

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
level (producer)	Logging level to use. The default value is INFO. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	INFO	String
logMask (producer)	If true, mask sensitive information like password or passphrase in the log.		Boolean
marker (producer)	An optional Marker name to use.		String
exchangeFormatter (advanced)	To use a custom exchange formatter.		ExchangeFormatter
maxChars (formatting)	Limits the number of characters logged per line.	10000	int
multiline (formatting)	If enabled then each information is outputted on a newline.	false	boolean
showAll (formatting)	Quick option for turning all options on. (multiline, maxChars has to be manually set if to be used).	false	boolean
showAllProperties (formatting)	Show all of the exchange properties (both internal and custom).	false	boolean

Name	Description	Default	Type
showBody (formatting)	Show the message body.	true	boolean
showBodyType (formatting)	Show the body Java type.	true	boolean
showCaughtException (formatting)	If the exchange has a caught exception, show the exception message (no stack trace). A caught exception is stored as a property on the exchange (using the key <code>org.apache.camel.Exchange#EXCEPTION_CAUGHT</code>) and for instance a <code>doCatch</code> can catch exceptions.	false	boolean
showException (formatting)	If the exchange has an exception, show the exception message (no stacktrace).	false	boolean
showExchangeId (formatting)	Show the unique exchange ID.	false	boolean
showExchangePattern (formatting)	Shows the Message Exchange Pattern (or MEP for short).	true	boolean
showFiles (formatting)	If enabled Camel will output files.	false	boolean
showFuture (formatting)	If enabled Camel will on Future objects wait for it to complete to obtain the payload to be logged.	false	boolean
showHeaders (formatting)	Show the message headers.	false	boolean
showProperties (formatting)	Show the exchange properties (only custom). Use <code>showAllProperties</code> to show both internal and custom properties.	false	boolean
showStackTrace (formatting)	Show the stack trace, if an exchange has an exception. Only effective if one of <code>showAll</code> , <code>showException</code> or <code>showCaughtException</code> are enabled.	false	boolean
showStreams (formatting)	Whether Camel should show stream bodies or not (eg such as <code>java.io.InputStream</code>). Beware if you enable this option then you may not be able later to access the message body as the stream have already been read by this logger. To remedy this you will have to use Stream Caching.	false	boolean

Name	Description	Default	Type
skipBodyLineSeparator (formatting)	Whether to skip line separators when logging the message body. This allows to log the message body in one line, setting this option to false will preserve any line separators from the body, which then will log the body as is.	true	boolean
style (formatting)	Sets the outputs style to use. Enum values: <ul style="list-style-type: none"> • Default • Tab • Fixed 	Default	OutputStyle

30.5. REGULAR LOGGER SAMPLE

In the route below we log the incoming orders at **DEBUG** level before the order is processed:

```
from("activemq:orders").to("log:com.mycompany.order?level=DEBUG").to("bean:processOrder");
```

Or using Spring XML to define the route:

```
<route>
  <from uri="activemq:orders"/>
  <to uri="log:com.mycompany.order?level=DEBUG"/>
  <to uri="bean:processOrder"/>
</route>
```

30.6. REGULAR LOGGER WITH FORMATTER SAMPLE

In the route below we log the incoming orders at **INFO** level before the order is processed.

```
from("activemq:orders").
  to("log:com.mycompany.order?showAll=true&multiline=true").to("bean:processOrder");
```

30.7. THROUGHPUT LOGGER WITH GROUPSIZE SAMPLE

In the route below we log the throughput of the incoming orders at **DEBUG** level grouped by 10 messages.

```
from("activemq:orders").
  to("log:com.mycompany.order?level=DEBUG&groupSize=10").to("bean:processOrder");
```

30.8. THROUGHPUT LOGGER WITH GROUPINTERVAL SAMPLE

This route will result in message stats logged every 10s, with an initial 60s delay and stats should be displayed even if there isn't any message traffic.

```
from("activemq:orders").
  to("log:com.mycompany.order?
  level=DEBUG&groupInterval=10000&groupDelay=60000&groupActiveOnly=false").to("bean:process
  Order");
```

The following will be logged:

```
"Received: 1000 new messages, with total 2000 so far. Last group took: 10000 millis which is: 100
  messages per second. average: 100"
```

30.9. MASKING SENSITIVE INFORMATION LIKE PASSWORD

You can enable security masking for logging by setting **logMask** flag to **true**. Note that this option also affects Log EIP.

To enable mask in Java DSL at CamelContext level:

```
camelContext.setLogMask(true);
```

And in XML:

```
<camelContext logMask="true">
```

You can also turn it on/off at endpoint level. To enable mask in Java DSL at endpoint level, add `logMask=true` option in the URI for the log endpoint:

```
from("direct:start").to("log:foo?logMask=true");
```

And in XML:

```
<route>
  <from uri="direct:foo"/>
  <to uri="log:foo?logMask=true"/>
</route>
```

org.apache.camel.support.processor.DefaultMaskingFormatter is used for the masking by default. If you want to use a custom masking formatter, put it into registry with the name **CamelCustomLogMask**. Note that the masking formatter must implement **org.apache.camel.spi.MaskingFormatter**.

30.10. FULL CUSTOMIZATION OF THE LOGGING OUTPUT

With the options outlined in the section, you can control much of the output of the logger. However, log lines will always follow this structure:

```
Exchange[Id:ID-machine-local-50656-1234567901234-1-2, ExchangePattern:InOut,
  Properties:{CamelToEndpoint=log://org.apache.camel.component.log.TEST?showAll=true,
  CamelCreatedTimestamp=Thu Mar 28 00:00:00 WET 2013},
  Headers:{breadcrumbId=ID-machine-local-50656-1234567901234-1-1}, BodyType:String, Body:Hello
  World, Out: null]
```

This format is unsuitable in some cases, perhaps because you need to...

- Filter the headers and properties that are printed, to strike a balance between insight and verbosity.
- Adjust the log message to whatever you deem most readable.
- Tailor log messages for digestion by log mining systems, e.g. Splunk.
- Print specific body types differently.

Whenever you require absolute customization, you can create a class that implements the interface. Within the **format(Exchange)** method you have access to the full Exchange, so you can select and extract the precise information you need, format it in a custom manner and return it. The return value will become the final log message.

You can have the Log component pick up your custom **ExchangeFormatter** in either of two ways:

Explicitly instantiating the LogComponent in your Registry:

```
<bean name="log" class="org.apache.camel.component.log.LogComponent">
  <property name="exchangeFormatter" ref="myCustomFormatter" />
</bean>
```

30.10.1. Convention over configuration

Simply by registering a bean with the name **logFormatter**; the Log Component is intelligent enough to pick it up automatically.

```
<bean name="logFormatter" class="com.xyz.MyCustomExchangeFormatter" />
```



NOTE

The **ExchangeFormatter** gets applied to **all Log endpoints within that Camel Context**. If you need different ExchangeFormatters for different endpoints, just instantiate the LogComponent as many times as needed, and use the relevant bean name as the endpoint prefix.

When using a custom log formatter, you can specify parameters in the log uri, which gets configured on the custom log formatter. Though when you do that you should define the "logFormatter" as prototype scoped so its not shared if you have different parameters, for example,

```
<bean name="logFormatter" class="com.xyz.MyCustomExchangeFormatter" scope="prototype"/>
```

And then we can have Camel routes using the log uri with different options:

```
<to uri="log:foo?param1=foo&param2=100"/>
<to uri="log:bar?param1=bar&param2=200"/>
```

30.11. SPRING BOOT AUTO-CONFIGURATION

When using log with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-log-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
camel.component.log.automated-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.log.enabled	Whether to enable auto configuration of the log component. This is enabled by default.		Boolean
camel.component.log.exchange-formatter	Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to DefaultExchangeFormatter. The option is a <code>org.apache.camel.spi.ExchangeFormatter</code> type.		ExchangeFormatter
camel.component.log.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean


CHAPTER 31. MAIL

Both producer and consumer are supported

The Mail component provides access to Email via Spring's Mail support and the underlying JavaMail system.

Maven users will need to add the following dependency to their **pom.xml** for this component:


```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-mail</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```



NOTE

POP3 or IMAP

POP3 has some limitations and end users are encouraged to use IMAP if possible.



NOTE

Using mock-mail for testing

You can use a mock framework for unit testing, which allows you to test without the need for a real mail server. However you should remember to not include the mock-mail when you go into production or other environments where you need to send mails to a real mail server. Just the presence of the mock-javamail.jar on the classpath means that it will kick in and avoid sending the mails.

31.1. URI FORMAT

Mail endpoints can have one of the following URI formats (for the protocols, SMTP, POP3, or IMAP, respectively):

```
smtp://[username@]host[:port][?options]
pop3://[username@]host[:port][?options]
imap://[username@]host[:port][?options]
```

The mail component also supports secure variants of these protocols (layered over SSL). You can enable the secure protocols by adding **s** to the scheme:

```
smtps://[username@]host[:port][?options]
pop3s://[username@]host[:port][?options]
imaps://[username@]host[:port][?options]
```

31.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level

- endpoint level

31.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

31.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

31.3. COMPONENT OPTIONS

The Mail component supports 43 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
closeFolder (consumer)	Whether the consumer should close the folder after polling. Setting this option to false and having disconnect=false as well, then the consumer keep the folder open between polls.	true	boolean

Name	Description	Default	Type
copyTo (consumer)	After processing a mail message, it can be copied to a mail folder with the given name. You can override this configuration value, with a header with the key <code>copyTo</code> , allowing you to copy messages to folder names configured at runtime.		String
decodeFilename (consumer)	If set to true, the <code>MimeUtility.decodeText</code> method will be used to decode the filename. This is similar to setting JVM system property <code>mail.mime.encodefilename</code> .	false	boolean
delete (consumer)	Deletes the messages after they have been processed. This is done by setting the DELETED flag on the mail message. If false, the SEEN flag is set instead. As of Camel 2.10 you can override this configuration option by setting a header with the key <code>delete</code> to determine if the mail should be deleted or not.	false	boolean
disconnect (consumer)	Whether the consumer should disconnect after polling. If enabled this forces Camel to connect on each poll.	false	boolean
handleFailedMessage (consumer)	If the mail consumer cannot retrieve a given mail message, then this option allows to handle the caused exception by the consumer's error handler. By enable the bridge error handler on the consumer, then the Camel routing error handler can handle the exception instead. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	boolean
mimeDecodeHeaders (consumer)	This option enables transparent MIME decoding and unfolding for mail headers.	false	boolean
moveTo (consumer)	After processing a mail message, it can be moved to a mail folder with the given name. You can override this configuration value, with a header with the key <code>moveTo</code> , allowing you to move messages to folder names configured at runtime.		String
peek (consumer)	Will mark the <code>javax.mail.Message</code> as peeked before processing the mail message. This applies to <code>IMAPMessage</code> messages types only. By using <code>peek</code> the mail will not be eager marked as SEEN on the mail server, which allows us to rollback the mail message if there is an error processing in Camel.	true	boolean

Name	Description	Default	Type
skipFailedMessage (consumer)	If the mail consumer cannot retrieve a given mail message, then this option allows to skip the message and move on to retrieve the next mail message. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	boolean
unseen (consumer)	Whether to limit by unseen mails only.	true	boolean
fetchSize (consumer (advanced))	Sets the maximum number of messages to consume during a poll. This can be used to avoid overloading a mail server, if a mailbox folder contains a lot of messages. Default value of -1 means no fetch size and all messages will be consumed. Setting the value to 0 is a special corner case, where Camel will not consume any messages at all.	-1	int
folderName (consumer (advanced))	The folder to poll.	INBOX	String
mapMailMessage (consumer (advanced))	Specifies whether Camel should map the received mail message to Camel body/headers/attachments. If set to true, the body of the mail message is mapped to the body of the Camel IN message, the mail headers are mapped to IN headers, and the attachments to Camel IN attachment message. If this option is set to false then the IN message contains a raw <code>javax.mail.Message</code> . You can retrieve this raw message by calling <code>exchange.getIn().getBody(javax.mail.Message.class)</code> .	true	boolean
bcc (producer)	Sets the BCC email address. Separate multiple email addresses with comma.		String
cc (producer)	Sets the CC email address. Separate multiple email addresses with comma.		String
from (producer)	The from email address.	camel@localhost	String

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
replyTo (producer)	The Reply-To recipients (the receivers of the response mail). Separate multiple email addresses with a comma.		String
subject (producer)	The Subject of the message being sent. Note: Setting the subject in the header takes precedence over this option.		String
to (producer)	Sets the To email address. Separate multiple email addresses with comma.		String
javaMailSender (producer (advanced))	To use a custom org.apache.camel.component.mail.JavaMailSender for sending emails.		JavaMailSender
additionalJavaMailProperties (advanced)	Sets additional java mail properties, that will append/override any default properties that is set based on all the other options. This is useful if you need to add some special options but want to keep the others as is.		Properties
alternativeBodyHeader (advanced)	Specifies the key to an IN message header that contains an alternative email body. For example, if you send emails in text/html format and want to provide an alternative mail body for non-HTML email clients, set the alternative mail body with this key as a header.	CamelMailAlternativeBody	String
attachmentsContentTransferEncodingResolver (advanced)	To use a custom AttachmentsContentTransferEncodingResolver to resolve what content-type-encoding to use for attachments.		AttachmentsContentTransferEncodingResolver

Name	Description	Default	Type
authenticator (advanced)	The authenticator for login. If set then the password and username are ignored. Can be used for tokens which can expire and therefore must be read dynamically.		MailAuthenticator
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
configuration (advanced)	Sets the Mail configuration.		MailConfiguration
connectionTimeout (advanced)	The connection timeout in milliseconds.	30000	int
contentType (advanced)	The mail message content type. Use text/html for HTML mails.	text/plain	String
contentTypeResolver (advanced)	Resolver to determine Content-Type for file attachments.		ContentTypeResolver
debugMode (advanced)	Enable debug mode on the underlying mail framework. The SUN Mail framework logs the debug messages to System.out by default.	false	boolean
ignoreUnsupportedCharset (advanced)	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	boolean
ignoreUriScheme (advanced)	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	boolean
javaMailProperties (advanced)	Sets the java mail options. Will clear any default properties and only use the properties provided for this method.		Properties

Name	Description	Default	Type
session (advanced)	Specifies the mail session that camel should use for all mail interactions. Useful in scenarios where mail sessions are created and managed by some other resource, such as a JavaEE container. When using a custom mail session, then the hostname and port from the mail session will be used (if configured on the session).		Session
useInlineAttachments (advanced)	Whether to use disposition inline or attachment.	false	boolean
headerFilterStrategy (filter)	To use a custom org.apache.camel.spi.HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
password (security)	The password for login. See also setAuthenticator(MailAuthenticator).		String
sslContextParameters (security)	To configure security using SSLContextParameters.		SSLContextParameters
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean
username (security)	The username for login. See also setAuthenticator(MailAuthenticator).		String

31.4. ENDPOINT OPTIONS

The Mail endpoint is configured using URI syntax:

```
imap:host:port
```

with the following path and query parameters:

31.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
host (common)	Required The mail server host name.		String
port (common)	The port number of the mail server.		int

31.4.2. Query Parameters (66 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
closeFolder (consumer)	Whether the consumer should close the folder after polling. Setting this option to false and having <code>disconnect=false</code> as well, then the consumer keep the folder open between polls.	true	boolean
copyTo (consumer)	After processing a mail message, it can be copied to a mail folder with the given name. You can override this configuration value, with a header with the key <code>copyTo</code> , allowing you to copy messages to folder names configured at runtime.		String
decodeFilename (consumer)	If set to true, the <code>MimeUtility.decodeText</code> method will be used to decode the filename. This is similar to setting JVM system property <code>mail.mime.encodefilename</code> .	false	boolean
delete (consumer)	Deletes the messages after they have been processed. This is done by setting the DELETED flag on the mail message. If false, the SEEN flag is set instead. As of Camel 2.10 you can override this configuration option by setting a header with the key <code>delete</code> to determine if the mail should be deleted or not.	false	boolean
disconnect (consumer)	Whether the consumer should disconnect after polling. If enabled this forces Camel to connect on each poll.	false	boolean
handleFailedMessage (consumer)	If the mail consumer cannot retrieve a given mail message, then this option allows to handle the caused exception by the consumer's error handler. By enable the bridge error handler on the consumer, then the Camel routing error handler can handle the exception instead. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	boolean

Name	Description	Default	Type
maxMessagesPerPoll (consumer)	Specifies the maximum number of messages to gather per poll. By default, no maximum is set. Can be used to set a limit of e.g. 1000 to avoid downloading thousands of files when the server starts up. Set a value of 0 or negative to disable this option.		int
mimeDecodeHeaders (consumer)	This option enables transparent MIME decoding and unfolding for mail headers.	false	boolean
moveTo (consumer)	After processing a mail message, it can be moved to a mail folder with the given name. You can override this configuration value, with a header with the key <code>moveTo</code> , allowing you to move messages to folder names configured at runtime.		String
peek (consumer)	Will mark the <code>javax.mail.Message</code> as peeked before processing the mail message. This applies to <code>IMAPMessage</code> messages types only. By using <code>peek</code> the mail will not be eager marked as <code>SEEN</code> on the mail server, which allows us to rollback the mail message if there is an error processing in Camel.	true	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
skipFailedMessage (consumer)	If the mail consumer cannot retrieve a given mail message, then this option allows to skip the message and move on to retrieve the next mail message. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	boolean
unseen (consumer)	Whether to limit by unseen mails only.	true	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at <code>WARN</code> or <code>ERROR</code> level and ignored.		<code>ExceptionHandler</code>

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
fetchSize (consumer (advanced))	Sets the maximum number of messages to consume during a poll. This can be used to avoid overloading a mail server, if a mailbox folder contains a lot of messages. Default value of -1 means no fetch size and all messages will be consumed. Setting the value to 0 is a special corner case, where Camel will not consume any messages at all.	-1	int
folderName (consumer (advanced))	The folder to poll.	INBOX	String
mailUidGenerator (consumer (advanced))	A pluggable MailUidGenerator that allows to use custom logic to generate UUID of the mail message.		MailUidGenerator
mapMailMessage (consumer (advanced))	Specifies whether Camel should map the received mail message to Camel body/headers/attachments. If set to true, the body of the mail message is mapped to the body of the Camel IN message, the mail headers are mapped to IN headers, and the attachments to Camel IN attachment message. If this option is set to false then the IN message contains a raw <code>javax.mail.Message</code> . You can retrieve this raw message by calling <code>exchange.getIn().getBody(javax.mail.Message.class)</code> .	true	boolean
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
postProcessAction (consumer (advanced))	Refers to an <code>MailBoxPostProcessAction</code> for doing post processing tasks on the mailbox once the normal processing ended.		MailBoxPostProcessAction

Name	Description	Default	Type
bcc (producer)	Sets the BCC email address. Separate multiple email addresses with comma.		String
cc (producer)	Sets the CC email address. Separate multiple email addresses with comma.		String
from (producer)	The from email address.	camel@localhost	String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
replyTo (producer)	The Reply-To recipients (the receivers of the response mail). Separate multiple email addresses with a comma.		String
subject (producer)	The Subject of the message being sent. Note: Setting the subject in the header takes precedence over this option.		String
to (producer)	Sets the To email address. Separate multiple email addresses with comma.		String
javaMailSender (producer (advanced))	To use a custom org.apache.camel.component.mail.JavaMailSender for sending emails.		JavaMailSender
additionalJavaMailProperties (advanced)	Sets additional java mail properties, that will append/override any default properties that is set based on all the other options. This is useful if you need to add some special options but want to keep the others as is.		Properties

Name	Description	Default	Type
alternativeBodyHeader (advanced)	Specifies the key to an IN message header that contains an alternative email body. For example, if you send emails in text/html format and want to provide an alternative mail body for non-HTML email clients, set the alternative mail body with this key as a header.	CamelMailAlternativeBody	String
attachmentsContentTransferEncodingResolver (advanced)	To use a custom AttachmentsContentTransferEncodingResolver to resolve what content-type-encoding to use for attachments.		AttachmentsContentTransferEncodingResolver
authenticator (advanced)	The authenticator for login. If set then the password and username are ignored. Can be used for tokens which can expire and therefore must be read dynamically.		MailAuthenticator
binding (advanced)	Sets the binding used to convert from a Camel message to and from a Mail message.		MailBinding
connectionTimeout (advanced)	The connection timeout in milliseconds.	30000	int
contentType (advanced)	The mail message content type. Use text/html for HTML mails.	text/plain	String
contentTypeResolver (advanced)	Resolver to determine Content-Type for file attachments.		ContentTypeResolver
debugMode (advanced)	Enable debug mode on the underlying mail framework. The SUN Mail framework logs the debug messages to System.out by default.	false	boolean
headerFilterStrategy (advanced)	To use a custom org.apache.camel.spi.HeaderFilterStrategy to filter headers.		HeaderFilterStrategy
ignoreUnsupportedCharset (advanced)	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	boolean

Name	Description	Default	Type
ignoreUriScheme (advanced)	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	boolean
javaMailProperties (advanced)	Sets the java mail options. Will clear any default properties and only use the properties provided for this method.		Properties
session (advanced)	Specifies the mail session that camel should use for all mail interactions. Useful in scenarios where mail sessions are created and managed by some other resource, such as a JavaEE container. When using a custom mail session, then the hostname and port from the mail session will be used (if configured on the session).		Session
useInlineAttachments (advanced)	Whether to use disposition inline or attachment.	false	boolean
idempotentRepository (filter)	A pluggable repository <code>org.apache.camel.spi.IdempotentRepository</code> which allows to cluster consuming from the same mailbox, and let the repository coordinate whether a mail message is valid for the consumer to process. By default no repository is in use.		IdempotentRepository
idempotentRepositoryRemoveOnCommit (filter)	When using idempotent repository, then when the mail message has been successfully processed and is committed, should the message id be removed from the idempotent repository (default) or be kept in the repository. By default its assumed the message id is unique and has no value to be kept in the repository, because the mail message will be marked as seen/moved or deleted to prevent it from being consumed again. And therefore having the message id stored in the idempotent repository has little value. However this option allows to store the message id, for whatever reason you may have.	true	boolean
searchTerm (filter)	Refers to a <code>javax.mail.search.SearchTerm</code> which allows to filter mails based on search criteria such as subject, body, from, sent after a certain date etc.		SearchTerm
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		int

Name	Description	Default	Type
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	60000	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object

Name	Description	Default	Type
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> • NANOSECONDS • MICROSECONDS • MILLISECONDS • SECONDS • MINUTES • HOURS • DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
password (security)	The password for login. See also setAuthenticator(MailAuthenticator).		String
sslContextParameters (security)	To configure security using SSLContextParameters.		SSLContextParameters
username (security)	The username for login. See also setAuthenticator(MailAuthenticator).		String
sortTerm (sort)	Sorting order for messages. Only natively supported for IMAP. Emulated to some degree when using POP3 or when IMAP server does not have the SORT capability.		SortTerm[]

31.4.3. Sample endpoints

Typically, you specify a URI with login credentials as follows (taking SMTP as an example):

```
smtp://[username@]host[:port][?password=somepwd]
```

Alternatively, it is possible to specify both the user name and the password as query options:

```
smtp://host[:port]?password=somepwd&username=someuser
```

For example:

```
smtp://mycompany.mailserver:30?password=tiger&username=scott
```

31.4.4. Component alias names

- IMAP
- IMAPs
- POP3s
- SMTP
- SMTPs

31.4.5. Default ports

Default port numbers are supported. If the port number is omitted, Camel determines the port number to use based on the protocol.

Protocol	Default Port Number
SMTP	25
SMTPS	465
POP3	110
POP3S	995
IMAP	143
IMAPS	993

31.5. SSL SUPPORT

The underlying mail framework is responsible for providing SSL support. You may either configure SSL/TLS support by completely specifying the necessary Java Mail API configuration options, or you may provide a configured `SSLContextParameters` through the component or endpoint configuration.

31.5.1. Using the JSSE Configuration Utility

The mail component supports SSL/TLS configuration through the [Camel JSSE Configuration Utility](#). This utility greatly decreases the amount of component specific code you need to write and is

configurable at the endpoint and component levels. The following examples demonstrate how to use the utility with the mail component.

Programmatic configuration of the endpoint

```

KeyStoreParameters ksp = new KeyStoreParameters();
ksp.setResource("/users/home/server/truststore.jks");
ksp.setPassword("keystorePassword");
TrustManagersParameters tmp = new TrustManagersParameters();
tmp.setKeyStore(ksp);
SSLContextParameters scp = new SSLContextParameters();
scp.setTrustManagers(tmp);
Registry registry = ...
registry.bind("sslContextParameters", scp);
...
from(...)
    .to("smtps://smtp.google.com?
username=user@gmail.com&password=password&sslContextParameters=#sslContextParameters");

```

Spring DSL based configuration of endpoint

```

...
<camel:sslContextParameters id="sslContextParameters">
  <camel:trustManagers>
    <camel:keyStore resource="/users/home/server/truststore.jks" password="keystorePassword"/>
  </camel:trustManagers>
</camel:sslContextParameters>...
...
<to uri="smtps://smtp.google.com?
username=user@gmail.com&password=password&sslContextParameters=#sslContextParameters"/
>...

```

31.5.2. Configuring JavaMail Directly

Camel uses Jakarta JavaMail, which only trusts certificates issued by well known Certificate Authorities (the default JVM trust configuration). If you issue your own certificates, you have to import the CA certificates into the JVM's Java trust/key store files, override the default JVM trust/key store files (see **SSLNOTES.txt** in JavaMail for details).

31.6. MAIL MESSAGE CONTENT

Camel uses the message exchange's IN body as the [MimeMessage](#) text content. The body is converted to `String.class`.

Camel copies all of the exchange's IN headers to the [MimeMessage](#) headers.

The subject of the [MimeMessage](#) can be configured using a header property on the IN message. The code below demonstrates this:

The same applies for other [MimeMessage](#) headers such as recipients, so you can use a header property as To:

When using the `MailProducer` to send the mail to server, you should be able to get the message id of the [MimeMessage](#) with the key `CamelMailMessageId` from the Camel message header.

31.7. HEADERS TAKE PRECEDENCE OVER PRE-CONFIGURED RECIPIENTS

The recipients specified in the message headers always take precedence over recipients pre-configured in the endpoint URI. The idea is that if you provide any recipients in the message headers, that is what you get. The recipients pre-configured in the endpoint URI are treated as a fallback.

In the sample code below, the email message is sent to **davsclaus@apache.org**, because it takes precedence over the pre-configured recipient, **info@mycompany.com**. Any **CC** and **BCC** settings in the endpoint URI are also ignored and those recipients will not receive any mail. The choice between headers and pre-configured settings is all or nothing: the mail component *either* takes the recipients exclusively from the headers or exclusively from the pre-configured settings. It is not possible to mix and match headers and pre-configured settings.

```
Map<String, Object> headers = new HashMap<String, Object>();
headers.put("to", "davsclaus@apache.org");

template.sendBodyAndHeaders("smtp://admin@localhost?to=info@mycompany.com", "Hello World",
headers);
```

31.8. MULTIPLE RECIPIENTS FOR EASIER CONFIGURATION

It is possible to set multiple recipients using a comma-separated or a semicolon-separated list. This applies both to header settings and to settings in an endpoint URI. For example:

```
Map<String, Object> headers = new HashMap<String, Object>();
headers.put("to", "davsclaus@apache.org ; jstrachan@apache.org ; ningjiang@apache.org");
```

The preceding example uses a semicolon, `;`, as the separator character.

31.9. SETTING SENDER NAME AND EMAIL

You can specify recipients in the format, **name <email>**, to include both the name and the email address of the recipient.

For example, you define the following headers on the a Message:

```
Map headers = new HashMap();
map.put("To", "Claus Ibsen <davsclaus@apache.org>");
map.put("From", "James Strachan <jstrachan@apache.org>");
map.put("Subject", "Camel is cool");
```

31.10. JAVAMAIL API (EX SUN JAVAMAIL)

[JavaMail API](#) is used under the hood for consuming and producing mails. We encourage end-users to consult these references when using either POP3 or IMAP protocol. Note particularly that POP3 has a much more limited set of features than IMAP.

- [JavaMail POP3 API](#)
- [JavaMail IMAP API](#)
- And generally about the [MAIL Flags](#)

31.11. SAMPLES

We start with a simple route that sends the messages received from a JMS queue as emails. The email account is the **admin** account on **mymailserver.com**.

```
from("jms://queue:subscription").to("smtp://admin@mymailserver.com?password=secret");
```

In the next sample, we poll a mailbox for new emails once every minute.

```
from("imap://admin@mymailserver.com?password=secret&unseen=true&delay=60000")
.to("seda://mails");
```

31.12. SENDING MAIL WITH ATTACHMENT SAMPLE



NOTE

Attachments are not support by all Camel components

The *Attachments API* is based on the Java Activation Framework and is generally only used by the Mail API. Since many of the other Camel components do not support attachments, the attachments could potentially be lost as they propagate along the route. The rule of thumb, therefore, is to add attachments just before sending a message to the mail endpoint.

The mail component supports attachments. In the sample below, we send a mail message containing a plain text message with a logo file attachment.

31.13. SSL SAMPLE

In this sample, we want to poll our Google mail inbox for mails. To download mail onto a local mail client, Google mail requires you to enable and configure SSL. This is done by logging into your Google mail account and changing your settings to allow IMAP access. Google have extensive documentation on how to do this.

```
from("imaps://imap.gmail.com?
username=YOUR_USERNAME@gmail.com&password=YOUR_PASSWORD"
+ "&delete=false&unseen=true&delay=60000").to("log:newmail");
```

The preceding route polls the Google mail inbox for new mails once every minute and logs the received messages to the **newmail** logger category.

Running the sample with **DEBUG** logging enabled, we can monitor the progress in the logs:

```
2008-05-08 06:32:09,640 DEBUG MailConsumer - Connecting to MailStore
imaps://imap.gmail.com:993 (SSL enabled), folder=INBOX
2008-05-08 06:32:11,203 DEBUG MailConsumer - Polling mailfolder: imaps://imap.gmail.com:993
(SSL enabled), folder=INBOX
2008-05-08 06:32:11,640 DEBUG MailConsumer - Fetching 1 messages. Total 1 messages.
2008-05-08 06:32:12,171 DEBUG MailConsumer - Processing message: messageNumber=[332],
from=[James Bond <007@mi5.co.uk>], to=YOUR_USERNAME@gmail.com], subject=[...
2008-05-08 06:32:12,187 INFO newmail - Exchange[MailMessage: messageNumber=[332], from=
[James Bond <007@mi5.co.uk>], to=YOUR_USERNAME@gmail.com], subject=[...
```

31.14. CONSUMING MAILS WITH ATTACHMENT SAMPLE

In this sample we poll a mailbox and store all attachments from the mails as files. First, we define a route to poll the mailbox. As this sample is based on google mail, it uses the same route as shown in the SSL sample:

```
from("imaps://imap.gmail.com?
username=YOUR_USERNAME@gmail.com&password=YOUR_PASSWORD"
+ "&delete=false&unseen=true&delay=60000").process(new MyMailProcessor());
```

Instead of logging the mail we use a processor where we can process the mail from java code:

```
public void process(Exchange exchange) throws Exception {
    // the API is a bit clunky so we need to loop
    AttachmentMessage attachmentMessage = exchange.getMessage(AttachmentMessage.class);
    Map<String, DataHandler> attachments = attachmentMessage.getAttachments();
    if (attachments.size() > 0) {
        for (String name : attachments.keySet()) {
            DataHandler dh = attachments.get(name);
            // get the file name
            String filename = dh.getName();

            // get the content and convert it to byte[]
            byte[] data = exchange.getContext().getTypeConverter()
                .convertTo(byte[].class, dh.getInputStream());

            // write the data to a file
            FileOutputStream out = new FileOutputStream(filename);
            out.write(data);
            out.flush();
            out.close();
        }
    }
}
```

As you can see the API to handle attachments is a bit clunky but it's there so you can get the **javax.activation.DataHandler** so you can handle the attachments using standard API.

31.15. HOW TO SPLIT A MAIL MESSAGE WITH ATTACHMENTS

In this example we consume mail messages which may have a number of attachments. What we want to do is to use the Splitter EIP per individual attachment, to process the attachments separately. For example if the mail message has 5 attachments, we want the Splitter to process five messages, each having a single attachment. To do this we need to provide a custom Expression to the Splitter where we provide a List<Message> that contains the five messages with the single attachment.

The code is provided out of the box in Camel 2.10 onwards in the **camel-mail** component. The code is in the class: **org.apache.camel.component.mail.SplitAttachmentsExpression**, which you can find in the source code [here](#).

In the Camel route you then need to use this Expression in the route as shown below:

If you use XML DSL then you need to declare a method call expression in the Splitter as shown below


```
<split>
  <method beanType="org.apache.camel.component.mail.SplitAttachmentsExpression"/>
  <to uri="mock:split"/>
</split>
```

You can also split the attachments as `byte[]` to be stored as the message body. This is done by creating the expression with boolean `true`

```
SplitAttachmentsExpression split = SplitAttachmentsExpression(true);
```

And then use the expression with the splitter EIP.

31.16. USING CUSTOM SEARCHTERM

You can configure a **searchTerm** on the **MailEndpoint** which allows you to filter out unwanted mails.

For example to filter mails to contain Camel in either Subject or Text you can do as follows:

```
<route>
  <from uri="imaps://mymailserver?
username=foo&password=secret&searchTerm.subjectOrBody=Camel"/>
  <to uri="bean:myBean"/>
</route>
```

Notice we use the **"searchTerm.subjectOrBody"** as parameter key to indicate that we want to search on mail subject or body, to contain the word "Camel".

The class **org.apache.camel.component.mail.SimpleSearchTerm** has a number of options you can configure:

Or to get the new unseen emails going 24 hours back in time you can do. Notice the "now-24h" syntax. See the table below for more details.

```
<route>
  <from uri="imaps://mymailserver?
username=foo&password=secret&searchTerm.fromSentDate=now-24h"/>
  <to uri="bean:myBean"/>
</route>
```

You can have multiple `searchTerm` in the endpoint uri configuration. They would then be combined together using AND operator, eg so both conditions must match. For example to get the last unseen emails going back 24 hours which has Camel in the mail subject you can do:

```
<route>
  <from uri="imaps://mymailserver?
username=foo&password=secret&searchTerm.subject=Camel&searchTerm.fromSentDate=now-
24h"/>
  <to uri="bean:myBean"/>
</route>
```

The **SimpleSearchTerm** is designed to be easily configurable from a POJO, so you can also configure it using a `<bean>` style in XML

```
<bean id="mySearchTerm" class="org.apache.camel.component.mail.SimpleSearchTerm">
```

```
<property name="subject" value="Order"/>
<property name="to" value="acme-order@acme.com"/>
<property name="fromSentDate" value="now"/>
</bean>
```

You can then refer to this bean, using `#beanId` in your Camel route as shown:

```
<route>
  <from uri="imaps://mymailserver?
username=foo&password=secret&searchTerm=#mySearchTerm"/>
  <to uri="bean:myBean"/>
</route>
```

In Java there is a builder class to build compound **SearchTerms** using the **org.apache.camel.component.mail.SearchTermBuilder** class. This allows you to build complex terms such as:

```
// we just want the unseen mails which is not spam
SearchTermBuilder builder = new SearchTermBuilder();

builder.unseen().body(Op.not, "Spam").subject(Op.not, "Spam")
  // which was sent from either foo or bar
  .from("foo@somewhere.com").from(Op.or, "bar@somewhere.com");
  // .. and we could continue building the terms

SearchTerm term = builder.build();
```

31.17. POLLING OPTIMIZATION

The parameter `maxMessagePerPoll` and `fetchSize` allow you to restrict the number message that should be processed for each poll. These parameters should help to prevent bad performance when working with folders that contain a lot of messages. In previous versions these parameters have been evaluated too late, so that big mailboxes could still cause performance problems. With Camel 3.1 these parameters are evaluated earlier during the poll to avoid these problems.

31.18. USING HEADERS WITH ADDITIONAL JAVA MAIL SENDER PROPERTIES

When sending mails, then you can provide dynamic java mail properties for the **JavaMailSender** from the Exchange as message headers with keys starting with **java.smtp..**

You can set any of the **java.smtp** properties which you can find in the Java Mail documentation.

For example to provide a dynamic uuid in **java.smtp.from** (SMTP MAIL command):

```
.setHeader("from", constant("reply2me@foo.com"));
.setHeader("java.smtp.from", method(UUID.class, "randomUUID"));
.to("smtp://mymailserver:1234");
```



NOTE

This is only supported when **not** using a custom **JavaMailSender**.

31.19. SPRING BOOT AUTO-CONFIGURATION

When using imap with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-mail-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 50 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.mail.additional-java-mail-properties</code>	Sets additional java mail properties, that will append/override any default properties that is set based on all the other options. This is useful if you need to add some special options but want to keep the others as is. The option is a <code>java.util.Properties</code> type.		Properties
<code>camel.component.mail.alternative-body-header</code>	Specifies the key to an IN message header that contains an alternative email body. For example, if you send emails in text/html format and want to provide an alternative mail body for non-HTML email clients, set the alternative mail body with this key as a header.	Camel MailAlternativeBody	String
<code>camel.component.mail.attachments-content-transfer-encoding-resolver</code>	To use a custom <code>AttachmentsContentTransferEncodingResolver</code> to resolve what content-type-encoding to use for attachments. The option is a <code>org.apache.camel.component.mail.AttachmentsContentTransferEncodingResolver</code> type.		<code>AttachmentsContentTransferEncodingResolver</code>
<code>camel.component.mail.authenticator</code>	The authenticator for login. If set then the password and username are ignored. Can be used for tokens which can expire and therefore must be read dynamically. The option is a <code>org.apache.camel.component.mail.MailAuthenticator</code> type.		<code>MailAuthenticator</code>
<code>camel.component.mail.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.mail.bcc</code>	Sets the BCC email address. Separate multiple email addresses with comma.		String
<code>camel.component.mail.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.mail.cc</code>	Sets the CC email address. Separate multiple email addresses with comma.		String
<code>camel.component.mail.close-folder</code>	Whether the consumer should close the folder after polling. Setting this option to false and having <code>disconnect=false</code> as well, then the consumer keep the folder open between polls.	true	Boolean
<code>camel.component.mail.configuration</code>	Sets the Mail configuration. The option is a <code>org.apache.camel.component.mail.MailConfiguration</code> type.		MailConfiguration
<code>camel.component.mail.connection-timeout</code>	The connection timeout in milliseconds.	30000	Integer
<code>camel.component.mail.content-type</code>	The mail message content type. Use text/html for HTML mails.	text/plain	String
<code>camel.component.mail.content-type-resolver</code>	Resolver to determine Content-Type for file attachments. The option is a <code>org.apache.camel.component.mail.ContentTypeResolver</code> type.		ContentTypeResolver
<code>camel.component.mail.copy-to</code>	After processing a mail message, it can be copied to a mail folder with the given name. You can override this configuration value, with a header with the key <code>copyTo</code> , allowing you to copy messages to folder names configured at runtime.		String
<code>camel.component.mail.debug-mode</code>	Enable debug mode on the underlying mail framework. The SUN Mail framework logs the debug messages to <code>System.out</code> by default.	false	Boolean

Name	Description	Default	Type
camel.component.mail.decode-filename	If set to true, the MimeUtility.decodeText method will be used to decode the filename. This is similar to setting JVM system property mail.mime.encodefilename.	false	Boolean
camel.component.mail.delete	Deletes the messages after they have been processed. This is done by setting the DELETED flag on the mail message. If false, the SEEN flag is set instead. As of Camel 2.10 you can override this configuration option by setting a header with the key delete to determine if the mail should be deleted or not.	false	Boolean
camel.component.mail.disconnect	Whether the consumer should disconnect after polling. If enabled this forces Camel to connect on each poll.	false	Boolean
camel.component.mail.enabled	Whether to enable auto configuration of the mail component. This is enabled by default.		Boolean
camel.component.mail.fetch-size	Sets the maximum number of messages to consume during a poll. This can be used to avoid overloading a mail server, if a mailbox folder contains a lot of messages. Default value of -1 means no fetch size and all messages will be consumed. Setting the value to 0 is a special corner case, where Camel will not consume any messages at all.	-1	Integer
camel.component.mail.folder-name	The folder to poll.	INBOX	String
camel.component.mail.from	The from email address.	camel@localhost	String
camel.component.mail.handle-failed-message	If the mail consumer cannot retrieve a given mail message, then this option allows to handle the caused exception by the consumer's error handler. By enable the bridge error handler on the consumer, then the Camel routing error handler can handle the exception instead. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	Boolean
camel.component.mail.header-filter-strategy	To use a custom org.apache.camel.spi.HeaderFilterStrategy to filter header to and from Camel message. The option is a org.apache.camel.spi.HeaderFilterStrategy type.		HeaderFilterStrategy

Name	Description	Default	Type
<code>camel.component.mail.ignore-unsupported-charset</code>	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	Boolean
<code>camel.component.mail.ignore-uri-scheme</code>	Option to let Camel ignore unsupported charset in the local JVM when sending mails. If the charset is unsupported then charset=XXX (where XXX represents the unsupported charset) is removed from the content-type and it relies on the platform default instead.	false	Boolean
<code>camel.component.mail.java-mail-properties</code>	Sets the java mail options. Will clear any default properties and only use the properties provided for this method. The option is a <code>java.util.Properties</code> type.		Properties
<code>camel.component.mail.java-mail-sender</code>	To use a custom <code>org.apache.camel.component.mail.JavaMailSender</code> for sending emails. The option is a <code>org.apache.camel.component.mail.JavaMailSender</code> type.		JavaMailSender
<code>camel.component.mail.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.mail.map-mail-message</code>	Specifies whether Camel should map the received mail message to Camel body/headers/attachments. If set to true, the body of the mail message is mapped to the body of the Camel IN message, the mail headers are mapped to IN headers, and the attachments to Camel IN attachment message. If this option is set to false then the IN message contains a raw <code>javax.mail.Message</code> . You can retrieve this raw message by calling <code>exchange.getIn().getBody(javax.mail.Message.class)</code> .	true	Boolean

Name	Description	Default	Type
camel.component.mail.mime-decode-headers	This option enables transparent MIME decoding and unfolding for mail headers.	false	Boolean
camel.component.mail.move-to	After processing a mail message, it can be moved to a mail folder with the given name. You can override this configuration value, with a header with the key <code>moveTo</code> , allowing you to move messages to folder names configured at runtime.		String
camel.component.mail.password	The password for login. See also <code>setAuthenticator(MailAuthenticator)</code> .		String
camel.component.mail.peek	Will mark the <code>javax.mail.Message</code> as peeked before processing the mail message. This applies to <code>IMAPMessage</code> messages types only. By using <code>peek</code> the mail will not be eager marked as <code>SEEN</code> on the mail server, which allows us to rollback the mail message if there is an error processing in Camel.	true	Boolean
camel.component.mail.reply-to	The Reply-To recipients (the receivers of the response mail). Separate multiple email addresses with a comma.		String
camel.component.mail.session	Specifies the mail session that camel should use for all mail interactions. Useful in scenarios where mail sessions are created and managed by some other resource, such as a JavaEE container. When using a custom mail session, then the hostname and port from the mail session will be used (if configured on the session). The option is a <code>javax.mail.Session</code> type.		Session
camel.component.mail.skip-failed-message	If the mail consumer cannot retrieve a given mail message, then this option allows to skip the message and move on to retrieve the next mail message. The default behavior would be the consumer throws an exception and no mails from the batch would be able to be routed by Camel.	false	Boolean
camel.component.mail.ssl-context-parameters	To configure security using <code>SSLContextParameters</code> . The option is a <code>org.apache.camel.support.jsse.SSLContextParameters</code> type.		<code>SSLContextParameters</code>
camel.component.mail.subject	The Subject of the message being sent. Note: Setting the subject in the header takes precedence over this option.		String

Name	Description	Default	Type
<code>camel.component.mail.to</code>	Sets the To email address. Separate multiple email addresses with comma.		String
<code>camel.component.mail.unseen</code>	Whether to limit by unseen mails only.	true	Boolean
<code>camel.component.mail.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	false	Boolean
<code>camel.component.mail.use-inline-attachments</code>	Whether to use disposition inline or attachment.	false	Boolean
<code>camel.component.mail.username</code>	The username for login. See also <code>setAuthenticator(MailAuthenticator)</code> .		String
<code>camel.dataformat.mime-multipart.binary-content</code>	Defines whether the content of binary parts in the MIME multipart is binary (true) or Base-64 encoded (false) Default is false.	false	Boolean
<code>camel.dataformat.mime-multipart.enabled</code>	Whether to enable auto configuration of the mime-multipart data format. This is enabled by default.		Boolean
<code>camel.dataformat.mime-multipart.headers-inline</code>	Defines whether the MIME-Multipart headers are part of the message body (true) or are set as Camel headers (false). Default is false.	false	Boolean
<code>camel.dataformat.mime-multipart.include-headers</code>	A regex that defines which Camel headers are also included as MIME headers into the MIME multipart. This will only work if <code>headersInline</code> is set to true. Default is to include no headers.		String
<code>camel.dataformat.mime-multipart.multipart-sub-type</code>	Specify the subtype of the MIME Multipart. Default is mixed.	mixed	String
<code>camel.dataformat.mime-multipart.multipart-without-attachment</code>	Defines whether a message without attachment is also marshaled into a MIME Multipart (with only one body part). Default is false.	false	Boolean

CHAPTER 32. MASTER

Only consumer is supported

The Camel-Master endpoint provides a way to ensure only a single consumer in a cluster consumes from a given endpoint; with automatic failover if that JVM dies.


This can be very useful if you need to consume from some legacy back end which either doesn't support concurrent consumption or due to commercial or stability reasons you can only have a single connection at any point in time.

32.1. USING THE MASTER ENDPOINT

Just prefix any camel endpoint with **master:someName:** where *someName* is a logical name and is used to acquire the master lock. e.g.

```
from("master:cheese:jms:foo").to("activemq:wine");
```

In this example, the master component ensures that the route is only active in one node, at any given time, in the cluster. So if there are 8 nodes in the cluster, then the master component will elect one route to be the leader, and only this route will be active, and hence only this route will consume messages from **jms:foo**. In case this route is stopped or unexpectedly terminated, then the master component will detect this, and re-elect another node to be active, which will then become active and start consuming messages from **jms:foo**.



NOTE

Apache ActiveMQ 5.x has such feature out of the box called [Exclusive Consumers](#).

32.2. URI FORMAT

```
master:namespace:endpoint[?options]
```

Where endpoint is any Camel endpoint you want to run in master/slave mode.

32.3. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

32.3.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

32.3.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

32.4. COMPONENT OPTIONS

The Master component supports 4 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
service (advanced)	Inject the service to use.		CamelClusterService
serviceSelector (advanced)	Inject the service selector used to lookup the CamelClusterService to use.		Selector

32.5. ENDPOINT OPTIONS

The Master endpoint is configured using URI syntax:

master:namespace:delegateUri

with the following path and query parameters:

32.5.1. Path Parameters (2 parameters)

Name	Description	Default	Type
namespace (consumer)	Required The name of the cluster namespace to use.		String
delegateUri (consumer)	Required The endpoint uri to use in master/slave mode.		String

32.5.2. Query Parameters (3 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

32.6. EXAMPLE

You can protect a clustered Camel application to only consume files from one active node.

```
// the file endpoint we want to consume from
String url = "file:target/inbox?delete=true";

// use the camel master component in the clustered group named myGroup
// to run a master/slave mode in the following Camel url
from("master:myGroup:" + url)
    .log(name + " - Received file: ${file:name}")
    .delay(delay)
    .log(name + " - Done file:  ${file:name}")
    .to("file:target/outbox");
```

The master component leverages CamelClusterService you can configure using

- **Java**

```
ZooKeeperClusterService service = new ZooKeeperClusterService();
service.setId("camel-node-1");
service.setNodes("myzk:2181");
service.setBasePath("/camel/cluster");

context.addService(service)
```

- **Xml (Spring/Blueprint)**

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring
    http://camel.apache.org/schema/spring/camel-spring.xsd">

  <bean id="cluster"
class="org.apache.camel.component.zookeeper.cluster.ZooKeeperClusterService">
    <property name="id" value="camel-node-1"/>
    <property name="basePath" value="/camel/cluster"/>
    <property name="nodes" value="myzk:2181"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/spring" autoStartup="false">
    ...
  </camelContext>

</beans>
```

- **Spring boot**

```
camel.component.zookeeper.cluster.service.enabled = true
camel.component.zookeeper.cluster.service.id      = camel-node-1
camel.component.zookeeper.cluster.service.base-path = /camel/cluster
camel.component.zookeeper.cluster.service.nodes   = myzk:2181
```

32.7. IMPLEMENTATIONS

Camel provides the following ClusterService implementations:

- camel-consul
- camel-file
- camel-infinispan
- camel-jgroups-raft
- camel-jgroups
- camel-kubernetes
- camel-zookeeper

32.8. SPRING BOOT AUTO-CONFIGURATION

When using master with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-master-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 5 options, which are listed below.

Name	Description	Default	Type
camel.component.master.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.master.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.master.enabled	Whether to enable auto configuration of the master component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component.master.service	Inject the service to use. The option is a <code>org.apache.camel.cluster.CamelClusterService</code> type.		<code>CamelClusterService</code>
camel.component.master.service-selector	Inject the service selector used to lookup the <code>CamelClusterService</code> to use. The option is a <code>org.apache.camel.cluster.CamelClusterService.Selector</code> type.		<code>CamelClusterService\$Selector</code>

CHAPTER 33. MLLP

Both producer and consumer are supported

The MLLP component is specifically designed to handle the nuances of the MLLP protocol and provide the functionality required by Healthcare providers to communicate with other systems using the MLLP protocol.

The MLLP component provides a simple configuration URI, automated HL7 acknowledgment generation and automatic acknowledgment interrogation.

The MLLP protocol does not typically use a large number of concurrent TCP connections - a single active TCP connection is the normal case. Therefore, the MLLP component uses a simple thread-per-connection model based on standard Java Sockets. This keeps the implementation simple and eliminates the dependencies on only Camel itself.

The component supports the following:

- A Camel consumer using a TCP Server
- A Camel producer using a TCP Client

The MLLP component use **byte[]** payloads, and relies on Camel type conversion to convert **byte[]** to other types.

Maven users will need to add the following dependency to their pom.xml for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-mlp</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

33.1. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

33.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

33.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

33.2. COMPONENT OPTIONS

The MLLP component supports 30 options, which are listed below.

Name	Description	Default	Type
autoAck (common)	Enable/Disable the automatic generation of a MLLP Acknowledgement MLLP Consumers only.	true	boolean
charsetName (common)	Sets the default charset to use.		String
configuration (common)	Sets the default configuration to use when creating MLLP endpoints.		MllpConfiguration
hl7Headers (common)	Enable/Disable the automatic generation of message headers from the HL7 Message MLLP Consumers only.	true	boolean
requireEndOfData (common)	Enable/Disable strict compliance to the MLLP standard. The MLLP standard specifies START_OF_BLOCKhl7 payloadEND_OF_BLOCKEND_OF_DATA, however, some systems do not send the final END_OF_DATA byte. This setting controls whether or not the final END_OF_DATA byte is required or optional.	true	boolean
stringPayload (common)	Enable/Disable converting the payload to a String. If enabled, HL7 Payloads received from external systems will be validated converted to a String. If the charsetName property is set, that character set will be used for the conversion. If the charsetName property is not set, the value of MSH-18 will be used to determine the appropriate character set. If MSH-18 is not set, then the default ISO-8859-1 character set will be use.	true	boolean

Name	Description	Default	Type
validatePayload (common)	Enable/Disable the validation of HL7 Payloads If enabled, HL7 Payloads received from external systems will be validated (see <code>HL7Util.generateInvalidPayloadExceptionMessage</code> for details on the validation). If and invalid payload is detected, a <code>MllpInvalidMessageException</code> (for consumers) or a <code>MllpInvalidAcknowledgementException</code> will be thrown.	false	boolean
acceptTimeout (consumer)	Timeout (in milliseconds) while waiting for a TCP connection TCP Server Only.	60000	int
backlog (consumer)	The maximum queue length for incoming connection indications (a request to connect) is set to the backlog parameter. If a connection indication arrives when the queue is full, the connection is refused.	5	Integer
bindRetryInterval (consumer)	TCP Server Only - The number of milliseconds to wait between bind attempts.	5000	int
bindTimeout (consumer)	TCP Server Only - The number of milliseconds to retry binding to a server port.	30000	int
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to receive incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. If disabled, the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions by logging them at WARN or ERROR level and ignored.	true	boolean
lenientBind (consumer)	TCP Server Only - Allow the endpoint to start before the TCP ServerSocket is bound. In some environments, it may be desirable to allow the endpoint to start before the TCP ServerSocket is bound.	false	boolean
maxConcurrentConsumers (consumer)	The maximum number of concurrent MLLP Consumer connections that will be allowed. If a new connection is received and the maximum is number are already established, the new connection will be reset immediately.	5	int
reuseAddress (consumer)	Enable/disable the <code>SO_REUSEADDR</code> socket option.	false	Boolean

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 	InOut	ExchangePattern
connectTimeout (producer)	Timeout (in milliseconds) for establishing for a TCP connection TCP Client only.	30000	int
idleTimeoutStrategy (producer)	decide what action to take when idle timeout occurs. Possible values are : RESET: set SO_LINGER to 0 and reset the socket CLOSE: close the socket gracefully default is RESET. Enum values: <ul style="list-style-type: none"> ● RESET ● CLOSE 	RESET	MllpIdleTimeoutStrategy
keepAlive (producer)	Enable/disable the SO_KEEPALIVE socket option.	true	Boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
tcpNoDelay (producer)	Enable/disable the TCP_NODELAY socket option.	true	Boolean

Name	Description	Default	Type
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
defaultCharset (advanced)	Set the default character set to use for byte to/from String conversions.	ISO-8859-1	String
logPhi (advanced)	Whether to log PHI.	true	Boolean
logPhiMaxBytes (advanced)	Set the maximum number of bytes of PHI that will be logged in a log entry.	5120	Integer
readTimeout (advanced)	The SO_TIMEOUT value (in milliseconds) used after the start of an MLLP frame has been received.	5000	int
receiveBufferSize (advanced)	Sets the SO_RCVBUF option to the specified value (in bytes).	8192	Integer
receiveTimeout (advanced)	The SO_TIMEOUT value (in milliseconds) used when waiting for the start of an MLLP frame.	15000	int
sendBufferSize (advanced)	Sets the SO_SNDBUF option to the specified value (in bytes).	8192	Integer
idleTimeout (tcp)	The approximate idle time allowed before the Client TCP Connection will be reset. A null value or a value less than or equal to zero will disable the idle timeout.		Integer

33.3. ENDPOINT OPTIONS

The MLLP endpoint is configured using URI syntax:

```
mlp:hostname:port
```

with the following path and query parameters:

33.3.1. Path Parameters (2 parameters)

Name	Description	Default	Type
hostname (common)	Required Hostname or IP for connection for the TCP connection. The default value is null, which means any local IP address.		String
port (common)	Required Port number for the TCP connection.		int

33.3.2. Query Parameters (26 parameters)

Name	Description	Default	Type
autoAck (common)	Enable/Disable the automatic generation of a MLLP Acknowledgement MLLP Consumers only.	true	boolean
charsetName (common)	Sets the default charset to use.		String
hl7Headers (common)	Enable/Disable the automatic generation of message headers from the HL7 Message MLLP Consumers only.	true	boolean
requireEndOfData (common)	Enable/Disable strict compliance to the MLLP standard. The MLLP standard specifies START_OF_BLOCKhl7 payload END_OF_BLOCK END_OF_DATA, however, some systems do not send the final END_OF_DATA byte. This setting controls whether or not the final END_OF_DATA byte is required or optional.	true	boolean
stringPayload (common)	Enable/Disable converting the payload to a String. If enabled, HL7 Payloads received from external systems will be validated converted to a String. If the charsetName property is set, that character set will be used for the conversion. If the charsetName property is not set, the value of MSH-18 will be used to determine the appropriate character set. If MSH-18 is not set, then the default ISO-8859-1 character set will be used.	true	boolean

Name	Description	Default	Type
validatePayload (common)	Enable/Disable the validation of HL7 Payloads If enabled, HL7 Payloads received from external systems will be validated (see <code>HL7Util.generateInvalidPayloadExceptionMessage</code> for details on the validation). If and invalid payload is detected, a <code>MllpInvalidMessageException</code> (for consumers) or a <code>MllpInvalidAcknowledgementException</code> will be thrown.	false	boolean
acceptTimeout (consumer)	Timeout (in milliseconds) while waiting for a TCP connection TCP Server Only.	60000	int
backlog (consumer)	The maximum queue length for incoming connection indications (a request to connect) is set to the backlog parameter. If a connection indication arrives when the queue is full, the connection is refused.	5	Integer
bindRetryInterval (consumer)	TCP Server Only - The number of milliseconds to wait between bind attempts.	5000	int
bindTimeout (consumer)	TCP Server Only - The number of milliseconds to retry binding to a server port.	30000	int
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to receive incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. If disabled, the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions by logging them at WARN or ERROR level and ignored.	true	boolean
lenientBind (consumer)	TCP Server Only - Allow the endpoint to start before the TCP ServerSocket is bound. In some environments, it may be desirable to allow the endpoint to start before the TCP ServerSocket is bound.	false	boolean
maxConcurrentConsumers (consumer)	The maximum number of concurrent MLLP Consumer connections that will be allowed. If a new connection is received and the maximum is number are already established, the new connection will be reset immediately.	5	int

Name	Description	Default	Type
reuseAddress (consumer)	Enable/disable the SO_REUSEADDR socket option.	false	Boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 	InOut	ExchangePattern
connectTimeout (producer)	Timeout (in milliseconds) for establishing for a TCP connection TCP Client only.	30000	int
idleTimeoutStrategy (producer)	decide what action to take when idle timeout occurs. Possible values are : RESET: set SO_LINGER to 0 and reset the socket CLOSE: close the socket gracefully default is RESET. Enum values: <ul style="list-style-type: none"> • RESET • CLOSE 	RESET	MillIdleTimeoutStrategy
keepAlive (producer)	Enable/disable the SO_KEEPALIVE socket option.	true	Boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
tcpNoDelay (producer)	Enable/disable the TCP_NODELAY socket option.	true	Boolean
readTimeout (advanced)	The SO_TIMEOUT value (in milliseconds) used after the start of an MLLP frame has been received.	5000	int
receiveBufferSize (advanced)	Sets the SO_RCVBUF option to the specified value (in bytes).	8192	Integer
receiveTimeout (advanced)	The SO_TIMEOUT value (in milliseconds) used when waiting for the start of an MLLP frame.	15000	int
sendBufferSize (advanced)	Sets the SO_SNDBUF option to the specified value (in bytes).	8192	Integer
idleTimeout (tcp)	The approximate idle time allowed before the Client TCP Connection will be reset. A null value or a value less than or equal to zero will disable the idle timeout.		Integer

33.4. MLLP CONSUMER

The MLLP Consumer supports receiving MLLP-framed messages and sending HL7 Acknowledgements. The MLLP Consumer can automatically generate the HL7 Acknowledgement (HL7 Application Acknowledgements only - AA, AE and AR), or the acknowledgement can be specified using the `CamelMllpAcknowledgement` exchange property. Additionally, the type of acknowledgement that will be generated can be controlled by setting the `CamelMllpAcknowledgementType` exchange property. The MLLP Consumer can read messages without sending any HL7 Acknowledgement if the automatic acknowledgement is disabled and exchange pattern is `InOnly`.

33.4.1. Message Headers

The MLLP Consumer adds these headers on the Camel message:

Key	Description
<code>CamelMllpLocalAddress</code>	The local TCP Address of the Socket
<code>CamelMllpRemoteAddress</code>	The local TCP Address of the Socket
<code>CamelMllpSendingApplication</code>	MSH-3 value
<code>CamelMllpSendingFacility</code>	MSH-4 value
<code>CamelMllpReceivingApplication</code>	MSH-5 value

CamelMllpReceivingFacility	MSH-6 value
CamelMllpTimestamp	MSH-7 value
CamelMllpSecurity	MSH-8 value
CamelMllpMessageType	MSH-9 value
CamelMllpEventType	MSH-9-1 value
CamelMllpTriggerEvent	MSH-9-2 value
CamelMllpMessageControllId	MSH-10 value
CamelMllpProcessingId	MSH-11 value
CamelMllpVersionId	MSH-12 value
CamelMllpCharset	MSH-18 value

All headers are String types. If a header value is missing, its value is null.

33.4.2. Exchange Properties

The type of acknowledgment the MLLP Consumer generates and state of the TCP Socket can be controlled by these properties on the Camel exchange:

Key	Type	Description
CamelMllpAcknowledgement	byte[]	If present, this property will be sent to client as the MLLP Acknowledgement
CamelMllpAcknowledgementString	String	If present and CamelMllpAcknowledgement is not present, this property will be sent to client as the MLLP Acknowledgement
CamelMllpAcknowledgementMsaText	String	If neither CamelMllpAcknowledgement or CamelMllpAcknowledgementString are present and autoAck is true, this property can be used to specify the contents of MSA-3 in the generated HL7 acknowledgement
CamelMllpAcknowledgementType	String	If neither CamelMllpAcknowledgement or CamelMllpAcknowledgementString are present and autoAck is true, this property can be used to specify the HL7 acknowledgement type (i.e. AA, AE, AR)
CamelMllpAutoAcknowledge	Boolean	Overrides the autoAck query parameter

Key	Type	Description
CamelMllpCloseConnectionBeforeSend	Boolean	If true, the Socket will be closed before sending data
CamelMllpResetConnectionBeforeSend	Boolean	If true, the Socket will be reset before sending data
CamelMllpCloseConnectionAfterSend	Boolean	If true, the Socket will be closed immediately after sending data
CamelMllpResetConnectionAfterSend	Boolean	If true, the Socket will be reset immediately after sending any data

33.5. MLLP PRODUCER

The MLLP Producer supports sending MLLP-framed messages and receiving HL7 Acknowledgements. The MLLP Producer interrogates the HL7 Acknowledgments and raises exceptions if a negative acknowledgement is received. The received acknowledgement is interrogated and an exception is raised in the event of a negative acknowledgement. The MLLP Producer can ignore acknowledgements when configured with InOnly exchange pattern.

33.5.1. Message Headers

The MLLP Producer adds these headers on the Camel message:

Key	Description
CamelMllpLocalAddress	The local TCP Address of the Socket
CamelMllpRemoteAddress	The remote TCP Address of the Socket
CamelMllpAcknowledgement	The HL7 Acknowledgment byte[] received
CamelMllpAcknowledgementString	The HL7 Acknowledgment received, converted to a String

33.5.2. Exchange Properties

The state of the TCP Socket can be controlled by these properties on the Camel exchange:

Key	Type	Description
CamelMllpCloseConnectionBeforeSend	Boolean	If true, the Socket will be closed before sending data

Key	Type	Description
CamelMllpResetConnectionBeforeSend	Boolean	If true, the Socket will be reset before sending data
CamelMllpCloseConnectionAfterSend	Boolean	If true, the Socket will be closed immediately after sending data
CamelMllpResetConnectionAfterSend	Boolean	If true, the Socket will be reset immediately after sending any data

33.6. SPRING BOOT AUTO-CONFIGURATION

When using mllp with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-mllp-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 31 options, which are listed below.

Name	Description	Default	Type
camel.component.mllp.accept-timeout	Timeout (in milliseconds) while waiting for a TCP connection TCP Server Only.	60000	Integer
camel.component.mllp.auto-ack	Enable/Disable the automatic generation of a MLLP Acknowledgement MLLP Consumers only.	true	Boolean
camel.component.mllp.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.mllp.backlog	The maximum queue length for incoming connection indications (a request to connect) is set to the backlog parameter. If a connection indication arrives when the queue is full, the connection is refused.	5	Integer

Name	Description	Default	Type
<code>camel.component.mllp.bind-retry-interval</code>	TCP Server Only - The number of milliseconds to wait between bind attempts.	5000	Integer
<code>camel.component.mllp.bind-timeout</code>	TCP Server Only - The number of milliseconds to retry binding to a server port.	30000	Integer
<code>camel.component.mllp.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to receive incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. If disabled, the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions by logging them at WARN or ERROR level and ignored.	true	Boolean
<code>camel.component.mllp.charset-name</code>	Sets the default charset to use.		String
<code>camel.component.mllp.configuration</code>	Sets the default configuration to use when creating MLLP endpoints. The option is a <code>org.apache.camel.component.mllp.MllpConfiguration</code> type.		MllpConfiguration
<code>camel.component.mllp.connect-timeout</code>	Timeout (in milliseconds) for establishing for a TCP connection TCP Client only.	30000	Integer
<code>camel.component.mllp.default-charset</code>	Set the default character set to use for byte to/from String conversions.	ISO-8859-1	String
<code>camel.component.mllp.enabled</code>	Whether to enable auto configuration of the mllp component. This is enabled by default.		Boolean
<code>camel.component.mllp.exchange-pattern</code>	Sets the exchange pattern when the consumer creates an exchange.		ExchangePattern
<code>camel.component.mllp.hl7-headers</code>	Enable/Disable the automatic generation of message headers from the HL7 Message MLLP Consumers only.	true	Boolean

Name	Description	Default	Type
<code>camel.component.mllp.idle-timeout</code>	The approximate idle time allowed before the Client TCP Connection will be reset. A null value or a value less than or equal to zero will disable the idle timeout.		Integer
<code>camel.component.mllp.idle-timeout-strategy</code>	decide what action to take when idle timeout occurs. Possible values are : RESET: set SO_LINGER to 0 and reset the socket CLOSE: close the socket gracefully default is RESET.		MllpIdleTimeoutStrategy
<code>camel.component.mllp.keep-alive</code>	Enable/disable the SO_KEEPALIVE socket option.	true	Boolean
<code>camel.component.mllp.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.mllp.lenient-bind</code>	TCP Server Only - Allow the endpoint to start before the TCP ServerSocket is bound. In some environments, it may be desirable to allow the endpoint to start before the TCP ServerSocket is bound.	false	Boolean
<code>camel.component.mllp.log-phi</code>	Whether to log PHI.	true	Boolean
<code>camel.component.mllp.log-phi-max-bytes</code>	Set the maximum number of bytes of PHI that will be logged in a log entry.	5120	Integer
<code>camel.component.mllp.max-concurrent-consumers</code>	The maximum number of concurrent MLLP Consumer connections that will be allowed. If a new connection is received and the maximum is number are already established, the new connection will be reset immediately.	5	Integer
<code>camel.component.mllp.read-timeout</code>	The SO_TIMEOUT value (in milliseconds) used after the start of an MLLP frame has been received.	5000	Integer

Name	Description	Default	Type
<code>camel.component.mllp.receive-buffer-size</code>	Sets the <code>SO_RCVBUF</code> option to the specified value (in bytes).	8192	Integer
<code>camel.component.mllp.receive-timeout</code>	The <code>SO_TIMEOUT</code> value (in milliseconds) used when waiting for the start of an MLLP frame.	15000	Integer
<code>camel.component.mllp.require-end-of-data</code>	Enable/Disable strict compliance to the MLLP standard. The MLLP standard specifies <code>START_OF_BLOCKhl7 payloadEND_OF_BLOCKEND_OF_DATA</code> , however, some systems do not send the final <code>END_OF_DATA</code> byte. This setting controls whether or not the final <code>END_OF_DATA</code> byte is required or optional.	true	Boolean
<code>camel.component.mllp.reuse-address</code>	Enable/disable the <code>SO_REUSEADDR</code> socket option.	false	Boolean
<code>camel.component.mllp.send-buffer-size</code>	Sets the <code>SO_SNDBUF</code> option to the specified value (in bytes).	8192	Integer
<code>camel.component.mllp.string-payload</code>	Enable/Disable converting the payload to a String. If enabled, HL7 Payloads received from external systems will be validated converted to a String. If the <code>charsetName</code> property is set, that character set will be used for the conversion. If the <code>charsetName</code> property is not set, the value of <code>MSH-18</code> will be used to determine the appropriate character set. If <code>MSH-18</code> is not set, then the default ISO-8859-1 character set will be used.	true	Boolean
<code>camel.component.mllp.tcp-no-delay</code>	Enable/disable the <code>TCP_NODELAY</code> socket option.	true	Boolean
<code>camel.component.mllp.validate-payload</code>	Enable/Disable the validation of HL7 Payloads. If enabled, HL7 Payloads received from external systems will be validated (see <code>HL7Util.generateInvalidPayloadExceptionMessage</code> for details on the validation). If an invalid payload is detected, a <code>MllpInvalidMessageException</code> (for consumers) or a <code>MllpInvalidAcknowledgementException</code> will be thrown.	false	Boolean

CHAPTER 34. MOCK

Only producer is supported

Testing of distributed and asynchronous processing is notoriously difficult. The [Mock](#), [Test](#) and [Dataset](#) endpoints work great with the Camel Testing Framework to simplify your unit and integration testing using [Enterprise Integration Patterns](#) and Camel's large range of Components together with the powerful Bean Integration.

The Mock component provides a powerful declarative testing mechanism, which is similar to [jMock](#) in that it allows declarative expectations to be created on any Mock endpoint before a test begins. Then the test is run, which typically fires messages to one or more endpoints, and finally the expectations can be asserted in a test case to ensure the system worked as expected.

This allows you to test various things like:

- The correct number of messages are received on each endpoint,
- The correct payloads are received, in the right order,
- Messages arrive on an endpoint in order, using some Expression to create an order testing function,
- Messages arrive match some kind of Predicate such as that specific headers have certain values, or that messages match some predicate, such as by evaluating an [XPath](#) or [XQuery](#) Expression.



NOTE

There is also the [Test endpoint](#) which is a Mock endpoint, but which uses a second endpoint to provide the list of expected message bodies and automatically sets up the Mock endpoint assertions. In other words, it's a Mock endpoint that automatically sets up its assertions from some sample messages in a File or [database](#), for example.



NOTE

Mock endpoints keep received Exchanges in memory indefinitely.

Remember that Mock is designed for testing. When you add Mock endpoints to a route, each Exchange sent to the endpoint will be stored (to allow for later validation) in memory until explicitly reset or the JVM is restarted. If you are sending high volume and/or large messages, this may cause excessive memory use. If your goal is to test deployable routes inline, consider using [NotifyBuilder](#) or [AdviceWith](#) in your tests instead of adding Mock endpoints to routes directly. There are two new options [retainFirst](#), and [retainLast](#) that can be used to limit the number of messages the Mock endpoints keep in memory.

34.1. URI FORMAT

```
mock:someName[?options]
```

Where **someName** can be any string that uniquely identifies the endpoint.

34.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

34.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

34.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

34.3. COMPONENT OPTIONS

The Mock component supports 4 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
log (producer)	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
exchangeFormatter (advanced)	Autowired Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to <code>DefaultExchangeFormatter</code> .		ExchangeFormatter

34.4. ENDPOINT OPTIONS

The Mock endpoint is configured using URI syntax:

```
mock:name
```

with the following path and query parameters:

34.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (producer)	Required Name of mock endpoint.		String

34.4.2. Query Parameters (12 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
assertPeriod (producer)	Sets a grace period after which the mock endpoint will re-assert to ensure the preliminary assertion is still valid. This is used for example to assert that exactly a number of messages arrives. For example if <code>expectedMessageCount(int)</code> was set to 5, then the assertion is satisfied when 5 or more message arrives. To ensure that exactly 5 messages arrives, then you would need to wait a little period to ensure no further message arrives. This is what you can use this method for. By default this period is disabled.		long
expectedCount (producer)	Specifies the expected number of message exchanges that should be received by this endpoint. Beware: If you want to expect that 0 messages, then take extra care, as 0 matches when the tests starts, so you need to set a assert period time to let the test run for a while to make sure there are still no messages arrived; for that use <code>setAssertPeriod(long)</code> . An alternative is to use <code>NotifyBuilder</code> , and use the notifier to know when Camel is done routing some messages, before you call the <code>assertIsSatisfied()</code> method on the mocks. This allows you to not use a fixed assert period, to speedup testing times. If you want to assert that exactly n'th message arrives to this mock endpoint, then see also the <code>setAssertPeriod(long)</code> method for further details.	-1	int
failFast (producer)	Sets whether <code>assertIsSatisfied()</code> should fail fast at the first detected failed expectation while it may otherwise wait for all expected messages to arrive before performing expectations verifications. Is by default true. Set to false to use behavior as in Camel 2.x.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
log (producer)	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	boolean
reportGroup (producer)	A number that is used to turn on throughput logging based on groups of the size.		int
resultMinimumWaitTime (producer)	Sets the minimum expected amount of time (in millis) the <code>assertIsSatisfied()</code> will wait on a latch until it is satisfied.		long
resultWaitTime (producer)	Sets the maximum amount of time (in millis) the <code>assertIsSatisfied()</code> will wait on a latch until it is satisfied.		long
retainFirst (producer)	Specifies to only retain the first n'th number of received Exchanges. This is used when testing with big data, to reduce memory consumption by not storing copies of every Exchange this mock endpoint receives. Important: When using this limitation, then the <code>getReceivedCounter()</code> will still return the actual number of received Exchanges. For example if we have received 5000 Exchanges, and have configured to only retain the first 10 Exchanges, then the <code>getReceivedCounter()</code> will still return 5000 but there is only the first 10 Exchanges in the <code>getExchanges()</code> and <code>getReceivedExchanges()</code> methods. When using this method, then some of the other expectation methods is not supported, for example the <code>expectedBodiesReceived(Object...)</code> sets a expectation on the first number of bodies received. You can configure both <code>setRetainFirst(int)</code> and <code>setRetainLast(int)</code> methods, to limit both the first and last received.	-1	int

Name	Description	Default	Type
retainLast (producer)	Specifies to only retain the last n'th number of received Exchanges. This is used when testing with big data, to reduce memory consumption by not storing copies of every Exchange this mock endpoint receives. Important: When using this limitation, then the <code>getReceivedCounter()</code> will still return the actual number of received Exchanges. For example if we have received 5000 Exchanges, and have configured to only retain the last 20 Exchanges, then the <code>getReceivedCounter()</code> will still return 5000 but there is only the last 20 Exchanges in the <code>getExchanges()</code> and <code>getReceivedExchanges()</code> methods. When using this method, then some of the other expectation methods is not supported, for example the <code>expectedBodiesReceived(Object...)</code> sets a expectation on the first number of bodies received. You can configure both <code>setRetainFirst(int)</code> and <code>setRetainLast(int)</code> methods, to limit both the first and last received.	-1	int
sleepForEmptyTest (producer)	Allows a sleep to be specified to wait to check that this endpoint really is empty when <code>expectedMessageCount(int)</code> is called with zero.		long
copyOnExchange (producer (advanced))	Sets whether to make a deep copy of the incoming Exchange when received at this mock endpoint. Is by default true.	true	boolean

34.5. SIMPLE EXAMPLE

Here's a simple example of Mock endpoint in use. First, the endpoint is resolved on the context. Then we set an expectation, and then, after the test has run, we assert that our expectations have been met:

```
MockEndpoint resultEndpoint = context.getEndpoint("mock:foo", MockEndpoint.class);

// set expectations
resultEndpoint.expectedMessageCount(2);

// send some messages

// now lets assert that the mock:foo endpoint received 2 messages
resultEndpoint.assertIsSatisfied();
```

You typically always call the method to test that the expectations were met after running a test.

Camel will by default wait 10 seconds when the `assertIsSatisfied()` is invoked. This can be configured by setting the `setResultWaitTime(millis)` method.

34.6. USING ASSERTPERIOD

When the assertion is satisfied then Camel will stop waiting and continue from the **assertIsSatisfied** method. That means if a new message arrives on the mock endpoint, just a bit later, that arrival will not affect the outcome of the assertion. Suppose you do want to test that no new messages arrives after a period thereafter, then you can do that by setting the **setAssertPeriod** method, for example:

```
MockEndpoint resultEndpoint = context.getEndpoint("mock:foo", MockEndpoint.class);
resultEndpoint.setAssertPeriod(5000);
resultEndpoint.expectedMessageCount(2);

// send some messages

// now lets assert that the mock:foo endpoint received 2 messages
resultEndpoint.assertIsSatisfied();
```

34.7. SETTING EXPECTATIONS

You can see from the Javadoc of [MockEndpoint](#) the various helper methods you can use to set expectations. The main methods are as follows:

Method	Description
expectedMessageCount(int)	To define the expected message count on the endpoint.
expectedMinimumMessageCount(int)	To define the minimum number of expected messages on the endpoint.
expectedBodiesReceived(...)	To define the expected bodies that should be received (in order).
expectedHeaderReceived(...)	To define the expected header that should be received
expectsAscending(Expression)	To add an expectation that messages are received in order, using the given Expression to compare messages.
expectsDescending(Expression)	To add an expectation that messages are received in order, using the given Expression to compare messages.
expectsNoDuplicates(Expression)	To add an expectation that no duplicate messages are received; using an Expression to calculate a unique identifier for each message. This could be something like the JMSMessageID if using JMS, or some unique reference number within the message.

Here's another example:

```
resultEndpoint.expectedBodiesReceived("firstMessageBody", "secondMessageBody",
"thirdMessageBody");
```

34.8. ADDING EXPECTATIONS TO SPECIFIC MESSAGES

In addition, you can use the `message(int messageIndex)` method to add assertions about a specific message that is received.

For example, to add expectations of the headers or body of the first message (using zero-based indexing like `java.util.List`), you can use the following code:

```
resultEndpoint.message(0).header("foo").isEqualTo("bar");
```

There are some examples of the Mock endpoint in use in the [camel-core processor tests](#).

34.9. MOCKING EXISTING ENDPOINTS

Camel now allows you to automatically mock existing endpoints in your Camel routes.



NOTE

How it works

The endpoints are still in action. What happens differently is that a `Mock` endpoint is injected and receives the message first and then delegates the message to the target endpoint. You can view this as a kind of intercept and delegate or endpoint listener.

Suppose you have the given route below:

Route

```
@Override
protected RouteBuilder createRouteBuilder() throws Exception {
    return new RouteBuilder() {
        @Override
        public void configure() throws Exception {
            from("direct:start").routeId("start")
                .to("direct:foo").to("log:foo").to("mock:result");

            from("direct:foo").routeId("foo")
                .transform(constant("Bye World"));
        }
    };
}
```

You can then use the `adviceWith` feature in Camel to mock all the endpoints in a given route from your unit test, as shown below:

`adviceWith` mocking all endpoints

```
@Test
public void testAdvisedMockEndpoints() throws Exception {
    // advice the start route using the inlined AdviceWith lambda style route builder
    // which has extended capabilities than the regular route builder
    AdviceWith.adviceWith(context, "start", a ->
        // mock all endpoints
```

```

a.mockEndpoints();

getMockEndpoint("mock:direct:start").expectedBodiesReceived("Hello World");
getMockEndpoint("mock:direct:foo").expectedBodiesReceived("Hello World");
getMockEndpoint("mock:log:foo").expectedBodiesReceived("Bye World");
getMockEndpoint("mock:result").expectedBodiesReceived("Bye World");

template.sendBody("direct:start", "Hello World");

assertMockEndpointsSatisfied();

// additional test to ensure correct endpoints in registry
assertNotNull(context.hasEndpoint("direct:start"));
assertNotNull(context.hasEndpoint("direct:foo"));
assertNotNull(context.hasEndpoint("log:foo"));
assertNotNull(context.hasEndpoint("mock:result"));
// all the endpoints was mocked
assertNotNull(context.hasEndpoint("mock:direct:start"));
assertNotNull(context.hasEndpoint("mock:direct:foo"));
assertNotNull(context.hasEndpoint("mock:log:foo"));
}

```

Notice that the mock endpoints is given the URI **mock:<endpoint>**, for example **mock:direct:foo**. Camel logs at **INFO** level the endpoints being mocked:

```
INFO Advised endpoint [direct://foo] with mock endpoint [mock:direct:foo]
```



NOTE

Mocked endpoints are without parameters

Endpoints which are mocked will have their parameters stripped off. For example the endpoint **log:foo?showAll=true** will be mocked to the following endpoint **mock:log:foo**. Notice the parameters have been removed.

Its also possible to only mock certain endpoints using a pattern. For example to mock all **log** endpoints you do as shown:

adviceWith mocking only log endpoints using a pattern

```

@Test
public void testAdvisedMockEndpointsWithPattern() throws Exception {
    // advice the start route using the inlined AdviceWith lambda style route builder
    // which has extended capabilities than the regular route builder
    AdviceWith.adviceWith(context, "start", a ->
        // mock only log endpoints
        a.mockEndpoints("log*"));

    // now we can refer to log:foo as a mock and set our expectations
    getMockEndpoint("mock:log:foo").expectedBodiesReceived("Bye World");

    getMockEndpoint("mock:result").expectedBodiesReceived("Bye World");

    template.sendBody("direct:start", "Hello World");
}

```


```

assertMockEndpointsSatisfied();

// additional test to ensure correct endpoints in registry
assertNotNull(context.hasEndpoint("direct:start"));
assertNotNull(context.hasEndpoint("direct:foo"));
assertNotNull(context.hasEndpoint("log:foo"));
assertNotNull(context.hasEndpoint("mock:result"));
// only the log:foo endpoint was mocked
assertNotNull(context.hasEndpoint("mock:log:foo"));
assertNull(context.hasEndpoint("mock:direct:start"));
assertNull(context.hasEndpoint("mock:direct:foo"));
}

```

The pattern supported can be a wildcard or a regular expression. See more details about this at Intercept as its the same matching function used by Camel.



NOTE

Mind that mocking endpoints causes the messages to be copied when they arrive on the mock. That means Camel will use more memory. This may not be suitable when you send in a lot of messages.

34.10. MOCKING EXISTING ENDPOINTS USING THE CAMEL-TEST COMPONENT

Instead of using the **adviceWith** to instruct Camel to mock endpoints, you can easily enable this behavior when using the **camel-test** Test Kit.

The same route can be tested as follows. Notice that we return **""** from the **isMockEndpoints** method, which tells Camel to mock all endpoints.

If you only want to mock all **log** endpoints you can return **"log*"** instead.

isMockEndpoints using camel-test kit

```

public class IsMockEndpointsJUnit4Test extends CamelTestSupport {

    @Override
    public String isMockEndpoints() {
        // override this method and return the pattern for which endpoints to mock.
        // use * to indicate all
        return "";
    }

    @Test
    public void testMockAllEndpoints() throws Exception {
        // notice we have automatic mocked all endpoints and the name of the endpoints is "mock:uri"
        getMockEndpoint("mock:direct:start").expectedBodiesReceived("Hello World");
        getMockEndpoint("mock:direct:foo").expectedBodiesReceived("Hello World");
        getMockEndpoint("mock:log:foo").expectedBodiesReceived("Bye World");
        getMockEndpoint("mock:result").expectedBodiesReceived("Bye World");

        template.sendBody("direct:start", "Hello World");
    }
}

```

```

assertMockEndpointsSatisfied();

// additional test to ensure correct endpoints in registry
assertNotNull(context.hasEndpoint("direct:start"));
assertNotNull(context.hasEndpoint("direct:foo"));
assertNotNull(context.hasEndpoint("log:foo"));
assertNotNull(context.hasEndpoint("mock:result"));
// all the endpoints was mocked
assertNotNull(context.hasEndpoint("mock:direct:start"));
assertNotNull(context.hasEndpoint("mock:direct:foo"));
assertNotNull(context.hasEndpoint("mock:log:foo"));
}

@Override
protected RouteBuilder createRouteBuilder() throws Exception {
    return new RouteBuilder() {
        @Override
        public void configure() throws Exception {
            from("direct:start").to("direct:foo").to("log:foo").to("mock:result");

            from("direct:foo").transform(constant("Bye World"));
        }
    };
}
}
}

```

34.11. MOCKING EXISTING ENDPOINTS WITH XML DSL

If you do not use the **camel-test** component for unit testing (as shown above) you can use a different approach when using XML files for routes.

The solution is to create a new XML file used by the unit test and then include the intended XML file which has the route you want to test.

Suppose we have the route in the **camel-route.xml** file:

camel-route.xml

```

<!-- this camel route is in the camel-route.xml file -->
<camelContext xmlns="http://camel.apache.org/schema/spring">

    <route>
        <from uri="direct:start"/>
        <to uri="direct:foo"/>
        <to uri="log:foo"/>
        <to uri="mock:result"/>
    </route>

    <route>
        <from uri="direct:foo"/>
        <transform>
            <constant>Bye World</constant>
        </transform>
    </route>

</camelContext>

```


Then we create a new XML file as follows, where we include the **camel-route.xml** file and define a spring bean with the class **org.apache.camel.impl.InterceptSendToMockEndpointStrategy** which tells Camel to mock all endpoints:

test-camel-route.xml

```
<!-- the Camel route is defined in another XML file -->
<import resource="camel-route.xml"/>

<!-- bean which enables mocking all endpoints -->
<bean id="mockAllEndpoints"
class="org.apache.camel.component.mock.InterceptSendToMockEndpointStrategy"/>
```

Then in your unit test you load the new XML file (**test-camel-route.xml**) instead of **camel-route.xml**.

To only mock all **Log** endpoints you can define the pattern in the constructor for the bean:

```
<bean id="mockAllEndpoints"
class="org.apache.camel.impl.InterceptSendToMockEndpointStrategy">
  <constructor-arg index="0" value="log*" />
</bean>
```

34.12. MOCKING ENDPOINTS AND SKIP SENDING TO ORIGINAL ENDPOINT

Sometimes you want to easily mock and skip sending to a certain endpoints. So the message is detoured and send to the mock endpoint only. You can now use the **mockEndpointsAndSkip** method using **AdviceWith**. The example below will skip sending to the two endpoints **"direct:foo"**, and **"direct:bar"**.

adviceWith mock and skip sending to endpoints

```
@Test
public void testAdvisedMockEndpointsWithSkip() throws Exception {
  // advice the first route using the inlined AdviceWith route builder
  // which has extended capabilities than the regular route builder
  AdviceWith.adviceWith(context.getRouteDefinitions().get(0), context, new
AdviceWithRouteBuilder() {
    @Override
    public void configure() throws Exception {
      // mock sending to direct:foo and direct:bar and skip send to it
      mockEndpointsAndSkip("direct:foo", "direct:bar");
    }
  });

  getMockEndpoint("mock:result").expectedBodiesReceived("Hello World");
  getMockEndpoint("mock:direct:foo").expectedMessageCount(1);
  getMockEndpoint("mock:direct:bar").expectedMessageCount(1);

  template.sendBody("direct:start", "Hello World");

  assertMockEndpointsSatisfied();

  // the message was not send to the direct:foo route and thus not sent to
  // the seda endpoint
```

```

    SedaEndpoint seda = context.getEndpoint("seda:foo", SedaEndpoint.class);
    assertEquals(0, seda.getCurrentQueueSize());
}

```

The same example using the Test Kit

isMockEndpointsAndSkip using camel-test kit

```

public class IsMockEndpointsAndSkipJUnit4Test extends CamelTestSupport {

    @Override
    public String isMockEndpointsAndSkip() {
        // override this method and return the pattern for which endpoints to mock,
        // and skip sending to the original endpoint.
        return "direct:foo";
    }

    @Test
    public void testMockEndpointAndSkip() throws Exception {
        // notice we have automatic mocked the direct:foo endpoints and the name of the endpoints is
        "mock:uri"
        getMockEndpoint("mock:result").expectedBodiesReceived("Hello World");
        getMockEndpoint("mock:direct:foo").expectedMessageCount(1);

        template.sendBody("direct:start", "Hello World");

        assertMockEndpointsSatisfied();

        // the message was not send to the direct:foo route and thus not sent to the seda endpoint
        SedaEndpoint seda = context.getEndpoint("seda:foo", SedaEndpoint.class);
        assertEquals(0, seda.getCurrentQueueSize());
    }

    @Override
    protected RouteBuilder createRouteBuilder() throws Exception {
        return new RouteBuilder() {
            @Override
            public void configure() throws Exception {
                from("direct:start").to("direct:foo").to("mock:result");

                from("direct:foo").transform(constant("Bye World")).to("seda:foo");
            }
        };
    }
}

```

34.13. LIMITING THE NUMBER OF MESSAGES TO KEEP

The [Mock](#) endpoints will by default keep a copy of every Exchange that it received. So if you test with a lot of messages, then it will consume memory.

We have introduced two options **retainFirst** and **retainLast** that can be used to specify to only keep N'th of the first and/or last Exchanges.

For example in the code below, we only want to retain a copy of the first 5 and last 5 Exchanges the mock receives.

```
MockEndpoint mock = getMockEndpoint("mock:data");
mock.setRetainFirst(5);
mock.setRetainLast(5);
mock.expectedMessageCount(2000);

mock.assertIsSatisfied();
```

Using this has some limitations. The **getExchanges()** and **getReceivedExchanges()** methods on the **MockEndpoint** will return only the retained copies of the Exchanges. So in the example above, the list will contain 10 Exchanges; the first five, and the last five.

The **retainFirst** and **retainLast** options also have limitations on which expectation methods you can use. For example the **expectedXXX** methods that work on message bodies, headers, etc. will only operate on the retained messages. In the example above they can test only the expectations on the 10 retained messages.

34.14. TESTING WITH ARRIVAL TIMES

The **Mock** endpoint stores the arrival time of the message as a property on the Exchange

```
Date time = exchange.getProperty(Exchange.RECEIVED_TIMESTAMP, Date.class);
```

You can use this information to know when the message arrived on the mock. But it also provides foundation to know the time interval between the previous and next message arrived on the mock. You can use this to set expectations using the **arrives** DSL on the Mock endpoint.

For example to say that the first message should arrive between 0-2 seconds before the next you can do:

```
mock.message(0).arrives().noLaterThan(2).seconds().beforeNext();
```

You can also define this as that 2nd message (0 index based) should arrive no later than 0-2 seconds after the previous:


```
mock.message(1).arrives().noLaterThan(2).seconds().afterPrevious();
```

You can also use **between** to set a lower bound. For example suppose that it should be between 1-4 seconds:

```
mock.message(1).arrives().between(1, 4).seconds().afterPrevious();
```

You can also set the expectation on all messages, for example to say that the gap between them should be at most 1 second:

```
mock.allMessages().arrives().noLaterThan(1).seconds().beforeNext();
```



NOTE

Time units

In the example above we use **seconds** as the time unit, but Camel offers **milliseconds**, and **minutes** as well.

34.15. SPRING BOOT AUTO-CONFIGURATION

When using mock with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-mock-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 5 options, which are listed below.

Name	Description	Default	Type
camel.component.mock.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.mock.enabled	Whether to enable auto configuration of the mock component. This is enabled by default.		Boolean
camel.component.mock.exchange-formatter	Sets a custom ExchangeFormatter to convert the Exchange to a String suitable for logging. If not specified, we default to DefaultExchangeFormatter. The option is a <code>org.apache.camel.spi.ExchangeFormatter</code> type.		ExchangeFormatter
camel.component.mock.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.mock.log	To turn on logging when the mock receives an incoming message. This will log only one time at INFO level for the incoming message. For more detailed logging then set the logger to DEBUG level for the <code>org.apache.camel.component.mock.MockEndpoint</code> class.	false	Boolean

CHAPTER 35. MONGODB

Both producer and consumer are supported

According to Wikipedia: "NoSQL is a movement promoting a loosely defined class of non-relational data stores that break with a long history of relational databases and ACID guarantees." NoSQL solutions have grown in popularity in the last few years, and major extremely-used sites and services such as Facebook, LinkedIn, Twitter, etc. are known to use them extensively to achieve scalability and agility.

Basically, NoSQL solutions differ from traditional RDBMS (Relational Database Management Systems) in that they don't use SQL as their query language and generally don't offer ACID-like transactional behaviour nor relational data. Instead, they are designed around the concept of flexible data structures and schemas (meaning that the traditional concept of a database table with a fixed schema is dropped), extreme scalability on commodity hardware and blazing-fast processing.

MongoDB is a very popular NoSQL solution and the camel-mongodb component integrates Camel with MongoDB allowing you to interact with MongoDB collections both as a producer (performing operations on the collection) and as a consumer (consuming documents from a MongoDB collection).

MongoDB revolves around the concepts of documents (not as is office documents, but rather hierarchical data defined in JSON/BSON) and collections. This component page will assume you are familiar with them. Otherwise, visit <http://www.mongodb.org/>.



NOTE

The MongoDB Camel component uses Mongo Java Driver 4.x.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-mongodb</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

35.1. URI FORMAT

```
mongodb:connectionBean?
database=databaseName&collection=collectionName&operation=operationName[&moreOptions...]
```

35.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

35.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

35.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

35.3. COMPONENT OPTIONS

The MongoDB component supports 4 options, which are listed below.

Name	Description	Default	Type
mongoConnection (common)	Autowired Shared client used for connection. All endpoints generated from the component will share this connection client.		MongoClient
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

35.4. ENDPOINT OPTIONS

The MongoDB endpoint is configured using URI syntax:

```
mongodb:connectionBean
```

with the following path and query parameters:

35.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
connectionBean (common)	Required Sets the connection bean reference used to lookup a client for connecting to a database.		String

35.4.2. Query Parameters (27 parameters)

Name	Description	Default	Type
collection (common)	Sets the name of the MongoDB collection to bind to this endpoint.		String
collectionIndex (common)	Sets the collection index (JSON FORMAT : <code>\\{ field1 : order1, field2 : order2}</code>).		String

Name	Description	Default	Type
createCollection (common)	Create collection during initialisation if it doesn't exist. Default is true.	true	boolean
database (common)	Sets the name of the MongoDB database to target.		String
hosts (common)	Host address of mongodb server in host:port format. It's possible also use more than one address, as comma separated list of hosts: host1:port1,host2:port2. If hosts parameter is specified, provided connectionBean is ignored.		String
mongoConnection (common)	Sets the connection bean used as a client for connecting to a database.		MongoClient
operation (common)	<p>Sets the operation this endpoint will execute against MongoDB.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● findById ● findOneByQuery ● findAll ● findDistinct ● insert ● save ● update ● remove ● bulkWrite ● aggregate ● getDbStats ● getColStats ● count ● command 		MongoDbOperation

Name	Description	Default	Type
outputType (common)	<p>Convert the output of the producer to the selected type : DocumentList Document or Mongolterable. DocumentList or Mongolterable applies to findAll and aggregate. Document applies to all other operations.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • DocumentList • Document • Mongolterable 		MongoDbOutputType
bridgeErrorHandler (consumer)	<p>Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.</p>	false	boolean
consumerType (consumer)	Consumer type.		String
exceptionHandler (consumer (advanced))	<p>To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.</p>		ExceptionHandler
exchangePattern (consumer (advanced))	<p>Sets the exchange pattern when the consumer creates an exchange.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 		ExchangePattern

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
cursorRegenerationDelay (advanced)	MongoDB tailable cursors will block until new data arrives. If no new data is inserted, after some time the cursor will be automatically freed and closed by the MongoDB server. The client is expected to regenerate the cursor if needed. This value specifies the time to wait before attempting to fetch a new cursor, and if the attempt fails, how long before the next attempt is made. Default value is 1000ms.	1000	long
dynamicity (advanced)	Sets whether this endpoint will attempt to dynamically resolve the target database and collection from the incoming Exchange properties. Can be used to override at runtime the database and collection specified on the otherwise static endpoint URI. It is disabled by default to boost performance. Enabling it will take a minimal performance hit.	false	boolean
readPreference (advanced)	Configure how MongoDB clients route read operations to the members of a replica set. Possible values are PRIMARY, PRIMARY_PREFERRED, SECONDARY, SECONDARY_PREFERRED or NEAREST. Enum values: <ul style="list-style-type: none"> ● PRIMARY ● PRIMARY_PREFERRED ● SECONDARY ● SECONDARY_PREFERRED ● NEAREST 	PRIMARY	String

Name	Description	Default	Type
writeConcern (advanced)	<p>Configure the connection bean with the level of acknowledgment requested from MongoDB for write operations to a standalone mongod, replicaset or cluster. Possible values are ACKNOWLEDGED, W1, W2, W3, UNACKNOWLEDGED, JOURNALED or MAJORITY.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● ACKNOWLEDGED ● W1 ● W2 ● W3 ● UNACKNOWLEDGED ● JOURNALED ● MAJORITY 	ACKNOWLEDGED	String
writeResultAsHeader (advanced)	In write operations, it determines whether instead of returning WriteResult as the body of the OUT message, we transfer the IN message to the OUT and attach the WriteResult as a header.	false	boolean
streamFilter (changeStream)	Filter condition for change streams consumer.		String
password (security)	User password for mongodb connection.		String
username (security)	Username for mongodb connection.		String
persistentId (tail)	One tail tracking collection can host many trackers for several tailable consumers. To keep them separate, each tracker should have its own unique persistentId.		String
persistentTailTracking (tail)	Enable persistent tail tracking, which is a mechanism to keep track of the last consumed message across system restarts. The next time the system is up, the endpoint will recover the cursor from the point where it last stopped slurping records.	false	boolean

Name	Description	Default	Type
tailTrackCollection (tail)	Collection where tail tracking information will be persisted. If not specified, <code>MongoDbTailTrackingConfig#DEFAULT_COLLECTION</code> will be used by default.		String
tailTrackDb (tail)	Indicates what database the tail tracking mechanism will persist to. If not specified, the current database will be picked by default. Dynamicity will not be taken into account even if enabled, i.e. the tail tracking database will not vary past endpoint initialisation.		String
tailTrackField (tail)	Field where the last tracked value will be placed. If not specified, <code>MongoDbTailTrackingConfig#DEFAULT_FIELD</code> will be used by default.		String
tailTrackIncreasingField (tail)	Correlation field in the incoming record which is of increasing nature and will be used to position the tailing cursor every time it is generated. The cursor will be (re)created with a query of type: <code>tailTrackIncreasingField greater than lastValue</code> (possibly recovered from persistent tail tracking). Can be of type Integer, Date, String, etc. NOTE: No support for dot notation at the current time, so the field should be at the top level of the document.		String

35.5. CONFIGURATION OF DATABASE IN SPRING XML

The following Spring XML creates a bean defining the connection to a MongoDB instance.

Since mongo java driver 3, the `WriteConcern` and `readPreference` options are not dynamically modifiable. They are defined in the `mongoClient` object

```
<beans xmlns="http://www.springframework.org/schema/beans"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:context="http://www.springframework.org/schema/context"
xmlns:mongo="http://www.springframework.org/schema/data/mongo"
xsi:schemaLocation="http://www.springframework.org/schema/context
http://www.springframework.org/schema/context/spring-context.xsd
http://www.springframework.org/schema/data/mongo
http://www.springframework.org/schema/data/mongo/spring-mongo.xsd
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">

  <mongo:mongo-client id="mongoBean" host="${mongo.url}" port="${mongo.port}"
credentials="${mongo.user}:${mongo.pass}@${mongo.dbname}">
    <mongo:client-options write-concern="NORMAL" />
  </mongo:mongo-client>
</beans>
```

35.6. SAMPLE ROUTE

The following route defined in Spring XML executes the operation **getDbStats** on a collection.

Get DB stats for specified collection

```
<route>
  <from uri="direct:start" />
  <!-- using bean 'mongoBean' defined above -->
  <to uri="mongodb:mongoBean?
database=${mongodb.database}&collection=${mongodb.collection}&operation=getDbStats"
/>
  <to uri="direct:result" />
</route>
```

35.7. MONGODB OPERATIONS - PRODUCER ENDPOINTS

35.7.1. Query operations

35.7.1.1. findById

This operation retrieves only one element from the collection whose `_id` field matches the content of the IN message body. The incoming object can be anything that has an equivalent to a **Bson** type. See <http://bsonspec.org/spec.html> and <http://www.mongodb.org/display/DOCS/Java+Types>.

```
from("direct:findById")
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findById")
  .to("mock:resultFindById");
```

Please, note that the default `_id` is treated by Mongo as and **ObjectId** type, so you may need to convert it properly.

```
from("direct:findById")
  .convertBodyTo(ObjectId.class)
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findById")
  .to("mock:resultFindById");
```



NOTE

Supports optional parameters

This operation supports projection operators. See [Specifying a fields filter \(projection\)](#).

35.7.1.2. findOneByQuery

Retrieve the first element from a collection matching a MongoDB query selector. **If the `CamelMongoDbCriteria` header is set, then its value is used as the query selector** If the **`CamelMongoDbCriteria`** header is *null*, then the IN message body is used as the query selector. In both cases, the query selector should be of type **Bson** or convertible to **Bson** (for instance, a JSON string or **HashMap**). See Type conversions for more info.

Create query selectors using the **Filters** provided by the MongoDB Driver.

35.7.1.3. Example without a query selector (returns the first document in a collection)

```
from("direct:findOneByQuery")
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findOneByQuery")
  .to("mock:resultFindOneByQuery");
```

35.7.1.4. Example with a query selector (returns the first matching document in a collection):

```
from("direct:findOneByQuery")
  .setHeader(MongoDbConstants.CRITERIA, constant(Filters.eq("name", "Raul Kripalani")))
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findOneByQuery")
  .to("mock:resultFindOneByQuery");
```



NOTE

Supports optional parameters

This operation supports projection operators and sort clauses. See [Specifying a fields filter \(projection\)](#), [Specifying a sort clause](#).

35.7.1.5. findAll

The **findAll** operation returns all documents matching a query, or none at all, in which case all documents contained in the collection are returned. **The query object is extracted CamelMongoDbCriteria header.** if the CamelMongoDbCriteria header is null the query object is extracted message body, i.e. it should be of type **Bson** or convertible to **Bson**. It can be a JSON String or a Hashmap. See [Type conversions](#) for more info.

35.7.1.5.1. Example without a query selector (returns all documents in a collection)

```
from("direct:findAll")
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findAll")
  .to("mock:resultFindAll");
```

35.7.1.5.2. Example with a query selector (returns all matching documents in a collection)

```
from("direct:findAll")
  .setHeader(MongoDbConstants.CRITERIA, Filters.eq("name", "Raul Kripalani"))
  .to("mongodb:myDb?database=flights&collection=tickets&operation=findAll")
  .to("mock:resultFindAll");
```

Paging and efficient retrieval is supported via the following headers:

Header key	Quick constant	Description (extracted from MongoDB API doc)	Expected type
CamelMongoDbNumToSkip	MongoDbConstants.NUM_TO_SKIP	Discards a given number of elements at the beginning of the cursor.	int/Integer

Header key	Quick constant	Description (extracted from MongoDB API doc)	Expected type
CamelMongoDbLimit	MongoDbConstants.LIMIT	Limits the number of elements returned.	int/Integer
CamelMongoDbBatchSize	MongoDbConstants.BATCH_SIZE	Limits the number of elements returned in one batch. A cursor typically fetches a batch of result objects and store them locally. If batchSize is positive, it represents the size of each batch of objects retrieved. It can be adjusted to optimize performance and limit data transfer. If batchSize is negative, it will limit of number objects returned, that fit within the max batch size limit (usually 4MB), and cursor will be closed. For example if batchSize is -10, then the server will return a maximum of 10 documents and as many as can fit in 4MB, then close the cursor. Note that this feature is different from limit() in that documents must fit within a maximum size, and it removes the need to send a request to close the cursor server-side. The batch size can be changed even after a cursor is iterated, in which case the setting will apply on the next batch retrieval.	int/Integer
CamelMongoDbAllowDiskUse	MongoDbConstants.ALLOW_DISK_USE	Sets allowDiskUse MongoDB flag. This is supported since MongoDB Server 4.3.1. Using this header with older MongoDB Server version can cause query to fail.	boolean/Boolean

35.7.1.5.3. Example with option *outputType=Mongolterable* and batch size


```

from("direct:findAll")
  .setHeader(MongoDbConstants.BATCH_SIZE).constant(10)
  .setHeader(MongoDbConstants.CRITERIA, constant(Filters.eq("name", "Raul Kripalani")))
  .to("mongodb:myDb?
database=flights&collection=tickets&operation=findAll&outputType=Mongolterable")
  .to("mock:resultFindAll");

```

The **findAll** operation will also return the following OUT headers to enable you to iterate through result pages if you are using paging:

Header key	Quick constant	Description (extracted from MongoDB API doc)	Data type
CamelMongoDbResultTotalSize	MongoDbConstants.RESULT_TOTAL_SIZE	Number of objects matching the query. This does not take limit/skip into consideration.	int/Integer
CamelMongoDbResultPageSize	MongoDbConstants.RESULT_PAGE_SIZE	Number of objects matching the query. This does not take limit/skip into consideration.	int/Integer



NOTE

Supports optional parameters

This operation supports projection operators and sort clauses. See [Specifying a fields filter \(projection\)](#), [Specifying a sort clause](#).

35.7.1.6. count

Returns the total number of objects in a collection, returning a Long as the OUT message body. The following example will count the number of records in the "dynamicCollectionName" collection. Notice how dynamicity is enabled, and as a result, the operation will not run against the "notableScientists" collection, but against the "dynamicCollectionName" collection.

```
// from("direct:count").to("mongodb:myDb?
database=tickets&collection=flights&operation=count&dynamicity=true");
Long result = template.requestBodyAndHeader("direct:count", "irrelevantBody",
MongoDbConstants.COLLECTION, "dynamicCollectionName");
assertTrue("Result is not of type Long", result instanceof Long);
```

You can provide a query **The query object is extracted CamelMongoDbCriteria header**. if the CamelMongoDbCriteria header is null the query object is extracted message body, i.e. it should be of type **Bson** or convertible to **Bson**, and operation will return the amount of documents matching this criteria.

```
Document query = ...
Long count = template.requestBodyAndHeader("direct:count", query,
MongoDbConstants.COLLECTION, "dynamicCollectionName");
```

35.7.1.7. Specifying a fields filter (projection)

Query operations will, by default, return the matching objects in their entirety (with all their fields). If your documents are large and you only require retrieving a subset of their fields, you can specify a field filter in all query operations, simply by setting the relevant **Bson** (or type convertible to **Bson**, such as a JSON String, Map, etc.) on the **CamelMongoDbFieldsProjection** header, constant shortcut: **MongoDbConstants.FIELDS_PROJECTION**.

Here is an example that uses MongoDB's **Projections** to simplify the creation of Bson. It retrieves all fields except **_id** and **boringField**:

```
// route: from("direct:findAll").to("mongodb:myDb?
database=flights&collection=tickets&operation=findAll")
Bson fieldProjection = Projection.exclude("_id", "boringField");
Object result = template.requestBodyAndHeader("direct:findAll", ObjectUtils.NULL,
MongoDbConstants.FIELDS_PROJECTION, fieldProjection);
```

Here is an example that uses MongoDB's **Projections** to simplify the creation of Bson. It retrieves all fields except **_id** and **boringField**:

```
// route: from("direct:findAll").to("mongodb:myDb?
database=flights&collection=tickets&operation=findAll")
Bson fieldProjection = Projection.exclude("_id", "boringField");
Object result = template.requestBodyAndHeader("direct:findAll", ObjectUtils.NULL,
MongoDbConstants.FIELDS_PROJECTION, fieldProjection);
```

35.7.1.8. Specifying a sort clause

There is often a requirement to fetch the min/max record from a collection based on sorting by a particular field that uses MongoDB's **Sorts** to simplify the creation of Bson. It retrieves all fields except **_id** and **boringField**:

```
// route: from("direct:findAll").to("mongodb:myDb?
database=flights&collection=tickets&operation=findAll")
Bson sorts = Sorts.descending("_id");
Object result = template.requestBodyAndHeader("direct:findAll", ObjectUtils.NULL,
MongoDbConstants.SORT_BY, sorts);
```

In a Camel route the SORT_BY header can be used with the `findOneByQuery` operation to achieve the same result. If the FIELDS_PROJECTION header is also specified the operation will return a single field/value pair that can be passed directly to another component (for example, a parameterized MyBatis SELECT query). This example demonstrates fetching the temporally newest document from a collection and reducing the result to a single field, based on the **documentTimestamp** field:

```
.from("direct:someTriggeringEvent")
.setHeader(MongoDbConstants.SORT_BY).constant(Sorts.descending("documentTimestamp"))
.setHeader(MongoDbConstants.FIELDS_PROJECTION).constant(Projection.include("documentTime
stamp"))
.setBody().constant("{}")
.to("mongodb:myDb?database=local&collection=myDemoCollection&operation=findOneByQuery")
.to("direct:aMyBatisParameterizedSelect");
```

35.7.2. Create/update operations

35.7.2.1. insert

Inserts a new object into the MongoDB collection, taken from the IN message body. Type conversion is attempted to turn it into **Document** or a **List**.

Two modes are supported: single insert and multiple insert. For multiple insert, the endpoint will expect a List, Array or Collections of objects of any type, as long as they are - or can be converted to - **Document**. Example:

```
from("direct:insert")
  .to("mongodb:myDb?database=flights&collection=tickets&operation=insert");
```

The operation will return a **WriteResult**, and depending on the **WriteConcern** or the value of the **invokeGetLastError** option, **getLastError()** would have been called already or not. If you want to access the ultimate result of the write operation, you need to retrieve the **CommandResult** by calling **getLastError()** or **getCachedLastError()** on the **WriteResult**. Then you can verify the result by calling **CommandResult.ok()**, **CommandResult.getErrorMessage()** and/or **CommandResult.getException()**.

Note that the new object's **_id** must be unique in the collection. If you don't specify the value, MongoDB will automatically generate one for you. But if you do specify it and it is not unique, the insert operation will fail (and for Camel to notice, you will need to enable **invokeGetLastError** or set a **WriteConcern** that waits for the write result).

This is not a limitation of the component, but it is how things work in MongoDB for higher throughput. If you are using a custom **_id**, you are expected to ensure at the application level that it is unique (and this is a good practice too).

OID(s) of the inserted record(s) is stored in the message header under **CamelMongoOid** key (**MongoDbConstants.OID** constant). The value stored is **org.bson.types.ObjectId** for single insert or **java.util.List<org.bson.types.ObjectId>** if multiple records have been inserted.

In MongoDB Java Driver 3.x the **insertOne** and **insertMany** operation return void. The Camel insert operation return the Document or List of Documents inserted. Note that each Documents are Updated by a new OID if need.

35.7.2.2. save

The save operation is equivalent to an *upsert* (UPdate, inSERT) operation, where the record will be updated, and if it doesn't exist, it will be inserted, all in one atomic operation. MongoDB will perform the matching based on the **_id** field.

Beware that in case of an update, the object is replaced entirely and the usage of [MongoDB's \\$modifiers](#) is not permitted. Therefore, if you want to manipulate the object if it already exists, you have two options:

1. perform a query to retrieve the entire object first along with all its fields (may not be efficient), alter it inside Camel and then save it.
2. use the update operation with [\\$modifiers](#), which will execute the update at the server-side instead. You can enable the upsert flag, in which case if an insert is required, MongoDB will apply the [\\$modifiers](#) to the filter query object and insert the result.

If the document to be saved does not contain the **_id** attribute, the operation will be an insert, and the new **_id** created will be placed in the **CamelMongoOid** header.

For example:


```
from("direct:insert")
  .to("mongodb:myDb?database=flights&collection=tickets&operation=save");
```

```
// route: from("direct:insert").to("mongodb:myDb?
database=flights&collection=tickets&operation=save");
org.bson.Document docForSave = new org.bson.Document();
docForSave.put("key", "value");
Object result = template.requestBody("direct:insert", docForSave);
```

35.7.2.3. update

Update one or multiple records on the collection. Requires a filter query and a update rules.

You can define the filter using `MongoDBConstants.CRITERIA` header as **Bson** and define the update rules as **Bson** in Body.



NOTE

Update after enrich

While defining the filter by using `MongoDBConstants.CRITERIA` header as **Bson** to query mongodb before you do update, you should notice you need to remove it from the resulting camel exchange during aggregation if you use enrich pattern with a aggregation strategy and then apply mongodb update. If you don't remove this header during aggregation and/or redefine `MongoDBConstants.CRITERIA` header before sending camel exchange to mongodb producer endpoint, you may end up with invalid camel exchange payload while updating mongodb.

The second way Require a `List<Bson>` as the IN message body containing exactly 2 elements:

- Element 1 (index 0) ⇒ filter query ⇒ determines what objects will be affected, same as a typical query object
- Element 2 (index 1) ⇒ update rules ⇒ how matched objects will be updated. All [modifier operations](#) from MongoDB are supported.



NOTE

Multiupdates

By default, MongoDB will only update 1 object even if multiple objects match the filter query. To instruct MongoDB to update **all** matching records, set the **CamelMongoDbMultiUpdate** IN message header to **true**.

A header with key **CamelMongoDbRecordsAffected** will be returned (**MongoDbConstants.RECORDS_AFFECTED** constant) with the number of records updated (copied from **WriteResult.getN()**).

Supports the following IN message headers:

Header key	Quick constant	Description (extracted from MongoDB API doc)	Expected type
CamelMongoDb MultiUpdate	MongoDbConstants.MULTIUPDATE	If the update should be applied to all objects matching. See http://www.mongodb.org/display/DOCS/Atomic+Operations	boolean/Boolean
CamelMongoDb Upsert	MongoDbConstants.UPSERT	If the database should create the element if it does not exist	boolean/Boolean

For example, the following will update **all** records whose `filterField` field equals `true` by setting the value of the `"scientist"` field to `"Darwin"`:

```
// route: from("direct:update").to("mongodb:myDb?
database=science&collection=notableScientists&operation=update");
List<Bson> body = new ArrayList<>();
Bson filterField = Filters.eq("filterField", true);
body.add(filterField);
BsonDocument updateObj = new BsonDocument().append("$set", new BsonDocument("scientist",
new BsonString("Darwin")));
body.add(updateObj);
Object result = template.requestBodyAndHeader("direct:update", body,
MongoDbConstants.MULTIUPDATE, true);
```

```
// route: from("direct:update").to("mongodb:myDb?
database=science&collection=notableScientists&operation=update");
Maps<String, Object> headers = new HashMap<>(2);
headers.add(MongoDbConstants.MULTIUPDATE, true);
headers.add(MongoDbConstants.FIELDS_FILTER, Filters.eq("filterField", true));
String updateObj = Updates.set("scientist", "Darwin");
Object result = template.requestBodyAndHeaders("direct:update", updateObj, headers);
```

```
// route: from("direct:update").to("mongodb:myDb?
database=science&collection=notableScientists&operation=update");
String updateObj = "[{"filterField": true}, {"$set": {"scientist": "Darwin"}}]";
Object result = template.requestBodyAndHeader("direct:update", updateObj,
MongoDbConstants.MULTIUPDATE, true);
```

35.7.3. Delete operations

35.7.3.1. remove

Remove matching records from the collection. The IN message body will act as the removal filter query, and is expected to be of type **DBObject** or a type convertible to it.

The following example will remove all objects whose field `'conditionField'` equals `true`, in the `science` database, `notableScientists` collection:

```
// route: from("direct:remove").to("mongodb:myDb?
database=science&collection=notableScientists&operation=remove");
Bson conditionField = Filters.eq("conditionField", true);
Object result = template.requestBody("direct:remove", conditionField);
```

A header with key **CamelMongoDbRecordsAffected** is returned (**MongoDbConstants.RECORDS_AFFECTED** constant) with type **int**, containing the number of records deleted (copied from **WriteResult.getN()**).

35.7.4. Bulk Write Operations

35.7.4.1. bulkWrite

Performs write operations in bulk with controls for order of execution. Requires a **List<WriteModel<Document>>** as the IN message body containing commands for insert, update, and delete operations.

The following example will insert a new scientist "Pierre Curie", update record with id "5" by setting the value of the "scientist" field to "Marie Curie" and delete record with id "3" :

```
// route: from("direct:bulkWrite").to("mongodb:myDb?
database=science&collection=notableScientists&operation=bulkWrite");
List<WriteModel<Document>> bulkOperations = Arrays.asList(
    new InsertOneModel<>(new Document("scientist", "Pierre Curie")),
    new UpdateOneModel<>(new Document("_id", "5"),
        new Document("$set", new Document("scientist", "Marie Curie"))),
    new DeleteOneModel<>(new Document("_id", "3")));
```

```
BulkWriteResult result = template.requestBody("direct:bulkWrite", bulkOperations,
BulkWriteResult.class);
```

By default, operations are executed in order and interrupted on the first write error without processing any remaining write operations in the list. To instruct MongoDB to continue to process remaining write operations in the list, set the **CamelMongoDbBulkOrdered** IN message header to **false**. Unordered operations are executed in parallel and this behavior is not guaranteed.

Header key	Quick constant	Description (extracted from MongoDB API doc)	Expected type
CamelMongoDbBulkOrdered	MongoDbConstants.BULK_ORDERED	Perform an ordered or unordered operation execution. Defaults to true.	boolean/Boolean

35.7.5. Other operations

35.7.5.1. aggregate

Perform a aggregation with the given pipeline contained in the body. **Aggregations could be long and heavy operations. Use with care.**

```
// route: from("direct:aggregate").to("mongodb:myDb?
```

```

database=science&collection=notableScientists&operation=aggregate");
List<Bson> aggregate = Arrays.asList(match(or(eq("scientist", "Darwin"), eq("scientist",
    group("$scientist", sum("count", 1))));
from("direct:aggregate")
    .setBody().constant(aggregate)
    .to("mongodb:myDb?database=science&collection=notableScientists&operation=aggregate")
    .to("mock:resultAggregate");

```

Supports the following IN message headers:

Header key	Quick constant	Description (extracted from MongoDB API doc)	Expected type
CamelMongoDbBatchSize	MongoDbConstants.BATCH_SIZE	Sets the number of documents to return per batch.	int/Integer
CamelMongoDbAllowDiskUse	MongoDbConstants.ALLOW_DISK_USE	Enable aggregation pipeline stages to write data to temporary files.	boolean/Boolean

By default a List of all results is returned. This can be heavy on memory depending on the size of the results. A safer alternative is to set your `outputType=Mongolterable`. The next Processor will see an iterable in the message body allowing it to step through the results one by one. Thus setting a batch size and returning an iterable allows for efficient retrieval and processing of the result.

An example would look like:

```

List<Bson> aggregate = Arrays.asList(match(or(eq("scientist", "Darwin"), eq("scientist",
    group("$scientist", sum("count", 1))));
from("direct:aggregate")
    .setHeader(MongoDbConstants.BATCH_SIZE).constant(10)
    .setBody().constant(aggregate)
    .to("mongodb:myDb?
database=science&collection=notableScientists&operation=aggregate&outputType=Mongolterable")
    .split(body())
    .streaming()
    .to("mock:resultAggregate");

```

Note that calling `.split(body())` is enough to send the entries down the route one-by-one, however it would still load all the entries into memory first. Calling `.streaming()` is thus required to load data into memory by batches.

35.7.5.2. getDbStats

Equivalent of running the `db.stats()` command in the MongoDB shell, which displays useful statistic figures about the database.

For example:

```

> db.stats();
{
  "db" : "test",

```

```

    "collections" : 7,
    "objects" : 719,
    "avgObjSize" : 59.73296244784423,
    "dataSize" : 42948,
    "storageSize" : 1000058880,
    "numExtents" : 9,
    "indexes" : 4,
    "indexSize" : 32704,
    "fileSize" : 1275068416,
    "nsSizeMB" : 16,
    "ok" : 1
  }

```

Usage example:

```

// from("direct:getDbStats").to("mongodb:myDb?
database=flights&collection=tickets&operation=getDbStats");
Object result = template.requestBody("direct:getDbStats", "irrelevantBody");
assertTrue("Result is not of type Document", result instanceof Document);

```

The operation will return a data structure similar to the one displayed in the shell, in the form of a **Document** in the OUT message body.

35.7.5.3. getColStats

Equivalent of running the **db.collection.stats()** command in the MongoDB shell, which displays useful statistic figures about the collection.

For example:

```

> db.camelTest.stats();
{
  "ns" : "test.camelTest",
  "count" : 100,
  "size" : 5792,
  "avgObjSize" : 57.92,
  "storageSize" : 20480,
  "numExtents" : 2,
  "nindexes" : 1,
  "lastExtentSize" : 16384,
  "paddingFactor" : 1,
  "flags" : 1,
  "totalIndexSize" : 8176,
  "indexSizes" : {
    "_id_" : 8176
  },
  "ok" : 1
}

```

Usage example:

```

// from("direct:getColStats").to("mongodb:myDb?
database=flights&collection=tickets&operation=getColStats");
Object result = template.requestBody("direct:getColStats", "irrelevantBody");
assertTrue("Result is not of type Document", result instanceof Document);

```

The operation will return a data structure similar to the one displayed in the shell, in the form of a **Document** in the OUT message body.

35.7.5.4. command

Run the body as a command on database. Useful for admin operation as getting host information, replication or sharding status.

Collection parameter is not use for this operation.

```
// route: from("command").to("mongodb:myDb?database=science&operation=command");
DBObject commandBody = new BasicDBObject("hostInfo", "1");
Object result = template.requestBody("direct:command", commandBody);
```

35.7.6. Dynamic operations

An Exchange can override the endpoint's fixed operation by setting the **CamelMongoDbOperation** header, defined by the **MongoDbConstants.OPERATION_HEADER** constant.

The values supported are determined by the **MongoDbOperation** enumeration and match the accepted values for the **operation** parameter on the endpoint URI.

For example:

```
// from("direct:insert").to("mongodb:myDb?database=flights&collection=tickets&operation=insert");
Object result = template.requestBodyAndHeader("direct:insert", "irrelevantBody",
MongoDbConstants.OPERATION_HEADER, "count");
assertTrue("Result is not of type Long", result instanceof Long);
```

35.8. CONSUMERS

There are several types of consumers:

1. Tailable Cursor Consumer
2. Change Streams Consumer

35.8.1. Tailable Cursor Consumer

MongoDB offers a mechanism to instantaneously consume ongoing data from a collection, by keeping the cursor open just like the **tail -f** command of *nix systems. This mechanism is significantly more efficient than a scheduled poll, due to the fact that the server pushes new data to the client as it becomes available, rather than making the client ping back at scheduled intervals to fetch new data. It also reduces otherwise redundant network traffic.

There is only one requisite to use tailable cursors: the collection must be a "capped collection", meaning that it will only hold N objects, and when the limit is reached, MongoDB flushes old objects in the same order they were originally inserted. For more information, please refer to <http://www.mongodb.org/display/DOCS/Tailable+Cursors>.

The Camel MongoDB component implements a tailable cursor consumer, making this feature available for you to use in your Camel routes. As new objects are inserted, MongoDB will push them as **Document** in natural order to your tailable cursor consumer, who will transform them to an Exchange and will trigger your route logic.

35.9. HOW THE TAILABLE CURSOR CONSUMER WORKS

To turn a cursor into a tailable cursor, a few special flags are to be signalled to MongoDB when first generating the cursor. Once created, the cursor will then stay open and will block upon calling the **MongoCursor.next()** method until new data arrives. However, the MongoDB server reserves itself the right to kill your cursor if new data doesn't appear after an indeterminate period. If you are interested to continue consuming new data, you have to regenerate the cursor. And to do so, you will have to remember the position where you left off or else you will start consuming from the top again.

The Camel MongoDB tailable cursor consumer takes care of all these tasks for you. You will just need to provide the key to some field in your data of increasing nature, which will act as a marker to position your cursor every time it is regenerated, e.g. a timestamp, a sequential ID, etc. It can be of any datatype supported by MongoDB. Date, Strings and Integers are found to work well. We call this mechanism "tail tracking" in the context of this component.

The consumer will remember the last value of this field and whenever the cursor is to be regenerated, it will run the query with a filter like: **increasingField > lastValue**, so that only unread data is consumed.

Setting the increasing field: Set the key of the increasing field on the endpoint URI **tailTrackingIncreasingField** option. In Camel 2.10, it must be a top-level field in your data, as nested navigation for this field is not yet supported. That is, the "timestamp" field is okay, but "nested.timestamp" will not work. Please open a ticket in the Camel JIRA if you do require support for nested increasing fields.

Cursor regeneration delay: One thing to note is that if new data is not already available upon initialisation, MongoDB will kill the cursor instantly. Since we don't want to overwhelm the server in this case, a **cursorRegenerationDelay** option has been introduced (with a default value of 1000ms.), which you can modify to suit your needs.

An example:

```
from("mongodb:myDb?
database=flights&collection=cancellations&tailTrackIncreasingField=departureTime")
  .id("tailableCursorConsumer1")
  .autoStartup(false)
  .to("mock:test");
```

The above route will consume from the "flights.cancellations" capped collection, using "departureTime" as the increasing field, with a default regeneration cursor delay of 1000ms.

35.10. PERSISTENT TAIL TRACKING

Standard tail tracking is volatile and the last value is only kept in memory. However, in practice you will need to restart your Camel container every now and then, but your last value would then be lost and your tailable cursor consumer would start consuming from the top again, very likely sending duplicate records into your route.

To overcome this situation, you can enable the **persistent tail tracking** feature to keep track of the last consumed increasing value in a special collection inside your MongoDB database too. When the consumer initialises again, it will restore the last tracked value and continue as if nothing happened.

The last read value is persisted on two occasions: every time the cursor is regenerated and when the consumer shuts down. We may consider persisting at regular intervals too in the future (flush every 5 seconds) for added robustness if the demand is there. To request this feature, please open a ticket in the Camel JIRA.

35.11. ENABLING PERSISTENT TAIL TRACKING

To enable this function, set at least the following options on the endpoint URI:

- **persistentTailTracking** option to **true**
- **persistentId** option to a unique identifier for this consumer, so that the same collection can be reused across many consumers

Additionally, you can set the **tailTrackDb**, **tailTrackCollection** and **tailTrackField** options to customise where the runtime information will be stored. Refer to the endpoint options table at the top of this page for descriptions of each option.

For example, the following route will consume from the "flights.cancellations" capped collection, using "departureTime" as the increasing field, with a default regeneration cursor delay of 1000ms, with persistent tail tracking turned on, and persisting under the "cancellationsTracker" id on the "flights.camelTailTracking", storing the last processed value under the "lastTrackingValue" field (**camelTailTracking** and **lastTrackingValue** are defaults).

```
from("mongodb:myDb?
database=flights&collection=cancellations&tailTrackIncreasingField=departureTime&persistentTailTrack
ng=true" +
"&persistentId=cancellationsTracker")
.id("tailableCursorConsumer2")
.autoStartup(false)
.to("mock:test");
```

Below is another example identical to the one above, but where the persistent tail tracking runtime information will be stored under the "trackers.camelTrackers" collection, in the "lastProcessedDepartureTime" field:

```
from("mongodb:myDb?
database=flights&collection=cancellations&tailTrackIncreasingField=departureTime&persistentTailTrack
ng=true" +
"&persistentId=cancellationsTracker&tailTrackDb=trackers&tailTrackCollection=camelTrackers" +
"&tailTrackField=lastProcessedDepartureTime")
.id("tailableCursorConsumer3")
.autoStartup(false)
.to("mock:test");
```

35.11.1. Change Streams Consumer

Change Streams allow applications to access real-time data changes without the complexity and risk of tailing the MongoDB oplog. Applications can use change streams to subscribe to all data changes on a collection and immediately react to them. Because change streams use the aggregation framework, applications can also filter for specific changes or transform the notifications at will. The exchange body will contain the full document of any change.

To configure Change Streams Consumer you need to specify **consumerType**, **database**, **collection** and optional JSON property **streamFilter** to filter events. That JSON property is standard MongoDB **\$match** aggregation. It could be easily specified using XML DSL configuration:

```
<route id="filterConsumer">
  <from uri="mongodb:myDb?
consumerType=changeStreams&database=flights&collection=tickets&streamFilter={
```

```
'$match':{'$or':[['fullDocument.stringValue': 'specificValue']]} }"/>
  <to uri="mock:test"/>
</route>
```

Java configuration:

```
from("mongodb:myDb?
consumerType=changeStreams&database=flights&collection=tickets&streamFilter={ '$match':{'$or':
[['fullDocument.stringValue': 'specificValue']]} }")
.to("mock:test");
```



NOTE

You can externalize the `streamFilter` value into a property placeholder which allows the endpoint URI parameters to be *cleaner* and easier to read.

The **changeStreams** consumer type will also return the following OUT headers:

Header key	Quick constant	Description (extracted from MongoDB API doc)	Data type
CamelMongoDbStreamOperationType	MongoDbConstants.STREAM_OPERATION_TYPE	The type of operation that occurred. Can be any of the following values: insert, delete, replace, update, drop, rename, dropDatabase, invalidate.	String
_id	MongoDbConstants.MONGO_ID	A document that contains the <code>_id</code> of the document created or modified by the insert, replace, delete, update operations (i.e. CRUD operations). For sharded collections, also displays the full shard key for the document. The <code>_id</code> field is not repeated if it is already a part of the shard key.	ObjectId

35.12. TYPE CONVERSIONS

The **MongoDbBasicConverters** type converter included with the camel-mongodb component provides the following conversions:

Name	From type	To type	How?
fromMapToDocument	Map	Document	constructs a new Document via the new Document(Map m) constructor.
fromDocumentToMap	Document	Map	Document already implements Map .

Name	From type	To type	How?
fromStringToDocument	String	Document	uses com.mongodb.Document.parse(String s) .
fromStringToObjectId	String	ObjectId	constructs a new ObjectId via the new ObjectId(s)
fromFileToDocument	File	Document	uses fromInputStreamToDocument under the hood
fromInputStreamToDocument	InputStream	Document	converts the inputstream bytes to a Document
fromStringToList	String	List<Bson>	uses org.bson.codecs.configuration.CodecRegistries to convert to BsonArray then to List<Bson>.

This type converter is auto-discovered, so you don't need to configure anything manually.

35.13. SPRING BOOT AUTO-CONFIGURATION

When using mongodb with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-mongodb-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 5 options, which are listed below.

Name	Description	Default	Type
camel.component.mongodb.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.mongodb.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.mongodb.enabled	Whether to enable auto configuration of the mongodb component. This is enabled by default.		Boolean
camel.component.mongodb.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.mongodb.mongo-connection	Shared client used for connection. All endpoints generated from the component will share this connection client. The option is a <code>com.mongodb.client.MongoClient</code> type.		MongoClient

CHAPTER 36. NETTY

Both producer and consumer are supported

The Netty component in Camel is a socket communication component, based on the [Netty](#) project version 4.

Netty is a NIO client server framework which enables quick and easy development of networkServerInitializerFactory applications such as protocol servers and clients.

Netty greatly simplifies and streamlines network programming such as TCP and UDP socket server.

This camel component supports both producer and consumer endpoints.

The Netty component has several options and allows fine-grained control of a number of TCP/UDP communication parameters (buffer sizes, keepAlives, tcpNoDelay, etc) and facilitates both In-Only and In-Out communication on a Camel route.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-netty</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

36.1. URI FORMAT

The URI scheme for a netty component is as follows

```
netty:tcp://0.0.0.0:99999[?options]
netty:udp://remotehost:99999[?options]
```

This component supports producer and consumer endpoints for both TCP and UDP.

36.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

36.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example, a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (**application.properties|yaml**), or directly with Java code.

36.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a *type safe* way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

36.3. COMPONENT OPTIONS

The Netty component supports 73 options, which are listed below.

Name	Description	Default	Type
configuration (common)	To use the NettyConfiguration as configuration when creating endpoints.		NettyConfiguration
disconnect (common)	Whether or not to disconnect(close) from Netty Channel right after use. Can be used for both consumer and producer.	false	boolean
keepAlive (common)	Setting to ensure socket is not closed due to inactivity.	true	boolean
reuseAddress (common)	Setting to facilitate socket multiplexing.	true	boolean
reuseChannel (common)	This option allows producers and consumers (in client mode) to reuse the same Netty Channel for the lifecycle of processing the Exchange. This is useful if you need to call a server multiple times in a Camel route and want to use the same network connection. When using this, the channel is not returned to the connection pool until the Exchange is done; or disconnected if the disconnect option is set to true. The reused Channel is stored on the Exchange as an exchange property with the key <code>NettyConstants#NETTY_CHANNEL</code> which allows you to obtain the channel during routing and use it as well.	false	boolean
sync (common)	Setting to set endpoint as one-way or request-response.	true	boolean

Name	Description	Default	Type
tcpNoDelay (common)	Setting to improve TCP protocol performance.	true	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
broadcast (consumer)	Setting to choose Multicast over UDP.	false	boolean
clientMode (consumer)	If the clientMode is true, netty consumer will connect the address as a TCP client.	false	boolean
reconnect (consumer)	Used only in clientMode in consumer, the consumer will attempt to reconnect on disconnection if this is enabled.	true	boolean
reconnectInterval (consumer)	Used if reconnect and clientMode is enabled. The interval in milli seconds to attempt reconnection.	10000	int
backlog (consumer (advanced))	Allows to configure a backlog for netty consumer (server). Note the backlog is just a best effort depending on the OS. Setting this option to a value such as 200, 500 or 1000, tells the TCP stack how long the accept queue can be If this option is not configured, then the backlog depends on OS setting.		int
bossCount (consumer (advanced))	When netty works on nio mode, it uses default bossCount parameter from Netty, which is 1. User can use this option to override the default bossCount from Netty.	1	int
bossGroup (consumer (advanced))	Set the BossGroup which could be used for handling the new connection of the server side across the NettyEndpoint.		EventLoopGroup
disconnectOnNoReply (consumer (advanced))	If sync is enabled then this option dictates NettyConsumer if it should disconnect where there is no reply to send back.	true	boolean

Name	Description	Default	Type
executorService (consumer (advanced))	To use the given EventExecutorGroup.		EventExecutorGroup
maximumPoolSize (consumer (advanced))	Sets a maximum thread pool size for the netty consumer ordered thread pool. The default size is 2 x cpu_core plus 1. Setting this value to eg 10 will then use 10 threads unless 2 x cpu_core plus 1 is a higher value, which then will override and be used. For example if there are 8 cores, then the consumer thread pool will be 17. This thread pool is used to route messages received from Netty by Camel. We use a separate thread pool to ensure ordering of messages and also in case some messages will block, then netty's worker threads (event loop) won't be affected.		int
nettyServerBootstrapFactory (consumer (advanced))	To use a custom NettyServerBootstrapFactory.		NettyServerBootstrapFactory
networkInterface (consumer (advanced))	When using UDP then this option can be used to specify a network interface by its name, such as eth0 to join a multicast group.		String
noReplyLogLevel (consumer (advanced))	If sync is enabled this option dictates NettyConsumer which logging level to use when logging a there is no reply to send back. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LoggingLevel

Name	Description	Default	Type
serverClosedChannelExceptionHandlerLogLevel (consumer (advanced))	<p>If the server (NettyConsumer) catches an <code>java.nio.channels.ClosedChannelException</code> then its logged using this logging level. This is used to avoid logging the closed channel exceptions, as clients can disconnect abruptly and then cause a flood of closed exceptions in the Netty server.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	DEBUG	LogLevel
serverExceptionHandlerLogLevel (consumer (advanced))	<p>If the server (NettyConsumer) catches an exception then its logged using this logging level.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LogLevel
serverInitializerFactory (consumer (advanced))	To use a custom <code>ServerInitializerFactory</code> .		<code>ServerInitializerFactory</code>
usingExecutorService (consumer (advanced))	Whether to use ordered thread pool, to ensure events are processed orderly on the same channel.	true	boolean
connectTimeout (producer)	Time to wait for a socket connection to be available. Value is in milliseconds.	10000	int

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
requestTimeout (producer)	Allows to use a timeout for the Netty producer when calling a remote server. By default no timeout is in use. The value is in milli seconds, so eg 30000 is 30 seconds. The requestTimeout is using Netty's ReadTimeoutHandler to trigger the timeout.		long
clientInitializerFactory (producer (advanced))	To use a custom ClientInitializerFactory.		ClientInitializerFactory
correlationManager (producer (advanced))	To use a custom correlation manager to manage how request and reply messages are mapped when using request/reply with the netty producer. This should only be used if you have a way to map requests together with replies such as if there is correlation ids in both the request and reply messages. This can be used if you want to multiplex concurrent messages on the same channel (aka connection) in netty. When doing this you must have a way to correlate the request and reply messages so you can store the right reply on the inflight Camel Exchange before its continued routed. We recommend extending the TimeoutCorrelationManagerSupport when you build custom correlation managers. This provides support for timeout and other complexities you otherwise would need to implement as well. See also the producerPoolEnabled option for more details.		NettyCamelStateCorrelationManager
lazyChannelCreation (producer (advanced))	Channels can be lazily created to avoid exceptions, if the remote server is not up and running when the Camel producer is started.	true	boolean

Name	Description	Default	Type
producerPoolEnabled (producer (advanced))	Whether producer pool is enabled or not. Important: If you turn this off then a single shared connection is used for the producer, also if you are doing request/reply. That means there is a potential issue with interleaved responses if replies comes back out-of-order. Therefore you need to have a correlation id in both the request and reply messages so you can properly correlate the replies to the Camel callback that is responsible for continue processing the message in Camel. To do this you need to implement <code>NettyCamelStateCorrelationManager</code> as correlation manager and configure it via the <code>correlationManager</code> option. See also the <code>correlationManager</code> option for more details.	true	boolean
producerPoolMaxIdle (producer (advanced))	Sets the cap on the number of idle instances in the pool.	100	int
producerPoolMaxTotal (producer (advanced))	Sets the cap on the number of objects that can be allocated by the pool (checked out to clients, or idle awaiting checkout) at a given time. Use a negative value for no limit.	-1	int
producerPoolMinEvictableIdle (producer (advanced))	Sets the minimum amount of time (value in millis) an object may sit idle in the pool before it is eligible for eviction by the idle object evictor.	30000 0	long
producerPoolMinIdle (producer (advanced))	Sets the minimum number of instances allowed in the producer pool before the evictor thread (if active) spawns new objects.		int
udpConnectionlessSending (producer (advanced))	This option supports connection less udp sending which is a real fire and forget. A connected udp send receive the <code>PortUnreachableException</code> if no one is listen on the receiving port.	false	boolean
useByteBuf (producer (advanced))	If the <code>useByteBuf</code> is true, netty producer will turn the message body into <code>ByteBuf</code> before sending it out.	false	boolean
hostnameVerification (security)	To enable/disable hostname verification on <code>SSLEngine</code> .	false	boolean

Name	Description	Default	Type
allowSerializedHeaders (advanced)	Only used for TCP when <code>transferExchange</code> is true. When set to true, serializable objects in headers and properties will be added to the exchange. Otherwise Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
channelGroup (advanced)	To use a explicit ChannelGroup.		ChannelGroup
nativeTransport (advanced)	Whether to use native transport instead of NIO. Native transport takes advantage of the host operating system and is only supported on some platforms. You need to add the netty JAR for the host operating system you are using. See more details at: .	false	boolean
options (advanced)	Allows to configure additional netty options using option. as prefix. For example <code>option.child.keepAlive=false</code> to set the netty option <code>child.keepAlive=false</code> . See the Netty documentation for possible options that can be used.		Map
receiveBufferSize (advanced)	The TCP/UDP buffer sizes to be used during inbound communication. Size is bytes.	65536	int
receiveBufferSizePredictor (advanced)	Configures the buffer size predictor. See details at Jetty documentation and this mail thread.		int
sendBufferSize (advanced)	The TCP/UDP buffer sizes to be used during outbound communication. Size is bytes.	65536	int
transferExchange (advanced)	Only used for TCP. You can transfer the exchange over the wire instead of just the body. The following fields are transferred: In body, Out body, fault body, In headers, Out headers, fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean

Name	Description	Default	Type
udpByteArrayCodec (advanced)	For UDP only. If enabled the using byte array codec instead of Java serialization protocol.	false	boolean
workerCount (advanced)	When netty works on nio mode, it uses default workerCount parameter from Netty (which is <code>cpu_core_threads x 2</code>). User can use this option to override the default workerCount from Netty.		int
workerGroup (advanced)	To use a explicit EventLoopGroup as the boss thread pool. For example to share a thread pool with multiple consumers or producers. By default each consumer or producer has their own worker pool with <code>2 x cpu count core threads</code> .		EventLoopGroup
allowDefaultCodec (codec)	The netty component installs a default codec if both, encoder/decoder is null and textline is false. Setting <code>allowDefaultCodec</code> to false prevents the netty component from installing a default codec as the first element in the filter chain.	true	boolean
autoAppendDelimiter (codec)	Whether or not to auto append missing end delimiter when sending using the textline codec.	true	boolean
decoderMaxLineLength (codec)	The max line length to use for the textline codec.	1024	int
decoders (codec)	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with <code>#</code> so Camel knows it should lookup.		List
delimiter (codec)	The delimiter to use for the textline codec. Possible values are LINE and NULL. Enum values: <ul style="list-style-type: none"> ● LINE ● NULL 	LINE	TextLineDelimiter
encoders (codec)	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with <code>#</code> so Camel knows it should lookup.		List

Name	Description	Default	Type
encoding (codec)	The encoding (a charset name) to use for the textline codec. If not provided, Camel will use the JVM default Charset.		String
textline (codec)	Only used for TCP. If no codec is specified, you can use this flag to indicate a text line based codec; if not specified or the value is false, then Object Serialization is assumed over TCP - however only Strings are allowed to be serialized by default.	false	boolean
enabledProtocols (security)	Which protocols to enable when using SSL.	TLSv1, TLSv1.1, TLSv1.2	String
keyStoreFile (security)	Client side certificate keystore to be used for encryption.		File
keyStoreFormat (security)	Keystore format to be used for payload encryption. Defaults to JKS if not set.		String
keyStoreResource (security)	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
needClientAuth (security)	Configures whether the server needs client authentication when using SSL.	false	boolean
passphrase (security)	Password setting to use in order to encrypt/decrypt payloads sent using SSH.		String
securityProvider (security)	Security provider to be used for payload encryption. Defaults to SunX509 if not set.		String
ssl (security)	Setting to specify whether SSL encryption is applied to this endpoint.	false	boolean
sslClientCertHeaders (security)	When enabled and in SSL mode, then the Netty consumer will enrich the Camel Message with headers having information about the client certificate such as subject name, issuer name, serial number, and the valid date range.	false	boolean
sslContextParameters (security)	To configure security using SSLContextParameters.		SSLContextParameters

Name	Description	Default	Type
sslHandler (security)	Reference to a class that could be used to return an SSL Handler.		SslHandler
trustStoreFile (security)	Server side certificate keystore to be used for encryption.		File
trustStoreResource (security)	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean

36.4. ENDPOINT OPTIONS

The Netty endpoint is configured using URI syntax:

```
netty:protocol://host:port
```

with the following path and query parameters:

36.4.1. Path Parameters (3 parameters)

Name	Description	Default	Type
protocol (common)	Required The protocol to use which can be tcp or udp. Enum values: <ul style="list-style-type: none"> • tcp • udp 		String
host (common)	Required The hostname. For the consumer the hostname is localhost or 0.0.0.0. For the producer the hostname is the remote host to connect to.		String
port (common)	Required The host port number.		int

36.4.2. Query Parameters (71 parameters)

Name	Description	Default	Type
disconnect (common)	Whether or not to disconnect(close) from Netty Channel right after use. Can be used for both consumer and producer.	false	boolean
keepAlive (common)	Setting to ensure socket is not closed due to inactivity.	true	boolean
reuseAddress (common)	Setting to facilitate socket multiplexing.	true	boolean
reuseChannel (common)	This option allows producers and consumers (in client mode) to reuse the same Netty Channel for the lifecycle of processing the Exchange. This is useful if you need to call a server multiple times in a Camel route and want to use the same network connection. When using this, the channel is not returned to the connection pool until the Exchange is done; or disconnected if the disconnect option is set to true. The reused Channel is stored on the Exchange as an exchange property with the key <code>NettyConstants#NETTY_CHANNEL</code> which allows you to obtain the channel during routing and use it as well.	false	boolean
sync (common)	Setting to set endpoint as one-way or request-response.	true	boolean
tcpNoDelay (common)	Setting to improve TCP protocol performance.	true	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
broadcast (consumer)	Setting to choose Multicast over UDP.	false	boolean
clientMode (consumer)	If the clientMode is true, netty consumer will connect the address as a TCP client.	false	boolean

Name	Description	Default	Type
reconnect (consumer)	Used only in clientMode in consumer, the consumer will attempt to reconnect on disconnection if this is enabled.	true	boolean
reconnectInterval (consumer)	Used if reconnect and clientMode is enabled. The interval in milli seconds to attempt reconnection.	10000	int
backlog (consumer (advanced))	Allows to configure a backlog for netty consumer (server). Note the backlog is just a best effort depending on the OS. Setting this option to a value such as 200, 500 or 1000, tells the TCP stack how long the accept queue can be. If this option is not configured, then the backlog depends on OS setting.		int
bossCount (consumer (advanced))	When netty works on nio mode, it uses default bossCount parameter from Netty, which is 1. User can use this option to override the default bossCount from Netty.	1	int
bossGroup (consumer (advanced))	Set the BossGroup which could be used for handling the new connection of the server side across the NettyEndpoint.		EventLoopGroup
disconnectOnNoReply (consumer (advanced))	If sync is enabled then this option dictates NettyConsumer if it should disconnect where there is no reply to send back.	true	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
nettyServerBootstrapFactory (consumer (advanced))	To use a custom NettyServerBootstrapFactory.		NettyServerBootstrapFactory

Name	Description	Default	Type
networkInterface (consumer (advanced))	When using UDP then this option can be used to specify a network interface by its name, such as eth0 to join a multicast group.		String
noReplyLogLevel (consumer (advanced))	<p>If sync is enabled this option dictates NettyConsumer which logging level to use when logging a there is no reply to send back.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LoggingLevel
serverClosedChannelExceptionCaughtLogLevel (consumer (advanced))	<p>If the server (NettyConsumer) catches an java.nio.channels.ClosedChannelException then its logged using this logging level. This is used to avoid logging the closed channel exceptions, as clients can disconnect abruptly and then cause a flood of closed exceptions in the Netty server.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	DEBUG	LoggingLevel

Name	Description	Default	Type
serverExceptionHandlerLogLevel (consumer (advanced))	<p>If the server (NettyConsumer) catches an exception then its logged using this logging level.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	WARN	LogLevel
serverInitializerFactory (consumer (advanced))	To use a custom ServerInitializerFactory.		ServerInitializerFactory
usingExecutorService (consumer (advanced))	Whether to use ordered thread pool, to ensure events are processed orderly on the same channel.	true	boolean
connectTimeout (producer)	Time to wait for a socket connection to be available. Value is in milliseconds.	10000	int
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
requestTimeout (producer)	Allows to use a timeout for the Netty producer when calling a remote server. By default no timeout is in use. The value is in milli seconds, so eg 30000 is 30 seconds. The requestTimeout is using Netty's ReadTimeoutHandler to trigger the timeout.		long
clientInitializerFactory (producer (advanced))	To use a custom ClientInitializerFactory.		ClientInitializerFactory

Name	Description	Default	Type
correlationManager (producer (advanced))	To use a custom correlation manager to manage how request and reply messages are mapped when using request/reply with the netty producer. This should only be used if you have a way to map requests together with replies such as if there is correlation ids in both the request and reply messages. This can be used if you want to multiplex concurrent messages on the same channel (aka connection) in netty. When doing this you must have a way to correlate the request and reply messages so you can store the right reply on the inflight Camel Exchange before its continued routed. We recommend extending the TimeoutCorrelationManagerSupport when you build custom correlation managers. This provides support for timeout and other complexities you otherwise would need to implement as well. See also the producerPoolEnabled option for more details.		NettyCamelStateCorrelationManager
lazyChannelCreation (producer (advanced))	Channels can be lazily created to avoid exceptions, if the remote server is not up and running when the Camel producer is started.	true	boolean
producerPoolEnabled (producer (advanced))	Whether producer pool is enabled or not. Important: If you turn this off then a single shared connection is used for the producer, also if you are doing request/reply. That means there is a potential issue with interleaved responses if replies comes back out-of-order. Therefore you need to have a correlation id in both the request and reply messages so you can properly correlate the replies to the Camel callback that is responsible for continue processing the message in Camel. To do this you need to implement NettyCamelStateCorrelationManager as correlation manager and configure it via the correlationManager option. See also the correlationManager option for more details.	true	boolean
producerPoolMaxIdle (producer (advanced))	Sets the cap on the number of idle instances in the pool.	100	int
producerPoolMaxTotal (producer (advanced))	Sets the cap on the number of objects that can be allocated by the pool (checked out to clients, or idle awaiting checkout) at a given time. Use a negative value for no limit.	-1	int

Name	Description	Default	Type
producerPoolMinEvictableIdle (producer (advanced))	Sets the minimum amount of time (value in millis) an object may sit idle in the pool before it is eligible for eviction by the idle object evictor.	30000 0	long
producerPoolMinIdle (producer (advanced))	Sets the minimum number of instances allowed in the producer pool before the evictor thread (if active) spawns new objects.		int
udpConnectionlessSending (producer (advanced))	This option supports connection less udp sending which is a real fire and forget. A connected udp send receive the PortUnreachableException if no one is listen on the receiving port.	false	boolean
useByteBuf (producer (advanced))	If the useByteBuf is true, netty producer will turn the message body into ByteBuf before sending it out.	false	boolean
hostnameVerification (security)	To enable/disable hostname verification on SSLEngine.	false	boolean
allowSerializedHeaders (advanced)	Only used for TCP when transferExchange is true. When set to true, serializable objects in headers and properties will be added to the exchange. Otherwise Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean
channelGroup (advanced)	To use a explicit ChannelGroup.		ChannelGroup
nativeTransport (advanced)	Whether to use native transport instead of NIO. Native transport takes advantage of the host operating system and is only supported on some platforms. You need to add the netty JAR for the host operating system you are using. See more details at: .	false	boolean
options (advanced)	Allows to configure additional netty options using option. as prefix. For example option.child.keepAlive=false to set the netty option child.keepAlive=false. See the Netty documentation for possible options that can be used.		Map
receiveBufferSize (advanced)	The TCP/UDP buffer sizes to be used during inbound communication. Size is bytes.	65536	int

Name	Description	Default	Type
receiveBufferSizePredictor (advanced)	Configures the buffer size predictor. See details at Jetty documentation and this mail thread.		int
sendBufferSize (advanced)	The TCP/UDP buffer sizes to be used during outbound communication. Size is bytes.	65536	int
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
transferExchange (advanced)	Only used for TCP. You can transfer the exchange over the wire instead of just the body. The following fields are transferred: In body, Out body, fault body, In headers, Out headers, fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	boolean
udpByteArrayCodec (advanced)	For UDP only. If enabled the using byte array codec instead of Java serialization protocol.	false	boolean
workerCount (advanced)	When netty works on nio mode, it uses default workerCount parameter from Netty (which is <code>cpu_core_threads x 2</code>). User can use this option to override the default workerCount from Netty.		int
workerGroup (advanced)	To use a explicit EventLoopGroup as the boss thread pool. For example to share a thread pool with multiple consumers or producers. By default each consumer or producer has their own worker pool with 2 x cpu count core threads.		EventLoopGroup
allowDefaultCodec (codec)	The netty component installs a default codec if both, encoder/decoder is null and textline is false. Setting allowDefaultCodec to false prevents the netty component from installing a default codec as the first element in the filter chain.	true	boolean
autoAppendDelimiter (codec)	Whether or not to auto append missing end delimiter when sending using the textline codec.	true	boolean
decoderMaxLineLength (codec)	The max line length to use for the textline codec.	1024	int

Name	Description	Default	Type
decoders (codec)	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.		List
delimiter (codec)	The delimiter to use for the textline codec. Possible values are LINE and NULL. Enum values: <ul style="list-style-type: none"> ● LINE ● NULL 	LINE	TextLineDelimiter
encoders (codec)	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.		List
encoding (codec)	The encoding (a charset name) to use for the textline codec. If not provided, Camel will use the JVM default Charset.		String
textline (codec)	Only used for TCP. If no codec is specified, you can use this flag to indicate a text line based codec; if not specified or the value is false, then Object Serialization is assumed over TCP - however only Strings are allowed to be serialized by default.	false	boolean
enabledProtocols (security)	Which protocols to enable when using SSL.	TLSv1, TLSv1.1, TLSv1.2	String
keyStoreFile (security)	Client side certificate keystore to be used for encryption.		File
keyStoreFormat (security)	Keystore format to be used for payload encryption. Defaults to JKS if not set.		String
keyStoreResource (security)	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
needClientAuth (security)	Configures whether the server needs client authentication when using SSL.	false	boolean

Name	Description	Default	Type
passphrase (security)	Password setting to use in order to encrypt/decrypt payloads sent using SSH.		String
securityProvider (security)	Security provider to be used for payload encryption. Defaults to SunX509 if not set.		String
ssl (security)	Setting to specify whether SSL encryption is applied to this endpoint.	false	boolean
sslClientCertHeaders (security)	When enabled and in SSL mode, then the Netty consumer will enrich the Camel Message with headers having information about the client certificate such as subject name, issuer name, serial number, and the valid date range.	false	boolean
sslContextParameters (security)	To configure security using SSLContextParameters.		SSLContextParameters
sslHandler (security)	Reference to a class that could be used to return an SSL Handler.		SslHandler
trustStoreFile (security)	Server side certificate keystore to be used for encryption.		File
trustStoreResource (security)	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String

36.5. REGISTRY BASED OPTIONS

Codec Handlers and SSL Keystores can be enlisted in the Registry, such as in the Spring XML file. The values that could be passed in, are the following:

Name	Description
passphrase	password setting to use in order to encrypt/decrypt payloads sent using SSH
keyStoreFormat	keystore format to be used for payload encryption. Defaults to "JKS" if not set
securityProvider	Security provider to be used for payload encryption. Defaults to "SunX509" if not set.
keyStoreFile	deprecated: Client side certificate keystore to be used for encryption

Name	Description
trustStoreFile	deprecated: Server side certificate keystore to be used for encryption
keyStoreResource	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with " classpath: ", " file: ", or " http: " to load the resource from different systems.
trustStoreResource	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with " classpath: ", " file: ", or " http: " to load the resource from different systems.
sslHandler	Reference to a class that could be used to return an SSL Handler
encoder	A custom ChannelHandler class that can be used to perform special marshalling of outbound payloads. Must override <code>io.netty.channel.ChannelInboundHandlerAdapter</code> .
encoders	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.
decoder	A custom ChannelHandler class that can be used to perform special marshalling of inbound payloads. Must override <code>io.netty.channel.ChannelOutboundHandlerAdapter</code> .
decoders	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.



NOTE

Read below about using non shareable encoders/decoders.

36.5.1. Using non shareable encoders or decoders

If your encoders or decoders are not shareable (e.g. they don't have the `@Shareable` class annotation), then your encoder/decoder must implement the **`org.apache.camel.component.netty.ChannelHandlerFactory`** interface, and return a new instance in the **`newChannelHandler`** method. This is to ensure the encoder/decoder can safely be used. If this is not the case, then the Netty component will log a WARN when an endpoint is created.

The Netty component offers a **`org.apache.camel.component.netty.ChannelHandlerFactories`** factory class, that has a number of commonly used methods.

36.6. SENDING MESSAGES TO/FROM A NETTY ENDPOINT

36.6.1. Netty Producer

In Producer mode, the component provides the ability to send payloads to a socket endpoint using either TCP or UDP protocols (with optional SSL support).

The producer mode supports both one-way and request-response based operations.

36.6.2. Netty Consumer

In Consumer mode, the component provides the ability to:

- listen on a specified socket using either TCP or UDP protocols (with optional SSL support),
- receive requests on the socket using text/xml, binary and serialized object based payloads and
- send them along on a route as message exchanges.

The consumer mode supports both one-way and request-response based operations.

36.7. EXAMPLES

36.7.1. A UDP Netty endpoint using Request-Reply and serialized object payload

Note that Object serialization is not allowed by default, and so a decoder must be configured.

```
@BindToRegistry("decoder")
public ChannelHandler getDecoder() throws Exception {
    return new DefaultChannelHandlerFactory() {
        @Override
        public ChannelHandler newChannelHandler() {
            return new DatagramPacketObjectDecoder(ClassResolvers.weakCachingResolver(null));
        }
    };
}

RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("netty:udp://0.0.0.0:5155?sync=true&decoders=#decoder")
            .process(new Processor() {
                public void process(Exchange exchange) throws Exception {
                    Poetry poetry = (Poetry) exchange.getIn().getBody();
                    // Process poetry in some way
                    exchange.getOut().setBody("Message received");
                }
            })
    }
};
```

36.7.2. A TCP based Netty consumer endpoint using One-way communication

```
RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("netty:tcp://0.0.0.0:5150")
            .to("mock:result");
    }
};
```

36.7.3. An SSL/TCP based Netty consumer endpoint using Request-Reply communication

Using the JSSE Configuration Utility

The Netty component supports SSL/TLS configuration through the [Camel JSSE Configuration Utility](#). This utility greatly decreases the amount of component specific code you need to write and is configurable at the endpoint and component levels. The following examples demonstrate how to use the utility with the Netty component.

Programmatic configuration of the component

```

KeyStoreParameters ksp = new KeyStoreParameters();
ksp.setResource("/users/home/server/keystore.jks");
ksp.setPassword("keystorePassword");

KeyManagersParameters kmp = new KeyManagersParameters();
kmp.setKeyStore(ksp);
kmp.setKeyPassword("keyPassword");

SSLContextParameters scp = new SSLContextParameters();
scp.setKeyManagers(kmp);

NettyComponent nettyComponent = getContext().getComponent("netty", NettyComponent.class);
nettyComponent.setSslContextParameters(scp);

```

Spring DSL based configuration of endpoint

```

...
<camel:sslContextParameters
  id="sslContextParameters">
  <camel:keyManagers
    keyPassword="keyPassword">
    <camel:keyStore
      resource="/users/home/server/keystore.jks"
      password="keystorePassword"/>
    </camel:keyManagers>
  </camel:sslContextParameters>...
...
<to uri="netty:tcp://0.0.0.0:5150?
sync=true&ssl=true&sslContextParameters=#sslContextParameters"/>
...

```

Using Basic SSL/TLS configuration on the Jetty Component

```

Registry registry = context.getRegistry();
registry.bind("password", "changeit");
registry.bind("ksf", new File("src/test/resources/keystore.jks"));
registry.bind("tsf", new File("src/test/resources/keystore.jks"));

context.addRoutes(new RouteBuilder() {
  public void configure() {
    String netty_ssl_endpoint =
      "netty:tcp://0.0.0.0:5150?sync=true&ssl=true&passphrase=#password"
      + "&keyStoreFile=#ksf&trustStoreFile=#tsf";

```

```
String return_string =
    "When You Go Home, Tell Them Of Us And Say,"
    + "For Your Tomorrow, We Gave Our Today.";

from(netty_ssl_endpoint)
    .process(new Processor() {
        public void process(Exchange exchange) throws Exception {
            exchange.getOut().setBody(return_string);
        }
    })
    });
```

Getting access to SSLSession and the client certificate

You can get access to the **javax.net.ssl.SSLSession** if you eg need to get details about the client certificate. When **ssl=true** then the Netty component will store the **SSLSession** as a header on the Camel Message as shown below:

```
SSLSession session = exchange.getIn().getHeader(NettyConstants.NETTY_SSL_SESSION,
    SSLSession.class);
// get the first certificate which is client certificate
javax.security.cert.X509Certificate cert = session.getPeerCertificateChain()[0];
Principal principal = cert.getSubjectDN();
```

Remember to set **needClientAuth=true** to authenticate the client, otherwise **SSLSession** cannot access information about the client certificate, and you may get an exception **javax.net.ssl.SSLPeerUnverifiedException: peer not authenticated**. You may also get this exception if the client certificate is expired or not valid etc.



NOTE

The option **sslClientCertHeaders** can be set to **true** which then enriches the Camel Message with headers having details about the client certificate. For example the subject name is readily available in the header **CamelNettySSLClientCertSubjectName**.

36.7.4. Using Multiple Codecs

In certain cases it may be necessary to add chains of encoders and decoders to the netty pipeline. To add multiple codecs to a camel netty endpoint the 'encoders' and 'decoders' uri parameters should be used. Like the 'encoder' and 'decoder' parameters they are used to supply references (lists of ChannelUpstreamHandlers and ChannelDownstreamHandlers) that should be added to the pipeline. Note that if encoders is specified then the encoder param will be ignored, similarly for decoders and the decoder param.



NOTE

Read further above about using non shareable encoders/decoders.

The lists of codecs need to be added to the Camel's registry so they can be resolved when the endpoint is created.

```
ChannelHandlerFactory lengthDecoder =
    ChannelHandlerFactories.newLengthFieldBasedFrameDecoder(1048576, 0, 4, 0, 4);
```

```

StringDecoder stringDecoder = new StringDecoder();
registry.bind("length-decoder", lengthDecoder);
registry.bind("string-decoder", stringDecoder);

LengthFieldPrepender lengthEncoder = new LengthFieldPrepender(4);
StringEncoder stringEncoder = new StringEncoder();
registry.bind("length-encoder", lengthEncoder);
registry.bind("string-encoder", stringEncoder);

List<ChannelHandler> decoders = new ArrayList<ChannelHandler>();
decoders.add(lengthDecoder);
decoders.add(stringDecoder);

List<ChannelHandler> encoders = new ArrayList<ChannelHandler>();
encoders.add(lengthEncoder);
encoders.add(stringEncoder);

registry.bind("encoders", encoders);
registry.bind("decoders", decoders);

```

Spring's native collections support can be used to specify the codec lists in an application context

```

<util:list id="decoders" list-class="java.util.LinkedList">
  <bean class="org.apache.camel.component.netty.ChannelHandlerFactories" factory-
method="newLengthFieldBasedFrameDecoder">
    <constructor-arg value="1048576"/>
    <constructor-arg value="0"/>
    <constructor-arg value="4"/>
    <constructor-arg value="0"/>
    <constructor-arg value="4"/>
  </bean>
  <bean class="io.netty.handler.codec.string.StringDecoder"/>
</util:list>

<util:list id="encoders" list-class="java.util.LinkedList">
  <bean class="io.netty.handler.codec.LengthFieldPrepender">
    <constructor-arg value="4"/>
  </bean>
  <bean class="io.netty.handler.codec.string.StringEncoder"/>
</util:list>

<bean id="length-encoder" class="io.netty.handler.codec.LengthFieldPrepender">
  <constructor-arg value="4"/>
</bean>
<bean id="string-encoder" class="io.netty.handler.codec.string.StringEncoder"/>

<bean id="length-decoder" class="org.apache.camel.component.netty.ChannelHandlerFactories"
factory-method="newLengthFieldBasedFrameDecoder">
  <constructor-arg value="1048576"/>
  <constructor-arg value="0"/>
  <constructor-arg value="4"/>
  <constructor-arg value="0"/>
  <constructor-arg value="4"/>
</bean>
<bean id="string-decoder" class="io.netty.handler.codec.string.StringDecoder"/>

```

The bean names can then be used in netty endpoint definitions either as a comma separated list or contained in a List e.g.

```
from("direct:multiple-codec").to("netty:tcp://0.0.0.0:{{port}}?encoders=#encoders&sync=false");

from("netty:tcp://0.0.0.0:{{port}}?decoders=#length-decoder,#string-decoder&sync=false").to("mock:multiple-codec");
```

or via XML.

```
<camelContext id="multiple-netty-codecs-context" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:multiple-codec"/>
    <to uri="netty:tcp://0.0.0.0:5150?encoders=#encoders&sync=false"/>
  </route>
  <route>
    <from uri="netty:tcp://0.0.0.0:5150?decoders=#length-decoder,#string-decoder&sync=false"/>
    <to uri="mock:multiple-codec"/>
  </route>
</camelContext>
```

36.8. CLOSING CHANNEL WHEN COMPLETE

When acting as a server you sometimes want to close the channel when, for example, a client conversion is finished.

You can do this by simply setting the endpoint option **disconnect=true**.

However you can also instruct Camel on a per message basis as follows.

To instruct Camel to close the channel, you should add a header with the key

CamelNettyCloseChannelWhenComplete set to a boolean **true** value.

For instance, the example below will close the channel after it has written the bye message back to the client:

```
from("netty:tcp://0.0.0.0:8080").process(new Processor() {
  public void process(Exchange exchange) throws Exception {
    String body = exchange.getIn().getBody(String.class);
    exchange.getOut().setBody("Bye " + body);
    // some condition which determines if we should close
    if (close) {
      exchange.getOut().setHeader(NettyConstants.NETTY_CLOSE_CHANNEL_WHEN_COMPLETE,
        true);
    }
  }
});
```

Adding custom channel pipeline factories to gain complete control over a created pipeline.

36.9. CUSTOM PIPELINE

Custom channel pipelines provide complete control to the user over the handler/interceptor chain by inserting custom handler(s), encoder(s) & decoder(s) without having to specify them in the Netty Endpoint URL in a very simple way.

In order to add a custom pipeline, a custom channel pipeline factory must be created and registered with the context via the context registry (Registry, or the camel-spring ApplicationContextRegistry etc).

A custom pipeline factory must be constructed as follows

- A Producer linked channel pipeline factory must extend the abstract class **ClientPipelineFactory**.
- A Consumer linked channel pipeline factory must extend the abstract class **ServerInitializerFactory**.
- The classes should override the `initChannel()` method in order to insert custom handler(s), encoder(s) and decoder(s). Not overriding the **initChannel()** method creates a pipeline with no handlers, encoders or decoders wired to the pipeline.

The example below shows how `ServerInitializerFactory` factory may be created

36.9.1. Using custom pipeline factory

```
public class SampleServerInitializerFactory extends ServerInitializerFactory {
    private int maxLineSize = 1024;

    protected void initChannel(Channel ch) throws Exception {
        ChannelPipeline channelPipeline = ch.pipeline();

        channelPipeline.addLast("encoder-SD", new StringEncoder(CharsetUtil.UTF_8));
        channelPipeline.addLast("decoder-DELIM", new DelimiterBasedFrameDecoder(maxLineSize,
true, Delimiters.lineDelimiter()));
        channelPipeline.addLast("decoder-SD", new StringDecoder(CharsetUtil.UTF_8));
        // here we add the default Camel ServerChannelHandler for the consumer, to allow Camel to
route the message etc.
        channelPipeline.addLast("handler", new ServerChannelHandler(consumer));
    }
}
```

The custom channel pipeline factory can then be added to the registry and instantiated/utilized on a camel route in the following way

```
Registry registry = camelContext.getRegistry();
ServerInitializerFactory factory = new TestServerInitializerFactory();
registry.bind("spf", factory);
context.addRoutes(new RouteBuilder() {
    public void configure() {
        String netty_ssl_endpoint =
            "netty:tcp://0.0.0.0:5150?serverInitializerFactory=#spf"
        String return_string =
            "When You Go Home, Tell Them Of Us And Say,"
            + "For Your Tomorrow, We Gave Our Today.";

        from(netty_ssl_endpoint)
            .process(new Processor() {
                public void process(Exchange exchange) throws Exception {
```



```

        exchange.getOut().setBody(return_string);
    }
}
});

```

36.10. REUSING NETTY BOSS AND WORKER THREAD POOLS

Netty has two kind of thread pools: boss and worker. By default each Netty consumer and producer has their private thread pools. If you want to reuse these thread pools among multiple consumers or producers then the thread pools must be created and enlisted in the Registry.


For example using Spring XML we can create a shared worker thread pool using the **NettyWorkerPoolBuilder** with 2 worker threads as shown below:

```

<!-- use the worker pool builder to help create the shared thread pool -->
<bean id="poolBuilder" class="org.apache.camel.component.netty.NettyWorkerPoolBuilder">
  <property name="workerCount" value="2"/>
</bean>

<!-- the shared worker thread pool -->
<bean id="sharedPool" class="org.jboss.netty.channel.socket.nio.WorkerPool"
  factory-bean="poolBuilder" factory-method="build" destroy-method="shutdown">
</bean>

```



NOTE

For boss thread pool there is a **org.apache.camel.component.netty.NettyServerBossPoolBuilder** builder for Netty consumers, and a **org.apache.camel.component.netty.NettyClientBossPoolBuilder** for the Netty producers.

Then in the Camel routes we can refer to this worker pools by configuring the **workerPool** option in the URI as shown below:

```

<route>
  <from uri="netty:tcp://0.0.0.0:5021?
textline=true&sync=true&workerPool=#sharedPool&usingExecutorService=false"/>
  <to uri="log:result"/>
  ...
</route>

```

And if we have another route we can refer to the shared worker pool:

```

<route>
  <from uri="netty:tcp://0.0.0.0:5022?
textline=true&sync=true&workerPool=#sharedPool&usingExecutorService=false"/>
  <to uri="log:result"/>
  ...
</route>

```

and so forth.

36.11. MULTIPLEXING CONCURRENT MESSAGES OVER A SINGLE CONNECTION WITH REQUEST/REPLY

When using Netty for request/reply messaging via the netty producer then by default each message is sent via a non-shared connection (pooled). This ensures that replies are automatically being able to map to the correct request thread for further routing in Camel. In other words correlation between request/reply messages happens out-of-the-box because the replies come back on the same connection that was used for sending the request; and this connection is not shared with others. When the response comes back, the connection is returned back to the connection pool, where it can be reused by others.

However if you want to multiplex concurrent request/responses on a single shared connection, then you need to turn off the connection pooling by setting **producerPoolEnabled=false**. Now this means there is a potential issue with interleaved responses if replies come back out-of-order. Therefore you need to have a correlation id in both the request and reply messages so you can properly correlate the replies to the Camel callback that is responsible for continuing processing the message in Camel. To do this you need to implement **NettyCamelStateCorrelationManager** as correlation manager and configure it via the **correlationManager=#myManager** option.



NOTE

We recommend extending the **TimeoutCorrelationManagerSupport** when you build custom correlation managers. This provides support for timeout and other complexities you otherwise would need to implement as well.

You can find an example with the Apache Camel source code in the examples directory under the **camel-example-netty-custom-correlation** directory.

36.12. SPRING BOOT AUTO-CONFIGURATION

When using netty with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-netty-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 74 options, which are listed below.

Name	Description	Default	Type
camel.component.netty.allow-default-codec	The netty component installs a default codec if both, encoder/decoder is null and textline is false. Setting allowDefaultCodec to false prevents the netty component from installing a default codec as the first element in the filter chain.	true	Boolean

Name	Description	Default	Type
camel.component.netty.allow-serialized-headers	Only used for TCP when transferExchange is true. When set to true, serializable objects in headers and properties will be added to the exchange. Otherwise Camel will exclude any non-serializable objects and log it at WARN level.	false	Boolean
camel.component.netty.auto-append-delimiter	Whether or not to auto append missing end delimiter when sending using the textline codec.	true	Boolean
camel.component.netty.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.netty.backlog	Allows to configure a backlog for netty consumer (server). Note the backlog is just a best effort depending on the OS. Setting this option to a value such as 200, 500 or 1000, tells the TCP stack how long the accept queue can be. If this option is not configured, then the backlog depends on OS setting.		Integer
camel.component.netty.boss-count	When netty works on nio mode, it uses default bossCount parameter from Netty, which is 1. User can use this option to override the default bossCount from Netty.	1	Integer
camel.component.netty.boss-group	Set the BossGroup which could be used for handling the new connection of the server side across the NettyEndpoint. The option is a io.netty.channel.EventLoopGroup type.		EventLoopGroup
camel.component.netty.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean

Name	Description	Default	Type
<code>camel.component.netty.broadcast</code>	Setting to choose Multicast over UDP.	false	Boolean
<code>camel.component.netty.channel-group</code>	To use a explicit ChannelGroup. The option is a <code>io.netty.channel.group.ChannelGroup</code> type.		ChannelGroup
<code>camel.component.netty.client-initializer-factory</code>	To use a custom ClientInitializerFactory. The option is a <code>org.apache.camel.component.netty.ClientInitializerFactory</code> type.		ClientInitializerFactory
<code>camel.component.netty.client-mode</code>	If the <code>clientMode</code> is true, netty consumer will connect the address as a TCP client.	false	Boolean
<code>camel.component.netty.configuration</code>	To use the NettyConfiguration as configuration when creating endpoints. The option is a <code>org.apache.camel.component.netty.NettyConfiguration</code> type.		NettyConfiguration
<code>camel.component.netty.connect-timeout</code>	Time to wait for a socket connection to be available. Value is in milliseconds.	10000	Integer
<code>camel.component.netty.correlation-manager</code>	To use a custom correlation manager to manage how request and reply messages are mapped when using request/reply with the netty producer. This should only be used if you have a way to map requests together with replies such as if there is correlation ids in both the request and reply messages. This can be used if you want to multiplex concurrent messages on the same channel (aka connection) in netty. When doing this you must have a way to correlate the request and reply messages so you can store the right reply on the inflight Camel Exchange before its continued routed. We recommend extending the <code>TimeoutCorrelationManagerSupport</code> when you build custom correlation managers. This provides support for timeout and other complexities you otherwise would need to implement as well. See also the <code>producerPoolEnabled</code> option for more details. The option is a <code>org.apache.camel.component.netty.NettyCamelStateCorrelationManager</code> type.		NettyCamelStateCorrelationManager
<code>camel.component.netty.decoder-max-line-length</code>	The max line length to use for the textline codec.	1024	Integer

Name	Description	Default	Type
camel.component.netty.decoders	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.		String
camel.component.netty.delimiter	The delimiter to use for the textline codec. Possible values are LINE and NULL.		TextLineDelimiter
camel.component.netty.disconnect	Whether or not to disconnect(close) from Netty Channel right after use. Can be used for both consumer and producer.	false	Boolean
camel.component.netty.disconnect-on-no-reply	If sync is enabled then this option dictates NettyConsumer if it should disconnect where there is no reply to send back.	true	Boolean
camel.component.netty.enabled	Whether to enable auto configuration of the netty component. This is enabled by default.		Boolean
camel.component.netty.enabled-protocols	Which protocols to enable when using SSL.	TLSv1, TLSv1.1, TLSv1.2	String
camel.component.netty.encoders	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.		String
camel.component.netty.encoding	The encoding (a charset name) to use for the textline codec. If not provided, Camel will use the JVM default Charset.		String
camel.component.netty.executor-service	To use the given EventExecutorGroup. The option is a io.netty.util.concurrent.EventExecutorGroup type.		EventExecutorGroup
camel.component.netty.hostname-verification	To enable/disable hostname verification on SSLEngine.	false	Boolean
camel.component.netty.keep-alive	Setting to ensure socket is not closed due to inactivity.	true	Boolean

Name	Description	Default	Type
<code>camel.component.netty.key-store-file</code>	Client side certificate keystore to be used for encryption.		File
<code>camel.component.netty.key-store-format</code>	Keystore format to be used for payload encryption. Defaults to JKS if not set.		String
<code>camel.component.netty.key-store-resource</code>	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.		String
<code>camel.component.netty.lazy-channel-creation</code>	Channels can be lazily created to avoid exceptions, if the remote server is not up and running when the Camel producer is started.	true	Boolean
<code>camel.component.netty.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.netty.maximum-pool-size</code>	Sets a maximum thread pool size for the netty consumer ordered thread pool. The default size is 2 x cpu_core plus 1. Setting this value to eg 10 will then use 10 threads unless 2 x cpu_core plus 1 is a higher value, which then will override and be used. For example if there are 8 cores, then the consumer thread pool will be 17. This thread pool is used to route messages received from Netty by Camel. We use a separate thread pool to ensure ordering of messages and also in case some messages will block, then netty's worker threads (event loop) won't be affected.		Integer
<code>camel.component.netty.native-transport</code>	Whether to use native transport instead of NIO. Native transport takes advantage of the host operating system and is only supported on some platforms. You need to add the netty JAR for the host operating system you are using. See more details at: .	false	Boolean

Name	Description	Default	Type
<code>camel.component.netty.need-client-auth</code>	Configures whether the server needs client authentication when using SSL.	false	Boolean
<code>camel.component.netty.netty-server-bootstrap-factory</code>	To use a custom <code>NettyServerBootstrapFactory</code> . The option is a <code>org.apache.camel.component.netty.NettyServerBootstrapFactory</code> type.		<code>NettyServerBootstrapFactory</code>
<code>camel.component.netty.network-interface</code>	When using UDP then this option can be used to specify a network interface by its name, such as <code>eth0</code> to join a multicast group.		String
<code>camel.component.netty.no-reply-log-level</code>	If sync is enabled this option dictates <code>NettyConsumer</code> which logging level to use when logging a there is no reply to send back.		LoggingLevel
<code>camel.component.netty.options</code>	Allows to configure additional netty options using option. as prefix. For example <code>option.child.keepAlive=false</code> to set the netty option <code>child.keepAlive=false</code> . See the Netty documentation for possible options that can be used.		Map
<code>camel.component.netty.passphrase</code>	Password setting to use in order to encrypt/decrypt payloads sent using SSH.		String
<code>camel.component.netty.producer-pool-enabled</code>	Whether producer pool is enabled or not. Important: If you turn this off then a single shared connection is used for the producer, also if you are doing request/reply. That means there is a potential issue with interleaved responses if replies comes back out-of-order. Therefore you need to have a correlation id in both the request and reply messages so you can properly correlate the replies to the Camel callback that is responsible for continue processing the message in Camel. To do this you need to implement <code>NettyCamelStateCorrelationManager</code> as correlation manager and configure it via the <code>correlationManager</code> option. See also the <code>correlationManager</code> option for more details.	true	Boolean
<code>camel.component.netty.producer-pool-max-idle</code>	Sets the cap on the number of idle instances in the pool.	100	Integer

Name	Description	Default	Type
<code>camel.component.netty.producer.pool-max-total</code>	Sets the cap on the number of objects that can be allocated by the pool (checked out to clients, or idle awaiting checkout) at a given time. Use a negative value for no limit.	-1	Integer
<code>camel.component.netty.producer.pool-min-evictable-idle</code>	Sets the minimum amount of time (value in millis) an object may sit idle in the pool before it is eligible for eviction by the idle object evictor.	30000 0	Long
<code>camel.component.netty.producer.pool-min-idle</code>	Sets the minimum number of instances allowed in the producer pool before the evictor thread (if active) spawns new objects.		Integer
<code>camel.component.netty.receive-buffer-size</code>	The TCP/UDP buffer sizes to be used during inbound communication. Size is bytes.	65536	Integer
<code>camel.component.netty.receive-buffer-size-predictor</code>	Configures the buffer size predictor. See details at Jetty documentation and this mail thread.		Integer
<code>camel.component.netty.reconnect</code>	Used only in clientMode in consumer, the consumer will attempt to reconnect on disconnection if this is enabled.	true	Boolean
<code>camel.component.netty.reconnect-interval</code>	Used if reconnect and clientMode is enabled. The interval in milli seconds to attempt reconnection.	10000	Integer
<code>camel.component.netty.request-timeout</code>	Allows to use a timeout for the Netty producer when calling a remote server. By default no timeout is in use. The value is in milli seconds, so eg 30000 is 30 seconds. The requestTimeout is using Netty's ReadTimeoutHandler to trigger the timeout.		Long
<code>camel.component.netty.reuse-address</code>	Setting to facilitate socket multiplexing.	true	Boolean

Name	Description	Default	Type
camel.component.netty.reuse-channel	This option allows producers and consumers (in client mode) to reuse the same Netty Channel for the lifecycle of processing the Exchange. This is useful if you need to call a server multiple times in a Camel route and want to use the same network connection. When using this, the channel is not returned to the connection pool until the Exchange is done; or disconnected if the disconnect option is set to true. The reused Channel is stored on the Exchange as an exchange property with the key <code>NettyConstants#NETTY_CHANNEL</code> which allows you to obtain the channel during routing and use it as well.	false	Boolean
camel.component.netty.security-provider	Security provider to be used for payload encryption. Defaults to SunX509 if not set.		String
camel.component.netty.send-buffer-size	The TCP/UDP buffer sizes to be used during outbound communication. Size is bytes.	65536	Integer
camel.component.netty.server-closed-channel-exception-caught-log-level	If the server (NettyConsumer) catches an <code>java.nio.channels.ClosedChannelException</code> then its logged using this logging level. This is used to avoid logging the closed channel exceptions, as clients can disconnect abruptly and then cause a flood of closed exceptions in the Netty server.		LogLevel
camel.component.netty.server-exception-caught-log-level	If the server (NettyConsumer) catches an exception then its logged using this logging level.		LogLevel
camel.component.netty.server-initializer-factory	To use a custom <code>ServerInitializerFactory</code> . The option is a <code>org.apache.camel.component.netty.ServerInitializerFactory</code> type.		<code>ServerInitializerFactory</code>
camel.component.netty.ssl	Setting to specify whether SSL encryption is applied to this endpoint.	false	Boolean
camel.component.netty.ssl-client-cert-headers	When enabled and in SSL mode, then the Netty consumer will enrich the Camel Message with headers having information about the client certificate such as subject name, issuer name, serial number, and the valid date range.	false	Boolean

Name	Description	Default	Type
<code>camel.component.netty.ssl-context-parameters</code>	To configure security using <code>SSLContextParameters</code> . The option is a <code>org.apache.camel.support.jsse.SSLContextParameters</code> type.		<code>SSLContextParameters</code>
<code>camel.component.netty.ssl-handler</code>	Reference to a class that could be used to return an SSL Handler. The option is a <code>io.netty.handler.ssl.SslHandler</code> type.		<code>SslHandler</code>
<code>camel.component.netty.sync</code>	Setting to set endpoint as one-way or request-response.	true	Boolean
<code>camel.component.netty.tcp-no-delay</code>	Setting to improve TCP protocol performance.	true	Boolean
<code>camel.component.netty.textline</code>	Only used for TCP. If no codec is specified, you can use this flag to indicate a text line based codec; if not specified or the value is false, then Object Serialization is assumed over TCP - however only Strings are allowed to be serialized by default.	false	Boolean
<code>camel.component.netty.transfer-exchange</code>	Only used for TCP. You can transfer the exchange over the wire instead of just the body. The following fields are transferred: In body, Out body, fault body, In headers, Out headers, fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	Boolean
<code>camel.component.netty.trust-store-file</code>	Server side certificate keystore to be used for encryption.		File
<code>camel.component.netty.trust-store-resource</code>	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with <code>classpath:</code> , <code>file:</code> , or <code>http:</code> to load the resource from different systems.		String
<code>camel.component.netty.udp-byte-array-codec</code>	For UDP only. If enabled the using byte array codec instead of Java serialization protocol.	false	Boolean
<code>camel.component.netty.udp-connectionless-sending</code>	This option supports connection less udp sending which is a real fire and forget. A connected udp send receive the <code>PortUnreachableException</code> if no one is listen on the receiving port.	false	Boolean

Name	Description	Default	Type
camel.component.netty.use-bytebuf	If the useByteBuf is true, netty producer will turn the message body into ByteBuf before sending it out.	false	Boolean
camel.component.netty.use-global-ssl-context-parameters	Enable usage of global SSL context parameters.	false	Boolean
camel.component.netty.using-executor-service	Whether to use ordered thread pool, to ensure events are processed orderly on the same channel.	true	Boolean
camel.component.netty.worker-count	When netty works on nio mode, it uses default workerCount parameter from Netty (which is <code>cpu_core_threads x 2</code>). User can use this option to override the default workerCount from Netty.		Integer
camel.component.netty.worker-group	To use a explicit EventLoopGroup as the boss thread pool. For example to share a thread pool with multiple consumers or producers. By default each consumer or producer has their own worker pool with <code>2 x cpu count core threads</code> . The option is a <code>io.netty.channel.EventLoopGroup</code> type.		EventLoopGroup

CHAPTER 37. PAHO

Both producer and consumer are supported

Paho component provides connector for the MQTT messaging protocol using the [Eclipse Paho library](#). Paho is one of the most popular MQTT libraries, so if you would like to integrate it with your Java project - Camel Paho connector is a way to go.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-paho</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

37.1. URI FORMAT

```
paho:topic[?options]
```

Where **topic** is the name of the topic.

37.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

37.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

37.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings.

In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

37.3. COMPONENT OPTIONS

The Paho component supports 31 options, which are listed below.

Name	Description	Default	Type
automaticReconnect (common)	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	boolean
brokerUrl (common)	The URL of the MQTT broker.	tcp://localhost:1883	String
cleanSession (common)	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	boolean
clientId (common)	MQTT client identifier. The identifier must be unique.		String
configuration (common)	To use the shared Paho configuration.		PahoConfiguration

Name	Description	Default	Type
connectionTimeout (common)	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	int
filePersistenceDirectory (common)	Base directory used by file persistence. Will by default use user directory.		String
keepAliveInterval (common)	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	int
maxInflight (common)	Sets the max inflight. please increase this value in a high traffic environment. The default value is 10.	10	int
maxReconnectDelay (common)	Get the maximum time (in millis) to wait between reconnects.	128000	int
mqttVersion (common)	Sets the MQTT version. The default action is to connect with version 3.1.1, and to fall back to 3.1 if that fails. Version 3.1.1 or 3.1 can be selected specifically, with no fall back, by using the MQTT_VERSION_3_1_1 or MQTT_VERSION_3_1 options respectively.		int
persistence (common)	Client persistence to be used - memory or file. Enum values: <ul style="list-style-type: none"> ● FILE ● MEMORY 	MEMORY	PahoPersistence
qos (common)	Client quality of service level (0-2).	2	int
retained (common)	Retain option.	false	boolean

Name	Description	Default	Type
serverURIs (common)	Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.		String
willPayload (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String
willQos (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		int

Name	Description	Default	Type
willRetained (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.	false	boolean
willTopic (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
client (advanced)	To use a shared Paho client.		MqttClient
customWebSocketHeaders (advanced)	Sets the Custom WebSocket Headers for the WebSocket Connection.		Properties
executorServiceTimeout (advanced)	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	int
httpsHostnameVerificationEnabled (security)	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	boolean
password (security)	Password to be used for authentication against the MQTT broker.		String
socketFactory (security)	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings.		SocketFactory
sslClientProps (security)	<p>Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is available. These properties are not used if a custom SocketFactory has been set. The following properties can be used:</p> <ul style="list-style-type: none"> com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS. com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the KeyManager to use. For example /mydir/etc/key.p12 com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords. com.ibm.ssl.keyStoreType Type of key store, for example PKCS12, JKS, or JCEKS. com.ibm.ssl.keyStoreProvider Key store provider, for example IBMJCE or IBMJCEFIPS. com.ibm.ssl.trustStore The name of the file that 		Properties

Name	Description	Default	Type
	<p>contains the KeyStore object that you want the TrustManager to use.</p> <p><code>com.ibm.ssl.trustStorePassword</code> The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method:</p> <p><code>com.ibm.micro.security.Password.obfuscate(char password)</code>. This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p><code>com.ibm.ssl.trustStoreType</code> The type of KeyStore object that you want the default TrustManager to use. Same possible values as <code>keyStoreType</code>.</p> <p><code>com.ibm.ssl.trustStoreProvider</code> Trust store provider, for example <code>IBMJCE</code> or <code>IBMJCEFIPS</code>.</p> <p><code>com.ibm.ssl.enabledCipherSuites</code> A list of which ciphers are enabled. Values are dependent on the provider, for example:</p> <p><code>SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA</code>.</p> <p><code>com.ibm.ssl.keyManager</code> Sets the algorithm that will be used to instantiate a <code>KeyManagerFactory</code> object instead of using the default algorithm available in the platform. Example values: <code>lbnX509</code> or <code>IBMJ9X509</code>.</p> <p><code>com.ibm.ssl.trustManager</code> Sets the algorithm that will be used to instantiate a <code>TrustManagerFactory</code> object instead of using the default algorithm available in the platform. Example values: <code>PKIX</code> or <code>IBMJ9X509</code>.</p>		
<code>sslHostnameVerifier</code> (security)	Sets the <code>HostnameVerifier</code> for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default <code>HostnameVerifier</code> .		<code>HostnameVerifier</code>
<code>userName</code> (security)	Username to be used for authentication against the MQTT broker.		String

37.4. ENDPOINT OPTIONS

The Paho endpoint is configured using URI syntax:

```
paho:topic
```

with the following path and query parameters:

37.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
topic (common)	Required Name of the topic.		String

37.4.2. Query Parameters (31 parameters)

Name	Description	Default	Type
automaticReconnect (common)	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	boolean
brokerUrl (common)	The URL of the MQTT broker.	tcp://localhost:1883	String
cleanSession (common)	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	boolean
clientId (common)	MQTT client identifier. The identifier must be unique.		String
connectionTimeout (common)	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	int

Name	Description	Default	Type
filePersistenceDirectory (common)	Base directory used by file persistence. Will by default use user directory.		String
keepAliveInterval (common)	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	int
maxInflight (common)	Sets the max inflight. please increase this value in a high traffic environment. The default value is 10.	10	int
maxReconnectDelay (common)	Get the maximum time (in millis) to wait between reconnects.	128000	int
mqttVersion (common)	Sets the MQTT version. The default action is to connect with version 3.1.1, and to fall back to 3.1 if that fails. Version 3.1.1 or 3.1 can be selected specifically, with no fall back, by using the MQTT_VERSION_3_1_1 or MQTT_VERSION_3_1 options respectively.		int
persistence (common)	Client persistence to be used - memory or file. Enum values: <ul style="list-style-type: none"> ● FILE ● MEMORY 	MEMORY	PahoPersistence
qos (common)	Client quality of service level (0-2).	2	int
retained (common)	Retain option.	false	boolean

Name	Description	Default	Type
serverURIs (common)	Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.		String
willPayload (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String
willQos (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		int

Name	Description	Default	Type
willRetained (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.	false	boolean
willTopic (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● <code>InOnly</code> ● <code>InOut</code> ● <code>InOptionalOut</code> 		<code>ExchangePattern</code>

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
client (advanced)	To use an existing mqtt client.		MqttClient
customWebSocketHeaders (advanced)	Sets the Custom WebSocket Headers for the WebSocket Connection.		Properties
executorServiceTimeout (advanced)	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	int
httpsHostnameVerificationEnabled (security)	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	boolean
password (security)	Password to be used for authentication against the MQTT broker.		String
socketFactory (security)	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings.		SocketFactory
sslClientProps (security)	Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is available. These properties are not used if a custom SocketFactory has been set. The following properties can be used: com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS. com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the KeyManager to use. For example /mydir/etc/key.p12 com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to		Properties

Name	Description	Default	Type
	<p>use. The password can either be in plain-text, or may be obfuscated using the static method: <code>com.ibm.micro.security.Password.obfuscate(char password)</code>. This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p><code>com.ibm.ssl.keyStoreType</code> Type of key store, for example PKCS12, JKS, or JCEKS.</p> <p><code>com.ibm.ssl.keyStoreProvider</code> Key store provider, for example IBMJCE or IBMJCEFIPS.</p> <p><code>com.ibm.ssl.trustStore</code> The name of the file that contains the KeyStore object that you want the TrustManager to use.</p> <p><code>com.ibm.ssl.trustStorePassword</code> The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method: <code>com.ibm.micro.security.Password.obfuscate(char password)</code>. This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p><code>com.ibm.ssl.trustStoreType</code> The type of KeyStore object that you want the default TrustManager to use. Same possible values as <code>keyStoreType</code>.</p> <p><code>com.ibm.ssl.trustStoreProvider</code> Trust store provider, for example IBMJCE or IBMJCEFIPS.</p> <p><code>com.ibm.ssl.enabledCipherSuites</code> A list of which ciphers are enabled. Values are dependent on the provider, for example: <code>SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA</code>.</p> <p><code>com.ibm.ssl.keyManager</code> Sets the algorithm that will be used to instantiate a KeyManagerFactory object instead of using the default algorithm available in the platform. Example values: <code>IbmX509</code> or <code>IBMJ9X509</code>.</p> <p><code>com.ibm.ssl.trustManager</code> Sets the algorithm that will be used to instantiate a TrustManagerFactory object instead of using the default algorithm available in the platform. Example values: <code>PKIX</code> or <code>IBMJ9X509</code>.</p>		
sslHostnameVerifier (security)	Sets the HostnameVerifier for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default HostnameVerifier.		HostnameVerifier
userName (security)	Username to be used for authentication against the MQTT broker.		String

37.5. HEADERS

The following headers are recognized by the Paho component:

Header	Java constant	Endpoint type	Value type	Description
CamelMqttTopic	PahoConstants.MQTT_TOPIC	Consumer	String	The name of the topic
CamelMqttQoS	PahoConstants.MQTT_QOS	Consumer	Integer	QualityOfService of the incoming message
CamelPahoOverrideTopic	PahoConstants.CAMEL_PAHO_OVERRIDE_TOPIC	Producer	String	Name of topic to override and send to instead of topic specified on endpoint

37.6. DEFAULT PAYLOAD TYPE

By default Camel Paho component operates on the binary payloads extracted out of (or put into) the MQTT message:

```
// Receive payload
byte[] payload = (byte[]) consumerTemplate.receiveBody("paho:topic");

// Send payload
byte[] payload = "message".getBytes();
producerTemplate.sendBody("paho:topic", payload);
```

But of course Camel build-in [type conversion API](#) can perform the automatic data type transformations for you. In the example below Camel automatically converts binary payload into **String** (and conversely):

```
// Receive payload
String payload = consumerTemplate.receiveBody("paho:topic", String.class);

// Send payload
String payload = "message";
producerTemplate.sendBody("paho:topic", payload);
```

37.7. SAMPLES

For example the following snippet reads messages from the MQTT broker installed on the same host as the Camel router:

```
from("paho:some/queue")
    .to("mock:test");
```

While the snippet below sends message to the MQTT broker:

```
from("direct:test")
    .to("paho:some/target/queue");
```

For example this is how to read messages from the remote MQTT broker:

```
from("paho:some/queue?brokerUrl=tcp://iot.eclipse.org:1883")
    .to("mock:test");
```

And here we override the default topic and set to a dynamic topic

```
from("direct:test")
    .setHeader(PahoConstants.CAMEL_PAHO_OVERRIDE_TOPIC, simple("${header.customerId}"))
    .to("paho:some/target/queue");
```

37.8. SPRING BOOT AUTO-CONFIGURATION

When using paho with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-paho-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 32 options, which are listed below.

Name	Description	Default	Type
camel.component.paho.automatic-reconnect	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	Boolean
camel.component.paho.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.paho.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.paho.broker-url	The URL of the MQTT broker.	tcp://localhost:1883	String
camel.component.paho.clean-session	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	Boolean
camel.component.paho.client	To use a shared Paho client. The option is a <code>org.eclipse.paho.client.mqttv3.MqttClient</code> type.		MqttClient
camel.component.paho.client-id	MQTT client identifier. The identifier must be unique.		String
camel.component.paho.configuration	To use the shared Paho configuration. The option is a <code>org.apache.camel.component.paho.PahoConfiguration</code> type.		PahoConfiguration
camel.component.paho.connection-timeout	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	Integer

Name	Description	Default	Type
<code>camel.component.paho.custom-web-socket-headers</code>	Sets the Custom WebSocket Headers for the WebSocket Connection. The option is a <code>java.util.Properties</code> type.		Properties
<code>camel.component.paho.enabled</code>	Whether to enable auto configuration of the paho component. This is enabled by default.		Boolean
<code>camel.component.paho.executor-service-timeout</code>	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	Integer
<code>camel.component.paho.file-persistence-directory</code>	Base directory used by file persistence. Will by default use user directory.		String
<code>camel.component.paho.https-hostname-verification-enabled</code>	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	Boolean
<code>camel.component.paho.keep-alive-interval</code>	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	Integer
<code>camel.component.paho.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

Name	Description	Default	Type
camel.component.paho.max-inflight	Sets the max inflight. please increase this value in a high traffic environment. The default value is 10.	10	Integer
camel.component.paho.max-reconnect-delay	Get the maximum time (in millis) to wait between reconnects.	128000	Integer
camel.component.paho.mqtt-version	Sets the MQTT version. The default action is to connect with version 3.1.1, and to fall back to 3.1 if that fails. Version 3.1.1 or 3.1 can be selected specifically, with no fall back, by using the MQTT_VERSION_3_1_1 or MQTT_VERSION_3_1 options respectively.		Integer
camel.component.paho.password	Password to be used for authentication against the MQTT broker.		String
camel.component.paho.persistence	Client persistence to be used - memory or file.		PahoPersistence
camel.component.paho.qos	Client quality of service level (0-2).	2	Integer
camel.component.paho.retained	Retain option.	false	Boolean

Name	Description	Default	Type
camel.component.paho.server-uris	Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.		String
camel.component.paho.socket-factory	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings. The option is a javax.net.SocketFactory type.		SocketFactory
camel.component.paho.ssl-client-props	Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is available. These properties are not used if a custom SocketFactory has been set. The following properties can be used: com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS. com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the		Properties

Name	Description	Default	Type
	<p>KeyManager to use. For example /mydir/etc/key.p12</p> <p>com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.keyStoreType Type of key store, for example PKCS12, JKS, or JCEKS.</p> <p>com.ibm.ssl.keyStoreProvider Key store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.trustStore The name of the file that contains the KeyStore object that you want the TrustManager to use.</p> <p>com.ibm.ssl.trustStorePassword The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.trustStoreType The type of KeyStore object that you want the default TrustManager to use. Same possible values as keyStoreType.</p> <p>com.ibm.ssl.trustStoreProvider Trust store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.enabledCipherSuites A list of which ciphers are enabled. Values are dependent on the provider, for example: SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA. com.ibm.ssl.keyManager Sets the algorithm that will be used to instantiate a KeyManagerFactory object instead of using the default algorithm available in the platform. Example values: IbmX509 or IBMJ9X509.</p> <p>com.ibm.ssl.trustManager Sets the algorithm that will be used to instantiate a TrustManagerFactory object instead of using the default algorithm available in the platform. Example values: PKIX or IBMJ9X509. The option is a java.util.Properties type.</p>		
camel.component.paho.ssl-hostname-verifier	Sets the HostnameVerifier for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default HostnameVerifier. The option is a javax.net.ssl.HostnameVerifier type.		HostnameVerifier

Name	Description	Default	Type
camel.component.paho.user-name	Username to be used for authentication against the MQTT broker.		String
camel.component.paho.will-payload	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String
camel.component.paho.will-qos	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		Integer
camel.component.paho.will-retained	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.	false	Boolean
camel.component.paho.will-topic	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to The byte payload for the message. The quality of service to publish the message at (0, 1 or 2). Whether or not the message should be retained.		String

CHAPTER 38. PAHO MQTT 5

Both producer and consumer are supported

Paho MQTT5 component provides connector for the MQTT messaging protocol using the [Eclipse Paho](#) library with MQTT v5. Paho is one of the most popular MQTT libraries, so if you would like to integrate it with your Java project – Camel Paho connector is a way to go.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-paho-mqtt5</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

38.1. URI FORMAT

```
paho-mqtt5:topic[?options]
```

Where **topic** is the name of the topic.

38.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

38.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

38.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

38.3. COMPONENT OPTIONS

The Paho MQTT 5 component supports 32 options, which are listed below.

Name	Description	Default	Type
automaticReconnect (common)	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	boolean
brokerUrl (common)	The URL of the MQTT broker.	tcp://localhost:1883	String
cleanStart (common)	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	boolean
clientId (common)	MQTT client identifier. The identifier must be unique.		String
configuration (common)	To use the shared Paho configuration.		PahoMqtt5Configuration

Name	Description	Default	Type
connectionTimeout (common)	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	int
filePersistenceDirectory (common)	Base directory used by file persistence. Will by default use user directory.		String
keepAliveInterval (common)	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	int
maxReconnectDelay (common)	Get the maximum time (in millis) to wait between reconnects.	12800 0	int
persistence (common)	Client persistence to be used - memory or file. Enum values: <ul style="list-style-type: none"> ● FILE ● MEMORY 	MEMORY	PahoMqtt5Persistence
qos (common)	Client quality of service level (0-2).	2	int
receiveMaximum (common)	Sets the Receive Maximum. This value represents the limit of QoS 1 and QoS 2 publications that the client is willing to process concurrently. There is no mechanism to limit the number of QoS 0 publications that the Server might try to send. The default value is 65535.	65535	int
retained (common)	Retain option.	false	boolean

Name	Description	Default	Type
serverURIs (common)	<p>Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.</p>		String
sessionExpiryInterval (common)	<p>Sets the Session Expiry Interval. This value, measured in seconds, defines the maximum time that the broker will maintain the session for once the client disconnects. Clients should only connect with a long Session Expiry interval if they intend to connect to the server at some later point in time. By default this value is -1 and so will not be sent, in this case, the session will not expire. If a 0 is sent, the session will end immediately once the Network Connection is closed. When the client has determined that it has no longer any use for the session, it should disconnect with a Session Expiry Interval set to 0.</p>	-1	long

Name	Description	Default	Type
willMqttProperties (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The MQTT properties set for the message.		MqttProperties
willPayload (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The byte payload for the message.		String
willQos (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The quality of service to publish the message at (0, 1 or 2).	1	int
willRetained (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. Whether or not the message should be retained.	false	boolean
willTopic (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
client (advanced)	To use a shared Paho client.		MqttClient
customWebSocketHeaders (advanced)	Sets the Custom WebSocket Headers for the WebSocket Connection.		Map
executorServiceTimeout (advanced)	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	int
httpsHostnameVerificationEnabled (security)	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	boolean
password (security)	Password to be used for authentication against the MQTT broker.		String
socketFactory (security)	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings.		SocketFactory
sslClientProps (security)	Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is		Properties

Name	Description	Default	Type
	<p>available. These properties are not used if a custom SocketFactory has been set. The following properties can be used: com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS.</p> <p>com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE</p> <p>com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the KeyManager to use. For example /mydir/etc/key.p12</p> <p>com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.keyStoreType Type of key store, for example PKCS12, JKS, or JCEKS.</p> <p>com.ibm.ssl.keyStoreProvider Key store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.trustStore The name of the file that contains the KeyStore object that you want the TrustManager to use.</p> <p>com.ibm.ssl.trustStorePassword The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.trustStoreType The type of KeyStore object that you want the default TrustManager to use. Same possible values as keyStoreType.</p> <p>com.ibm.ssl.trustStoreProvider Trust store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.enabledCipherSuites A list of which ciphers are enabled. Values are dependent on the provider, for example: SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA. com.ibm.ssl.keyManager Sets the algorithm that will be used to instantiate a KeyManagerFactory object instead of using the default algorithm available in the platform. Example values: IbmX509 or IBMJ9X509.</p> <p>com.ibm.ssl.trustManager Sets the algorithm that will be used to instantiate a TrustManagerFactory object instead of using the default algorithm available in the platform. Example values: PKIX or IBMJ9X509.</p>		

Name	Description	Default	Type
sslHostnameVerifier (security)	Sets the HostnameVerifier for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default HostnameVerifier.		HostnameVerifier
userName (security)	Username to be used for authentication against the MQTT broker.		String

38.4. ENDPOINT OPTIONS

The Paho MQTT 5 endpoint is configured using URI syntax:

```
paho-mqtt5:topic
```

with the following path and query parameters:

38.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
topic (common)	Required Name of the topic.		String

38.4.2. Query Parameters (32 parameters)

Name	Description	Default	Type
automaticReconnect (common)	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	boolean
brokerUrl (common)	The URL of the MQTT broker.	tcp://localhost:1883	String

Name	Description	Default	Type
cleanStart (common)	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	boolean
clientId (common)	MQTT client identifier. The identifier must be unique.		String
connectionTimeout (common)	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	int
filePersistenceDirectory (common)	Base directory used by file persistence. Will by default use user directory.		String
keepAliveInterval (common)	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	int
maxReconnectDelay (common)	Get the maximum time (in millis) to wait between reconnects.	12800 0	int

Name	Description	Default	Type
persistence (common)	Client persistence to be used - memory or file. Enum values: <ul style="list-style-type: none"> • FILE • MEMORY 	MEMORY	PahoMqtt5Persistence
qos (common)	Client quality of service level (0-2).	2	int
receiveMaximum (common)	Sets the Receive Maximum. This value represents the limit of QoS 1 and QoS 2 publications that the client is willing to process concurrently. There is no mechanism to limit the number of QoS 0 publications that the Server might try to send. The default value is 65535.	65535	int
retained (common)	Retain option.	false	boolean

Name	Description	Default	Type
serverURIs (common)	Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.		String
sessionExpiryInterval (common)	Sets the Session Expiry Interval. This value, measured in seconds, defines the maximum time that the broker will maintain the session for once the client disconnects. Clients should only connect with a long Session Expiry interval if they intend to connect to the server at some later point in time. By default this value is -1 and so will not be sent, in this case, the session will not expire. If a 0 is sent, the session will end immediately once the Network Connection is closed. When the client has determined that it has no longer any use for the session, it should disconnect with a Session Expiry Interval set to 0.	-1	long

Name	Description	Default	Type
willMqttProperties (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The MQTT properties set for the message.		MqttProperties
willPayload (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The byte payload for the message.		String
willQos (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The quality of service to publish the message at (0, 1 or 2).	1	int
willRetained (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. Whether or not the message should be retained.	false	boolean
willTopic (common)	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
client (advanced)	To use an existing mqtt client.		MqttClient
customWebSocketHeaders (advanced)	Sets the Custom WebSocket Headers for the WebSocket Connection.		Map
executorServiceTimeout (advanced)	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	int
httpsHostnameVerificationEnabled (security)	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	boolean
password (security)	Password to be used for authentication against the MQTT broker.		String
socketFactory (security)	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings.		SocketFactory
sslClientProps (security)	Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is		Properties

Name	Description	Default	Type
	<p>available. These properties are not used if a custom SocketFactory has been set. The following properties can be used: com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS.</p> <p>com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE</p> <p>com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the KeyManager to use. For example /mydir/etc/key.p12</p> <p>com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.keyStoreType Type of key store, for example PKCS12, JKS, or JCEKS.</p> <p>com.ibm.ssl.keyStoreProvider Key store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.trustStore The name of the file that contains the KeyStore object that you want the TrustManager to use.</p> <p>com.ibm.ssl.trustStorePassword The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method: com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords.</p> <p>com.ibm.ssl.trustStoreType The type of KeyStore object that you want the default TrustManager to use. Same possible values as keyStoreType.</p> <p>com.ibm.ssl.trustStoreProvider Trust store provider, for example IBMJCE or IBMJCEFIPS.</p> <p>com.ibm.ssl.enabledCipherSuites A list of which ciphers are enabled. Values are dependent on the provider, for example: SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA. com.ibm.ssl.keyManager Sets the algorithm that will be used to instantiate a KeyManagerFactory object instead of using the default algorithm available in the platform. Example values: IbmX509 or IBMJ9X509.</p> <p>com.ibm.ssl.trustManager Sets the algorithm that will be used to instantiate a TrustManagerFactory object instead of using the default algorithm available in the platform. Example values: PKIX or IBMJ9X509.</p>		

Name	Description	Default	Type
sslHostnameVerifier (security)	Sets the HostnameVerifier for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default HostnameVerifier.		HostnameVerifier
userName (security)	Username to be used for authentication against the MQTT broker.		String

38.5. HEADERS

The following headers are recognized by the Paho component:

Header	Java constant	Endpoint type	Value type	Description
CamelMqttTopic	PahoConstants.MQTT_TOPIC	Consumer	String	The name of the topic
CamelMqttQoS	PahoConstants.MQTT_QOS	Consumer	Integer	QualityOfService of the incoming message
CamelPahoOverrideTopic	PahoConstants.CAMEL_PAHO_OVERRIDE_TOPIC	Producer	String	Name of topic to override and send to instead of topic specified on endpoint

38.6. DEFAULT PAYLOAD TYPE

By default Camel Paho component operates on the binary payloads extracted out of (or put into) the MQTT message:

```
// Receive payload
byte[] payload = (byte[]) consumerTemplate.receiveBody("paho:topic");

// Send payload
byte[] payload = "message".getBytes();
producerTemplate.sendBody("paho:topic", payload);
```

But of course Camel build-in [type conversion API](#) can perform the automatic data type transformations for you. In the example below Camel automatically converts binary payload into **String** (and conversely):

```
// Receive payload
String payload = consumerTemplate.receiveBody("paho:topic", String.class);

// Send payload
String payload = "message";
producerTemplate.sendBody("paho:topic", payload);
```

-

38.7. SAMPLES

For example the following snippet reads messages from the MQTT broker installed on the same host as the Camel router:

```
from("paho:some/queue")
  .to("mock:test");
```

While the snippet below sends message to the MQTT broker:

```
from("direct:test")
  .to("paho:some/target/queue");
```

For example this is how to read messages from the remote MQTT broker:

```
from("paho:some/queue?brokerUrl=tcp://iot.eclipse.org:1883")
  .to("mock:test");
```

And here we override the default topic and set to a dynamic topic

```
from("direct:test")
  .setHeader(PahoConstants.CAMEL_PAHO_OVERRIDE_TOPIC, simple("${header.customerId}"))
  .to("paho:some/target/queue");
```

38.8. SPRING BOOT AUTO-CONFIGURATION

When using paho-mqtt5 with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-paho-mqtt5-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 33 options, which are listed below.

Name	Description	Default	Type
		t	

Name	Description	Default	Type
camel.component.paho-mqtt5.automatic-reconnect	Sets whether the client will automatically attempt to reconnect to the server if the connection is lost. If set to false, the client will not attempt to automatically reconnect to the server in the event that the connection is lost. If set to true, in the event that the connection is lost, the client will attempt to reconnect to the server. It will initially wait 1 second before it attempts to reconnect, for every failed reconnect attempt, the delay will double until it is at 2 minutes at which point the delay will stay at 2 minutes.	true	Boolean
camel.component.paho-mqtt5.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.paho-mqtt5.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.paho-mqtt5.broker-url	The URL of the MQTT broker.	tcp://localhost:1883	String

Name	Description	Default	Type
<code>camel.component.paho-mqtt5.clean-start</code>	Sets whether the client and server should remember state across restarts and reconnects. If set to false both the client and server will maintain state across restarts of the client, the server and the connection. As state is maintained: Message delivery will be reliable meeting the specified QOS even if the client, server or connection are restarted. The server will treat a subscription as durable. If set to true the client and server will not maintain state across restarts of the client, the server or the connection. This means Message delivery to the specified QOS cannot be maintained if the client, server or connection are restarted The server will treat a subscription as non-durable.	true	Boolean
<code>camel.component.paho-mqtt5.client</code>	To use a shared Paho client. The option is a <code>org.eclipse.paho.mqttv5.client.MqttClient</code> type.		MqttClient
<code>camel.component.paho-mqtt5.client-id</code>	MQTT client identifier. The identifier must be unique.		String
<code>camel.component.paho-mqtt5.configuration</code>	To use the shared Paho configuration. The option is a <code>org.apache.camel.component.paho.mqtt5.PahoMqtt5Configuration</code> type.		PahoMqtt5Configuration
<code>camel.component.paho-mqtt5.connection-timeout</code>	Sets the connection timeout value. This value, measured in seconds, defines the maximum time interval the client will wait for the network connection to the MQTT server to be established. The default timeout is 30 seconds. A value of 0 disables timeout processing meaning the client will wait until the network connection is made successfully or fails.	30	Integer
<code>camel.component.paho-mqtt5.custom-web-socket-headers</code>	Sets the Custom WebSocket Headers for the WebSocket Connection.		Map
<code>camel.component.paho-mqtt5.enabled</code>	Whether to enable auto configuration of the paho-mqtt5 component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.paho-mqtt5.executor-service-timeout</code>	Set the time in seconds that the executor service should wait when terminating before forcefully terminating. It is not recommended to change this value unless you are absolutely sure that you need to.	1	Integer
<code>camel.component.paho-mqtt5.file-persistence-directory</code>	Base directory used by file persistence. Will by default use user directory.		String
<code>camel.component.paho-mqtt5.https-hostname-verification-enabled</code>	Whether SSL HostnameVerifier is enabled or not. The default value is true.	true	Boolean
<code>camel.component.paho-mqtt5.keep-alive-interval</code>	Sets the keep alive interval. This value, measured in seconds, defines the maximum time interval between messages sent or received. It enables the client to detect if the server is no longer available, without having to wait for the TCP/IP timeout. The client will ensure that at least one message travels across the network within each keep alive period. In the absence of a data-related message during the time period, the client sends a very small ping message, which the server will acknowledge. A value of 0 disables keepalive processing in the client. The default value is 60 seconds.	60	Integer
<code>camel.component.paho-mqtt5.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.paho-mqtt5.max-reconnect-delay</code>	Get the maximum time (in millis) to wait between reconnects.	12800 0	Integer

Name	Description	Default	Type
<code>camel.component.paho-mqtt5.password</code>	Password to be used for authentication against the MQTT broker.		String
<code>camel.component.paho-mqtt5.persistence</code>	Client persistence to be used - memory or file.		PahoMqtt5Persistence
<code>camel.component.paho-mqtt5.qos</code>	Client quality of service level (0-2).	2	Integer
<code>camel.component.paho-mqtt5.receive-maximum</code>	Sets the Receive Maximum. This value represents the limit of QoS 1 and QoS 2 publications that the client is willing to process concurrently. There is no mechanism to limit the number of QoS 0 publications that the Server might try to send. The default value is 65535.	65535	Integer
<code>camel.component.paho-mqtt5.retained</code>	Retain option.	false	Boolean

Name	Description	Default	Type
camel.component.paho-mqtt5.server-uris	Set a list of one or more serverURIs the client may connect to. Multiple servers can be separated by comma. Each serverURI specifies the address of a server that the client may connect to. Two types of connection are supported tcp:// for a TCP connection and ssl:// for a TCP connection secured by SSL/TLS. For example: tcp://localhost:1883 ssl://localhost:8883 If the port is not specified, it will default to 1883 for tcp:// URIs, and 8883 for ssl:// URIs. If serverURIs is set then it overrides the serverURI parameter passed in on the constructor of the MQTT client. When an attempt to connect is initiated the client will start with the first serverURI in the list and work through the list until a connection is established with a server. If a connection cannot be made to any of the servers then the connect attempt fails. Specifying a list of servers that a client may connect to has several uses: High Availability and reliable message delivery Some MQTT servers support a high availability feature where two or more equal MQTT servers share state. An MQTT client can connect to any of the equal servers and be assured that messages are reliably delivered and durable subscriptions are maintained no matter which server the client connects to. The cleansession flag must be set to false if durable subscriptions and/or reliable message delivery is required. Hunt List A set of servers may be specified that are not equal (as in the high availability option). As no state is shared across the servers reliable message delivery and durable subscriptions are not valid. The cleansession flag must be set to true if the hunt list mode is used.		String
camel.component.paho-mqtt5.session-expiry-interval	Sets the Session Expiry Interval. This value, measured in seconds, defines the maximum time that the broker will maintain the session for once the client disconnects. Clients should only connect with a long Session Expiry interval if they intend to connect to the server at some later point in time. By default this value is -1 and so will not be sent, in this case, the session will not expire. If a 0 is sent, the session will end immediately once the Network Connection is closed. When the client has determined that it has no longer any use for the session, it should disconnect with a Session Expiry Interval set to 0.	-1	Long

Name	Description	Default	Type
camel.component.paho-mqtt5.socket-factory	Sets the SocketFactory to use. This allows an application to apply its own policies around the creation of network sockets. If using an SSL connection, an SSLSocketFactory can be used to supply application-specific security settings. The option is a javax.net.SocketFactory type.		SocketFactory
camel.component.paho-mqtt5.ssl-client-props	<p>Sets the SSL properties for the connection. Note that these properties are only valid if an implementation of the Java Secure Socket Extensions (JSSE) is available. These properties are not used if a custom SocketFactory has been set. The following properties can be used:</p> <ul style="list-style-type: none"> com.ibm.ssl.protocol One of: SSL, SSLv3, TLS, TLSv1, SSL_TLS. com.ibm.ssl.contextProvider Underlying JSSE provider. For example IBMJSSE2 or SunJSSE com.ibm.ssl.keyStore The name of the file that contains the KeyStore object that you want the KeyManager to use. For example /mydir/etc/key.p12 com.ibm.ssl.keyStorePassword The password for the KeyStore object that you want the KeyManager to use. The password can either be in plain-text, or may be obfuscated using the static method: <ul style="list-style-type: none"> com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords. com.ibm.ssl.keyStoreType Type of key store, for example PKCS12, JKS, or JCEKS. com.ibm.ssl.keyStoreProvider Key store provider, for example IBMJCE or IBMJCEFIPS. com.ibm.ssl.trustStore The name of the file that contains the KeyStore object that you want the TrustManager to use. com.ibm.ssl.trustStorePassword The password for the TrustStore object that you want the TrustManager to use. The password can either be in plain-text, or may be obfuscated using the static method: <ul style="list-style-type: none"> com.ibm.micro.security.Password.obfuscate(char password). This obfuscates the password using a simple and insecure XOR and Base64 encoding mechanism. Note that this is only a simple scrambler to obfuscate clear-text passwords. com.ibm.ssl.trustStoreType The type of KeyStore object that you want the default TrustManager to use. Same possible values as keyStoreType. com.ibm.ssl.trustStoreProvider Trust store provider, for example IBMJCE or IBMJCEFIPS. com.ibm.ssl.enabledCipherSuites A list of which ciphers are enabled. Values are dependent on the 		Properties

Name	Description	Default	Type
	<p>provider, for example: SSL_RSA_WITH_AES_128_CBC_SHA;SSL_RSA_WITH_3DES_EDE_CBC_SHA. com.ibm.ssl.keyManager</p> <p>Sets the algorithm that will be used to instantiate a KeyManagerFactory object instead of using the default algorithm available in the platform. Example values: IbmX509 or IBMJ9X509.</p> <p>com.ibm.ssl.trustManager Sets the algorithm that will be used to instantiate a TrustManagerFactory object instead of using the default algorithm available in the platform. Example values: PKIX or IBMJ9X509. The option is a java.util.Properties type.</p>		
camel.component.paho-mqtt5.ssl-hostname-verifyer	Sets the HostnameVerifier for the SSL connection. Note that it will be used after handshake on a connection and you should do actions by yourself when hostname is verified error. There is no default HostnameVerifier. The option is a javax.net.ssl.HostnameVerifier type.		HostnameVerifier
camel.component.paho-mqtt5.user-name	Username to be used for authentication against the MQTT broker.		String
camel.component.paho-mqtt5.will-mqtt-properties	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The MQTT properties set for the message. The option is a org.eclipse.paho.mqttv5.common.packet.MqttProperties type.		MqttProperties
camel.component.paho-mqtt5.will-payload	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The byte payload for the message.		String
camel.component.paho-mqtt5.will-qos	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The quality of service to publish the message at (0, 1 or 2).	1	Integer
camel.component.paho-mqtt5.will-retained	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. Whether or not the message should be retained.	false	Boolean

Name	Description	Default	Type
camel.component.paho-mqtt5.will-topic	Sets the Last Will and Testament (LWT) for the connection. In the event that this client unexpectedly loses its connection to the server, the server will publish a message to itself using the supplied details. The topic to publish to.		String

CHAPTER 39. QUARTZ

Only consumer is supported

The Quartz component provides a scheduled delivery of messages using the [Quartz Scheduler 2.x](#). Each endpoint represents a different timer (in Quartz terms, a Trigger and JobDetail).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-quartz</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

39.1. URI FORMAT

```
quartz://timerName?options
quartz://groupName/timerName?options
quartz://groupName/timerName?cron=expression
quartz://timerName?cron=expression
```

The component uses either a **CronTrigger** or a **SimpleTrigger**. If no cron expression is provided, the component uses a simple trigger. If no **groupName** is provided, the quartz component uses the **Camel** group name.

39.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

39.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

39.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

39.3. COMPONENT OPTIONS

The Quartz component supports 13 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
enableJmx (consumer)	Whether to enable Quartz JMX which allows to manage the Quartz scheduler from JMX. This options is default true.	true	boolean
prefixInstanceName (consumer)	Whether to prefix the Quartz Scheduler instance name with the CamelContext name. This is enabled by default, to let each CamelContext use its own Quartz scheduler instance by default. You can set this option to false to reuse Quartz scheduler instances between multiple CamelContext's.	true	boolean
prefixJobNameWithEndpointId (consumer)	Whether to prefix the quartz job with the endpoint id. This option is default false.	false	boolean
properties (consumer)	Properties to configure the Quartz scheduler.		Map
propertiesFile (consumer)	File name of the properties to load from the classpath.		String
propertiesRef (consumer)	References to an existing Properties or Map to lookup in the registry to use for configuring quartz.		String

Name	Description	Default	Type
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
scheduler (advanced)	To use the custom configured Quartz scheduler, instead of creating a new Scheduler.		Scheduler
schedulerFactory (advanced)	To use the custom SchedulerFactory which is used to create the Scheduler.		SchedulerFactory
autoStartScheduler (scheduler)	Whether or not the scheduler should be auto started. This options is default true.	true	boolean
interruptJobsOnShutdown (scheduler)	Whether to interrupt jobs on shutdown which forces the scheduler to shutdown quicker and attempt to interrupt any running jobs. If this is enabled then any running jobs can fail due to being interrupted. When a job is interrupted then Camel will mark the exchange to stop continue routing and set <code>java.util.concurrent.RejectedExecutionException</code> as caused exception. Therefore use this with care, as its often better to allow Camel jobs to complete and shutdown gracefully.	false	boolean
startDelayedSeconds (scheduler)	Seconds to wait before starting the quartz scheduler.		int

39.4. ENDPOINT OPTIONS

The Quartz endpoint is configured using URI syntax:

```
quartz:groupName/triggerName
```

with the following path and query parameters:

39.4.1. Path Parameters (2 parameters)

Name	Description	Default	Type
groupName (consumer)	The quartz group name to use. The combination of group name and trigger name should be unique.	Camel	String

Name	Description	Default	Type
triggerName (consumer)	Required The quartz trigger name to use. The combination of group name and trigger name should be unique.		String

39.4.2. Query Parameters (17 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
cron (consumer)	Specifies a cron expression to define when to trigger.		String
deleteJob (consumer)	If set to true, then the trigger automatically delete when route stop. Else if set to false, it will remain in scheduler. When set to false, it will also mean user may reuse pre-configured trigger with camel Uri. Just ensure the names match. Notice you cannot have both <code>deleteJob</code> and <code>pauseJob</code> set to true.	true	boolean
durableJob (consumer)	Whether or not the job should remain stored after it is orphaned (no triggers point to it).	false	boolean
pauseJob (consumer)	If set to true, then the trigger automatically pauses when route stop. Else if set to false, it will remain in scheduler. When set to false, it will also mean user may reuse pre-configured trigger with camel Uri. Just ensure the names match. Notice you cannot have both <code>deleteJob</code> and <code>pauseJob</code> set to true.	false	boolean
recoverableJob (consumer)	Instructs the scheduler whether or not the job should be re-executed if a 'recovery' or 'fail-over' situation is encountered.	false	boolean
stateful (consumer)	Uses a Quartz <code>PersistJobDataAfterExecution</code> and <code>DisallowConcurrentExecution</code> instead of the default job.	false	boolean

Name	Description	Default	Type
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
customCalendar (advanced)	Specifies a custom calendar to avoid specific range of date.		Calendar
jobParameters (advanced)	To configure additional options on the job.		Map
prefixJobNameWithEndpointId (advanced)	Whether the job name should be prefixed with endpoint id.	false	boolean
triggerParameters (advanced)	To configure additional options on the trigger.		Map
usingFixedCamelContextName (advanced)	If it is true, JobDataMap uses the CamelContext name directly to reference the CamelContext, if it is false, JobDataMap uses use the CamelContext management name which could be changed during the deploy time.	false	boolean
autoStartScheduler (scheduler)	Whether or not the scheduler should be auto started.	true	boolean
startDelayedSeconds (scheduler)	Seconds to wait before starting the quartz scheduler.		int
triggerStartDelay (scheduler)	In case of scheduler has already started, we want the trigger start slightly after current time to ensure endpoint is fully started before the job kicks in. Negative value shifts trigger start time in the past.	500	long

39.4.3. Configuring quartz.properties file

By default Quartz will look for a **quartz.properties** file in the **org/quartz** directory of the classpath. If you are using WAR deployments this means just drop the quartz.properties in **WEB-INF/classes/org/quartz**.

However the Camel [Quartz](#) component also allows you to configure properties:

Parameter	Default	Type	Description
properties	null	Properties	You can configure a java.util.Properties instance.
propertiesFile	null	String	File name of the properties to load from the classpath

To do this you can configure this in Spring XML as follows

```
<bean id="quartz" class="org.apache.camel.component.quartz.QuartzComponent">
  <property name="propertiesFile" value="com/mycompany/myquartz.properties"/>
</bean>
```

39.5. ENABLING QUARTZ SCHEDULER IN JMX

You need to configure the quartz scheduler properties to enable JMX.

That is typically setting the option **"org.quartz.scheduler.jmx.export"** to a **true** value in the configuration file.

This option is set to true by default, unless explicitly disabled.

39.6. STARTING THE QUARTZ SCHEDULER

The [Quartz](#) component offers an option to let the Quartz scheduler be started delayed, or not auto started at all.

This is an example:

```
<bean id="quartz" class="org.apache.camel.component.quartz.QuartzComponent">
  <property name="startDelayedSeconds" value="5"/>
</bean>
```

39.7. CLUSTERING

If you use Quartz in clustered mode, e.g. the **JobStore** is clustered. Then the [Quartz](#) component will not pause/remove triggers when a node is being stopped/shutdown. This allows the trigger to keep running on the other nodes in the cluster.



NOTE

When running in clustered node no checking is done to ensure unique job name/group for endpoints.

39.8. MESSAGE HEADERS

Camel adds the getters from the Quartz Execution Context as header values. The following headers are added:

calendar, fireTime, jobDetail, jobInstance, jobRunTime, mergedJobDataMap, nextFireTime, previousFireTime, refireCount, result, scheduledFireTime, scheduler, trigger, triggerName, triggerGroup.

The **fireTime** header contains the **java.util.Date** of when the exchange was fired.

39.9. USING CRON TRIGGERS

Quartz supports [Cron-like expressions](#) for specifying timers in a handy format. You can use these expressions in the **cron** URI parameter; though to preserve valid URI encoding we allow **+** to be used instead of spaces.

For example, the following will fire a message every five minutes starting at 12pm (noon) to 6pm on weekdays:

```
from("quartz://myGroup/myTimerName?cron=0+0/5+12-18+?+*+MON-FRI")
.to("activemq:Totally.Rocks");
```

which is equivalent to using the cron expression

```
0 0/5 12-18 ? * MON-FRI
```

The following table shows the URI character encodings we use to preserve valid URI syntax:

URI Character	Cron character
+	<i>Space</i>

39.10. SPECIFYING TIME ZONE

The Quartz Scheduler allows you to configure time zone per trigger. For example to use a timezone of your country, then you can do as follows:

```
quartz://groupName/timerName?cron=0+0/5+12-18+?+*+MON-FRI&trigger.timeZone=Europe/Stockholm
```

The **timeZone** value is the values accepted by **java.util.TimeZone**.

39.11. CONFIGURING MISFIRE INSTRUCTIONS

The quartz scheduler can be configured with a misfire instruction to handle misfire situations for the trigger. The concrete trigger type that you are using will have defined a set of additional **MISFIRE_INSTRUCTION_XXX** constants that may be set as this property's value.

For example to configure the simple trigger to use misfire instruction 4:

```
quartz://myGroup/myTimerName?trigger.repeatInterval=2000&trigger.misfireInstruction=4
```

And likewise you can configure the cron trigger with one of its misfire instructions as well:

```
quartz://myGroup/myTimerName?cron=0/2+*+*+*+*+*?&trigger.misfireInstruction=2
```

The simple and cron triggers has the following misfire instructions representative:

39.11.1. SimpleTrigger.MISFIRE_INSTRUCTION_FIRE_NOW = 1 (default)

Instructs the Scheduler that upon a mis-fire situation, the SimpleTrigger wants to be fired now by Scheduler.

This instruction should typically only be used for 'one-shot' (non-repeating) Triggers. If it is used on a trigger with a repeat count > 0 then it is equivalent to the instruction MISFIRE_INSTRUCTION_RESCHEDULE_NOW_WITH_REMAINING_REPEAT_COUNT.

39.11.2. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NOW_WITH_EXISTING_REPEAT_COUNT = 2

Instructs the Scheduler that upon a mis-fire situation, the SimpleTrigger wants to be re-scheduled to 'now' (even if the associated Calendar excludes 'now') with the repeat count left as-is. This does obey the Trigger end-time however, so if 'now' is after the end-time the Trigger will not fire again.

Use of this instruction causes the trigger to 'forget' the start-time and repeat-count that it was originally setup with (this is only an issue if you for some reason wanted to be able to tell what the original values were at some later time).

39.11.3. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NOW_WITH_REMAINING_REPEAT_COUNT = 3

Instructs the Scheduler that upon a mis-fire situation, the SimpleTrigger wants to be re-scheduled to 'now' (even if the associated Calendar excludes 'now') with the repeat count set to what it would be, if it had not missed any firings. This does obey the Trigger end-time however, so if 'now' is after the end-time the Trigger will not fire again.

Use of this instruction causes the trigger to 'forget' the start-time and repeat-count that it was originally setup with. Instead, the repeat count on the trigger will be changed to whatever the remaining repeat count is (this is only an issue if you for some reason wanted to be able to tell what the original values were at some later time).

This instruction could cause the Trigger to go to the 'COMPLETE' state after firing 'now', if all the repeat-fire-times were missed.

39.11.4. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NEXT_WITH_REMAINING_REPEAT_COUNT = 4

Instructs the Scheduler that upon a mis-fire situation, the SimpleTrigger wants to be re-scheduled to the next scheduled time after 'now' - taking into account any associated Calendar and with the repeat count set to what it would be, if it had not missed any firings.



NOTE

This instruction could cause the Trigger to go directly to the 'COMPLETE' state if all fire-times were missed.

39.11.5. SimpleTrigger.MISFIRE_INSTRUCTION_RESCHEDULE_NEXT_WITH_EXISTING_COUNT = 5

Instructs the Scheduler that upon a mis-fire situation, the SimpleTrigger wants to be re-scheduled to the next scheduled time after 'now' - taking into account any associated Calendar, and with the repeat count left unchanged.



NOTE

This instruction could cause the Trigger to go directly to the 'COMPLETE' state if the end-time of the trigger has arrived.

39.11.6. CronTrigger.MISFIRE_INSTRUCTION_FIRE_ONCE_NOW = 1 (default)

Instructs the Scheduler that upon a mis-fire situation, the CronTrigger wants to be fired now by Scheduler.

39.11.7. CronTrigger.MISFIRE_INSTRUCTION_DO_NOTHING = 2

Instructs the Scheduler that upon a mis-fire situation, the CronTrigger wants to have its next-fire-time updated to the next time in the schedule after the current time (taking into account any associated Calendar but it does not want to be fired now).

39.12. USING QUARTZSCHEDULEDPOLLCONSUMERSCHEDULER

The [Quartz](#) component provides a Polling Consumer scheduler which allows to use cron based scheduling for Polling Consumer such as the File and FTP consumers.

For example to use a cron based expression to poll for files every 2nd second, then a Camel route can be define simply as:

```
from("file:inbox?scheduler=quartz&scheduler.cron=0/2+*+*+*+*+*?")
    .to("bean:process");
```

Notice we define the **scheduler=quartz** to instruct Camel to use the [Quartz](#) based scheduler. Then we use **scheduler.xxx** options to configure the scheduler. The [Quartz](#) scheduler requires the cron option to be set.

The following options is supported:

Parameter	Default	Type	Description
quartzScheduler	null	org.quartz.Scheduler	To use a custom Quartz scheduler. If none configure then the shared scheduler from the component is used.
cron	null	String	Mandatory: To define the cron expression for triggering the polls.
triggerId	null	String	To specify the trigger id. If none provided then an UUID is generated and used.

Parameter	Default	Type	Description
triggerGroup	QuartzScheduledPollConsumerScheduler	String	To specify the trigger group.
timeZone	Default	Timezone	The time zone to use for the CRON trigger.



IMPORTANT

Remember configuring these options from the endpoint URIs must be prefixed with **scheduler**.

For example to configure the trigger id and group:

```
from("file:inbox?scheduler=quartz&scheduler.cron=0/2+*+*+*+*+?
&scheduler.triggerId=myId&scheduler.triggerGroup=myGroup")
.to("bean:process");
```

There is also a CRON scheduler in Spring, so you can use the following as well:

```
from("file:inbox?scheduler=spring&scheduler.cron=0/2+*+*+*+*+?")
.to("bean:process");
```

39.13. CRON COMPONENT SUPPORT

The Quartz component can be used as implementation of the Camel Cron component.

Maven users will need to add the following additional dependency to their **pom.xml**:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-cron</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

Users can then use the cron component instead of the quartz component, as in the following route:

```
from("cron://name?schedule=0+0/5+12-18+?+*+MON-FRI")
.to("activemq:Totally.Rocks");
```

39.14. SPRING BOOT AUTO-CONFIGURATION

When using quartz with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-quartz-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 14 options, which are listed below.

Name	Description	Default	Type
camel.component.quartz.auto-start-scheduler	Whether or not the scheduler should be auto started. This options is default true.	true	Boolean
camel.component.quartz.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.quartz.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.quartz.enable-jmx	Whether to enable Quartz JMX which allows to manage the Quartz scheduler from JMX. This options is default true.	true	Boolean
camel.component.quartz.enabled	Whether to enable auto configuration of the quartz component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.quartz.interrupt-jobs-on-shutdown</code>	Whether to interrupt jobs on shutdown which forces the scheduler to shutdown quicker and attempt to interrupt any running jobs. If this is enabled then any running jobs can fail due to being interrupted. When a job is interrupted then Camel will mark the exchange to stop continue routing and set <code>java.util.concurrent.RejectedExecutionException</code> as caused exception. Therefore use this with care, as its often better to allow Camel jobs to complete and shutdown gracefully.	false	Boolean
<code>camel.component.quartz.prefix-instance-name</code>	Whether to prefix the Quartz Scheduler instance name with the CamelContext name. This is enabled by default, to let each CamelContext use its own Quartz scheduler instance by default. You can set this option to false to reuse Quartz scheduler instances between multiple CamelContext's.	true	Boolean
<code>camel.component.quartz.prefix-job-name-with-endpoint-id</code>	Whether to prefix the quartz job with the endpoint id. This option is default false.	false	Boolean
<code>camel.component.quartz.properties</code>	Properties to configure the Quartz scheduler.		Map
<code>camel.component.quartz.properties-file</code>	File name of the properties to load from the classpath.		String
<code>camel.component.quartz.properties-ref</code>	References to an existing Properties or Map to lookup in the registry to use for configuring quartz.		String
<code>camel.component.quartz.scheduler</code>	To use the custom configured Quartz scheduler, instead of creating a new Scheduler. The option is a <code>org.quartz.Scheduler</code> type.		Scheduler
<code>camel.component.quartz.scheduler-factory</code>	To use the custom SchedulerFactory which is used to create the Scheduler. The option is a <code>org.quartz.SchedulerFactory</code> type.		SchedulerFactory
<code>camel.component.quartz.start-delayed-seconds</code>	Seconds to wait before starting the quartz scheduler.		Integer

CHAPTER 40. REF

Both producer and consumer are supported

The Ref component is used for lookup of existing endpoints bound in the Registry.

40.1. URI FORMAT

```
ref:someName[?options]
```

Where **someName** is the name of an endpoint in the Registry (usually, but not always, the Spring registry). If you are using the Spring registry, **someName** would be the bean ID of an endpoint in the Spring registry.

40.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

40.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

40.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

40.3. COMPONENT OPTIONS

The Ref component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

40.4. ENDPOINT OPTIONS

The Ref endpoint is configured using URI syntax:

```
ref:name
```

with the following path and query parameters:

40.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required Name of endpoint to lookup in the registry.		String

40.4.2. Query Parameters (4 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● <code>InOnly</code> ● <code>InOut</code> ● <code>InOptionalOut</code> 		<code>ExchangePattern</code>
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

40.5. RUNTIME LOOKUP

This component can be used when you need dynamic discovery of endpoints in the Registry where you can compute the URI at runtime. Then you can look up the endpoint using the following code:

```
// lookup the endpoint
String myEndpointRef = "bigspenderOrder";
Endpoint endpoint = context.getEndpoint("ref:" + myEndpointRef);

Producer producer = endpoint.createProducer();
Exchange exchange = producer.createExchange();
```

```
exchange.getIn().setBody(payloadToSend);
// send the exchange
producer.process(exchange);
```

And you could have a list of endpoints defined in the Registry such as:

```
<camelContext id="camel" xmlns="http://activemq.apache.org/camel/schema/spring">
  <endpoint id="normalOrder" uri="activemq:order.slow"/>
  <endpoint id="bigspenderOrder" uri="activemq:order.high"/>
</camelContext>
```

40.6. SAMPLE

In the sample below we use the **ref:** in the URI to reference the endpoint with the spring ID, **endpoint2:**

You could, of course, have used the **ref** attribute instead:

```
<to uri="ref:endpoint2"/>
```

Which is the more common way to write it.

40.7. SPRING BOOT AUTO-CONFIGURATION

When using ref with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-ref-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.ref.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.ref.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.ref.enabled	Whether to enable auto configuration of the ref component. This is enabled by default.		Boolean
camel.component.ref.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 41. REST

Both producer and consumer are supported

The REST component allows to define REST endpoints (consumer) using the Rest DSL and plugin to other Camel components as the REST transport.

The rest component can also be used as a client (producer) to call REST services.

41.1. URI FORMAT

```
rest://method:path[:uriTemplate]?[options]
```

41.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

41.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

41.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

41.3. COMPONENT OPTIONS

The REST component supports 8 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
consumerComponentName (consumer)	The Camel Rest component to use for (consumer) the REST transport, such as jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestConsumerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
apiDoc (producer)	The swagger api doc resource to use. The resource is loaded from classpath by default and must be in JSON format.		String
componentName (producer)	Deprecated The Camel Rest component to use for (producer) the REST transport, such as http, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestProducerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
host (producer)	Host and port of HTTP service to use (override host in swagger schema).		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
producerComponentName (producer)	The Camel Rest component to use for (producer) the REST transport, such as http, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestProducerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

41.4. ENDPOINT OPTIONS

The REST endpoint is configured using URI syntax:

```
rest:method:path:uriTemplate
```

with the following path and query parameters:

41.4.1. Path Parameters (3 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
method (common)	<p>Required HTTP method to use.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • get • post • put • delete • patch • head • trace • connect • options 		String
path (common)	Required The base path.		String
uriTemplate (common)	The uri template.		String

41.4.2. Query Parameters (16 parameters)

Name	Description	Default	Type
consumes (common)	Media type such as: 'text/xml', or 'application/json' this REST service accepts. By default we accept all kinds of types.		String
inType (common)	To declare the incoming POJO binding type as a FQN class name.		String
outType (common)	To declare the outgoing POJO binding type as a FQN class name.		String
produces (common)	Media type such as: 'text/xml', or 'application/json' this REST service returns.		String
routeId (common)	Name of the route this REST services creates.		String

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
consumerComponentName (consumer)	The Camel Rest component to use for (consumer) the REST transport, such as jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestConsumerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
description (consumer)	Human description to document this REST service.		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● <code>InOnly</code> ● <code>InOut</code> ● <code>InOptionalOut</code> 		<code>ExchangePattern</code>
apiDoc (producer)	The openapi api doc resource to use. The resource is loaded from classpath by default and must be in JSON format.		String

Name	Description	Default	Type
bindingMode (producer)	<p>Configures the binding mode for the producer. If set to anything other than 'off' the producer will try to convert the body of the incoming message from inType to the json or xml, and the response from json or xml to outType.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • auto • off • json • xml • json_xml 		RestBindingMode
host (producer)	Host and port of HTTP service to use (override host in openapi schema).		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
producerComponentName (producer)	The Camel Rest component to use for (producer) the REST transport, such as http, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestProducerFactory is registered in the registry. If either one is found, then that is being used.		String
queryParameters (producer)	Query parameters for the HTTP service to call. The query parameters can contain multiple parameters separated by ampersand such such as foo=123&bar=456.		String

41.5. SUPPORTED REST COMPONENTS

The following components support rest consumer (Rest DSL):

- camel-servlet

The following components support rest producer:

- camel-http

41.6. PATH AND URITEMPLATE SYNTAX

The path and uriTemplate option is defined using a REST syntax where you define the REST context path using support for parameters.



NOTE

If no uriTemplate is configured then path option works the same way. It does not matter if you configure only path or if you configure both options. Though configuring both a path and uriTemplate is a more common practice with REST.

The following is a Camel route using a path only

```
from("rest:get:hello")
  .transform().constant("Bye World");
```

And the following route uses a parameter which is mapped to a Camel header with the key "me".

```
from("rest:get:hello/{me}")
  .transform().simple("Bye ${header.me}");
```

The following examples have configured a base path as "hello" and then have two REST services configured using uriTemplates.

```
from("rest:get:hello/{me}")
  .transform().simple("Hi ${header.me}");

from("rest:get:hello/french/{me}")
  .transform().simple("Bonjour ${header.me}");
```

41.7. REST PRODUCER EXAMPLES

You can use the rest component to call REST services like any other Camel component.

For example to call a REST service on using **hello/{me}** you can do

```
from("direct:start")
  .to("rest:get:hello/{me}");
```

And then the dynamic value **{me}** is mapped to Camel message with the same name. So to call this REST service you can send an empty message body and a header as shown:

```
template.sendBodyAndHeader("direct:start", null, "me", "Donald Duck");
```

The Rest producer needs to know the hostname and port of the REST service, which you can configure using the host option as shown:


```
from("direct:start")
  .to("rest:get:hello/{me}?host=myserver:8080/foo");
```

Instead of using the host option, you can configure the host on the **restConfiguration** as shown:

```
restConfiguration().host("myserver:8080/foo");

from("direct:start")
  .to("rest:get:hello/{me}");
```

You can use the **producerComponent** to select which Camel component to use as the HTTP client, for example to use http you can do:

```
restConfiguration().host("myserver:8080/foo").producerComponent("http");

from("direct:start")
  .to("rest:get:hello/{me}");
```

41.8. REST PRODUCER BINDING

The REST producer supports binding using JSON or XML like the rest-dsl does.

For example to use jetty with json binding mode turned on you can configure this in the rest configuration:

```
restConfiguration().component("jetty").host("localhost").port(8080).bindingMode(RestBindingMode.json)
;

from("direct:start")
  .to("rest:post:user");
```

Then when calling the REST service using rest producer it will automatic bind any POJOs to json before calling the REST service:

```
UserPojo user = new UserPojo();
user.setId(123);
user.setName("Donald Duck");

template.sendBody("direct:start", user);
```

In the example above we send a POJO instance **UserPojo** as the message body. And because we have turned on JSON binding in the rest configuration, then the POJO will be marshalled from POJO to JSON before calling the REST service.

However if you want to also perform binding for the response message (eg what the REST service send back as response) you would need to configure the **outType** option to specify what is the classname of the POJO to unmarshal from JSON to POJO.

For example if the REST service returns a JSON payload that binds to **com.foo.MyResponsePojo** you can configure this as shown:

```
restConfiguration().component("jetty").host("localhost").port(8080).bindingMode(RestBindingMode.json)
```

```

;
from("direct:start")
.to("rest:post:user?outType=com.foo.MyResponsePojo");

```



NOTE

You must configure **outType** option if you want POJO binding to happen for the response messages received from calling the REST service.

41.9. MORE EXAMPLES

See Rest DSL which offers more examples and how you can use the Rest DSL to define those in a nicer RESTful way.

There is a **camel-example-servlet-rest-tomcat** example in the Apache Camel distribution, that demonstrates how to use the Rest DSL with SERVLET as transport that can be deployed on Apache Tomcat, or similar web containers.

41.10. SPRING BOOT AUTO-CONFIGURATION

When using rest with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-rest-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 12 options, which are listed below.

Name	Description	Default	Type
camel.component.rest-api.automated-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
camel.component.rest-api-bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.rest-api.enabled	Whether to enable auto configuration of the rest-api component. This is enabled by default.		Boolean
camel.component.rest-api-doc	The swagger api doc resource to use. The resource is loaded from classpath by default and must be in JSON format.		String
camel.component.rest.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.rest.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.rest.consumer-component-name	The Camel Rest component to use for (consumer) the REST transport, such as jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestConsumerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
camel.component.rest.enabled	Whether to enable auto configuration of the rest component. This is enabled by default.		Boolean
camel.component.rest.host	Host and port of HTTP service to use (override host in swagger schema).		String

Name	Description	Default	Type
<code>camel.component.rest.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.rest.producer-component-name</code>	The Camel Rest component to use for (producer) the REST transport, such as http, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestProducerFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.component.rest.component-name</code>	Deprecated The Camel Rest component to use for (producer) the REST transport, such as http, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a <code>org.apache.camel.spi.RestProducerFactory</code> is registered in the registry. If either one is found, then that is being used.		String

CHAPTER 42. SAGA

Only producer is supported

The Saga component provides a bridge to execute custom actions within a route using the Saga EIP.

The component should be used for advanced tasks, such as deciding to complete or compensate a Saga with `completionMode` set to **MANUAL**.

Refer to the Saga EIP documentation for help on using sagas in common scenarios.

42.1. URI FORMAT

saga:action

42.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

42.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (`application.properties|yaml`), or directly with Java code.

42.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

42.3. COMPONENT OPTIONS

The Saga component supports 2 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

42.4. ENDPOINT OPTIONS

The Saga endpoint is configured using URI syntax:

```
saga:action
```

with the following path and query parameters:

42.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
action (producer)	Required Action to execute (complete or compensate). Enum values: <ul style="list-style-type: none"> ● COMPLETE ● COMPENSATE 		SagaEndpointAction

42.4.2. Query Parameters (1 parameters)

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

42.5. SPRING BOOT AUTO-CONFIGURATION

When using saga with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-saga-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
camel.component.saga.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.saga.enabled	Whether to enable auto configuration of the saga component. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.component.saga.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 43. SALESFORCE

Both producer and consumer are supported

This component supports producer and consumer endpoints to communicate with Salesforce using Java DTOs.

There is a companion maven plugin Camel Salesforce Plugin that generates these DTOs (see further below).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-salesforce</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```



NOTE

Developers wishing to contribute to the component are instructed to look at the [README.md](#) file on instructions on how to get started and setup your environment for running integration tests.

43.1. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

43.1.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

43.1.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

43.2. COMPONENT OPTIONS

The Salesforce component supports 90 options, which are listed below.

Name	Description	Default	Type
apexMethod (common)	APEX method name.		String
apexQueryParams (common)	Query params for APEX method.		Map
apiVersion (common)	Salesforce API version.	53.0	String
backoffIncrement (common)	Backoff interval increment for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	1000	long
batchId (common)	Bulk API Batch ID.		String
contentType (common)	Bulk API content type, one of XML, CSV, ZIP_XML, ZIP_CSV. Enum values: <ul style="list-style-type: none"> ● XML ● CSV ● JSON ● ZIP_XML ● ZIP_CSV ● ZIP_JSON 		ContentType
defaultReplayId (common)	Default replayId setting if no value is found in initialReplayIdMap.	-1	Long
fallBackReplayId (common)	ReplayId to fall back to after an Invalid Replay Id response.	-1	Long

Name	Description	Default	Type
format (common)	<p>Payload format to use for Salesforce API calls, either JSON or XML, defaults to JSON. As of Camel 3.12, this option only applies to the Raw operation.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • JSON • XML 		PayloadFormat
httpClient (common)	Custom Jetty Http Client to use to connect to Salesforce.		SalesforceHttpClient
httpClientConnectionTimeout (common)	Connection timeout used by the HttpClient when connecting to the Salesforce server.	60000	long
httpClientIdleTimeout (common)	Timeout used by the HttpClient when waiting for response from the Salesforce server.	10000	long
httpClientMaxContentLength (common)	Max content length of an HTTP response.		Integer
httpClientRequestBufferSize (common)	HTTP request buffer size. May need to be increased for large SOQL queries.	8192	Integer
includeDetails (common)	Include details in Salesforce Analytics report, defaults to false.		Boolean
initialReplayIdMap (common)	Replay IDs to start from per channel name.		Map
instanceId (common)	Salesforce Analytics report execution instance ID.		String
jobId (common)	Bulk API Job ID.		String
limit (common)	Limit on number of returned records. Applicable to some of the API, check the Salesforce documentation.		Integer
locator (common)	Locator provided by Salesforce Bulk 2.0 API for use in getting results for a Query job.		String

Name	Description	Default	Type
maxBackoff (common)	Maximum backoff interval for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	30000	long
maxRecords (common)	The maximum number of records to retrieve per set of results for a Bulk 2.0 Query. The request is still subject to the size limits. If you are working with a very large number of query results, you may experience a timeout before receiving all the data from Salesforce. To prevent a timeout, specify the maximum number of records your client is expecting to receive in the maxRecords parameter. This splits the results into smaller sets with this value as the maximum size.		Integer
notFoundBehaviour (common)	Sets the behaviour of 404 not found status received from Salesforce API. Should the body be set to NULL NotFoundBehaviour#NULL or should an exception be signaled on the exchange NotFoundBehaviour#EXCEPTION - the default. Enum values: <ul style="list-style-type: none"> ● EXCEPTION ● NULL 	EXCEPTION	NotFoundBehaviour
notifyForFields (common)	Notify for fields, options are ALL, REFERENCED, SELECT, WHERE. Enum values: <ul style="list-style-type: none"> ● ALL ● REFERENCED ● SELECT ● WHERE 		NotifyForFieldsEnum
notifyForOperationCreate (common)	Notify for create operation, defaults to false (API version = 29.0).		Boolean
notifyForOperationDelete (common)	Notify for delete operation, defaults to false (API version = 29.0).		Boolean

Name	Description	Default	Type
notifyForOperations (common)	<p>Notify for operations, options are ALL, CREATE, EXTENDED, UPDATE (API version 29.0).</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● ALL ● CREATE ● EXTENDED ● UPDATE 		NotifyForOperationsEnum
notifyForOperationUndelete (common)	Notify for un-delete operation, defaults to false (API version = 29.0).		Boolean
notifyForOperationUpdate (common)	Notify for update operation, defaults to false (API version = 29.0).		Boolean
objectMapper (common)	Custom Jackson ObjectMapper to use when serializing/deserializing Salesforce objects.		ObjectMapper
packages (common)	In what packages are the generated DTO classes. Typically the classes would be generated using camel-salesforce-maven-plugin. Set it if using the generated DTOs to gain the benefit of using short SObject names in parameters/header values. Multiple packages can be separated by comma.		String
pkChunking (common)	Use PK Chunking. Only for use in original Bulk API. Bulk 2.0 API performs PK chunking automatically, if necessary.		Boolean
pkChunkingChunkSize (common)	Chunk size for use with PK Chunking. If unspecified, salesforce default is 100,000. Maximum size is 250,000.		Integer
pkChunkingParent (common)	Specifies the parent object when you're enabling PK chunking for queries on sharing objects. The chunks are based on the parent object's records rather than the sharing object's records. For example, when querying on AccountShare, specify Account as the parent object. PK chunking is supported for sharing objects as long as the parent object is supported.		String

Name	Description	Default	Type
pkChunkingStartRow (common)	Specifies the 15-character or 18-character record ID to be used as the lower boundary for the first chunk. Use this parameter to specify a starting ID when restarting a job that failed between batches.		String
queryLocator (common)	Query Locator provided by salesforce for use when a query results in more records than can be retrieved in a single call. Use this value in a subsequent call to retrieve additional records.		String
rawPayload (common)	Use raw payload String for request and response (either JSON or XML depending on format), instead of DTOs, false by default.	false	boolean
reportId (common)	Salesforce1 Analytics report Id.		String
reportMetadata (common)	Salesforce1 Analytics report metadata for filtering.		ReportMetadata
resultId (common)	Bulk API Result ID.		String
sObjectBlobFieldName (common)	SObject blob field name.		String
sObjectClass (common)	Fully qualified SObject class name, usually generated using camel-salesforce-maven-plugin.		String
sObjectFields (common)	SObject fields to retrieve.		String
sObjectId (common)	SObject ID if required by API.		String
sObjectIdName (common)	SObject external ID field name.		String
sObjectIdValue (common)	SObject external ID field value.		String
sObjectName (common)	SObject name if required or supported by API.		String
sObjectQuery (common)	Salesforce SOQL query string.		String

Name	Description	Default	Type
sObjectSearch (common)	Salesforce SOSL search string.		String
updateTopic (common)	Whether to update an existing Push Topic when using the Streaming API, defaults to false.	false	boolean
config (common (advanced))	Global endpoint configuration - use to set values that are common to all endpoints.		SalesforceEndpointConfig
httpClientProperties (common (advanced))	Used to set any properties that can be configured on the underlying HTTP client. Have a look at properties of SalesforceHttpClient and the Jetty HttpClient for all available options.		Map
longPollingTransportProperties (common (advanced))	Used to set any properties that can be configured on the LongPollingTransport used by the BayeuxClient (CometD) used by the streaming api.		Map
workerPoolMaxSize (common (advanced))	Maximum size of the thread pool used to handle HTTP responses.	20	int
workerPoolSize (common (advanced))	Size of the thread pool used to handle HTTP responses.	10	int
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
allOrNone (producer)	Composite API option to indicate to rollback all records if any are not successful.	false	boolean
apexUrl (producer)	APEX method URL.		String
compositeMethod (producer)	Composite (raw) method.		String

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
rawHttpHeaders (producer)	Comma separated list of message headers to include as HTTP parameters for Raw operation.		String
rawMethod (producer)	HTTP method to use for the Raw operation.		String
rawPath (producer)	The portion of the endpoint URL after the domain name. E.g., <code>'/services/data/v52.0/subjects/Account/'</code> .		String
rawQueryParameters (producer)	Comma separated list of message headers to include as query parameters for Raw operation. Do not url-encode values as this will be done automatically.		String
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
httpProxyExcludedAddresses (proxy)	A list of addresses for which HTTP proxy server should not be used.		Set
httpProxyHost (proxy)	Hostname of the HTTP proxy server to use.		String
httpProxyIncludedAddresses (proxy)	A list of addresses for which HTTP proxy server should be used.		Set
httpProxyPort (proxy)	Port number of the HTTP proxy server to use.		Integer

Name	Description	Default	Type
httpProxySocks4 (proxy)	If set to true the configures the HTTP proxy to use as a SOCKS4 proxy.	false	boolean
authenticationType (security)	Explicit authentication method to be used, one of USERNAME_PASSWORD, REFRESH_TOKEN or JWT. Salesforce component can auto-determine the authentication method to use from the properties set, set this property to eliminate any ambiguity. Enum values: <ul style="list-style-type: none"> ● USERNAME_PASSWORD ● REFRESH_TOKEN ● JWT 		AuthenticationType
clientId (security)	Required OAuth Consumer Key of the connected app configured in the Salesforce instance setup. Typically a connected app needs to be configured but one can be provided by installing a package.		String
clientSecret (security)	OAuth Consumer Secret of the connected app configured in the Salesforce instance setup.		String
httpProxyAuthUri (security)	Used in authentication against the HTTP proxy server, needs to match the URI of the proxy server in order for the httpProxyUsername and httpProxyPassword to be used for authentication.		String
httpProxyPassword (security)	Password to use to authenticate against the HTTP proxy server.		String
httpProxyRealm (security)	Realm of the proxy server, used in preemptive Basic/Digest authentication methods against the HTTP proxy server.		String
httpProxySecure (security)	If set to false disables the use of TLS when accessing the HTTP proxy.	true	boolean
httpProxyUseDigestAuth (security)	If set to true Digest authentication will be used when authenticating to the HTTP proxy, otherwise Basic authorization method will be used.	false	boolean
httpProxyUsername (security)	Username to use to authenticate against the HTTP proxy server.		String

Name	Description	Default	Type
instanceUrl (security)	URL of the Salesforce instance used after authentication, by default received from Salesforce on successful authentication.		String
jwtAudience (security)	Value to use for the Audience claim (aud) when using OAuth JWT flow. If not set, the login URL will be used, which is appropriate in most cases.		String
keystore (security)	KeyStore parameters to use in OAuth JWT flow. The KeyStore should contain only one entry with private key and certificate. Salesforce does not verify the certificate chain, so this can easily be a selfsigned certificate. Make sure that you upload the certificate to the corresponding connected app.		KeyStoreParameters
lazyLogin (security)	If set to true prevents the component from authenticating to Salesforce with the start of the component. You would generally set this to the (default) false and authenticate early and be immediately aware of any authentication issues.	false	boolean
loginConfig (security)	All authentication configuration in one nested bean, all properties set there can be set directly on the component as well.		SalesforceLoginConfig
loginUrl (security)	Required URL of the Salesforce instance used for authentication, by default set to https://login.salesforce.com .	https://login.salesforce.com	String
password (security)	Password used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows. Make sure that you append security token to the end of the password if using one.		String

Name	Description	Default	Type
refreshToken (security)	Refresh token already obtained in the refresh token OAuth flow. One needs to setup a web application and configure a callback URL to receive the refresh token, or configure using the builtin callback at https://login.salesforce.com/services/oauth2/success or https://test.salesforce.com/services/oauth2/success and then retrieve the refresh_token from the URL at the end of the flow. Note that in development organizations Salesforce allows hosting the callback web application at localhost.		String
sslContextParameters (security)	SSL parameters to use, see SSLContextParameters class for all available options.		SSLContextParameters
useGlobalSslContextParameters (security)	Enable usage of global SSL context parameters.	false	boolean
userName (security)	Username used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows.		String

43.3. ENDPOINT OPTIONS

The Salesforce endpoint is configured using URI syntax:

```
salesforce:operationName:topicName
```

with the following path and query parameters:

43.3.1. Path Parameters (2 parameters)

Name	Description	Default	Type
operationName (producer)	The operation to use. Enum values: <ul style="list-style-type: none"> • getVersions • getResources • getGlobalObjects 		OperationName

Name	Description	Default	Type
	<ul style="list-style-type: none"> ● getBasicInfo ● getDescription ● getObject ● createObject ● updateObject ● deleteObject ● getObjectWithId ● upsertObject ● deleteObjectWithId ● getBlobField ● query ● queryMore ● queryAll ● search ● apexCall ● recent ● createJob ● getJob ● closeJob ● abortJob ● createBatch ● getBatch ● getAllBatches ● getRequest ● getResults ● createBatchQuery ● getQueryResultIds ● getQueryResult ● getRecentReports ● getReportDescription ● executeSyncReport ● executeAsyncReport ● getReportInstances 		

Name	Description <ul style="list-style-type: none"> ● getReportResults ● limits 	Default	Type
	<ul style="list-style-type: none"> ● approval ● approvals ● composite-tree ● composite-batch ● composite ● compositeRetrieveSObjectCollections ● compositeCreateSObjectCollections ● compositeUpdateSObjectCollections ● compositeUpsertSObjectCollections ● compositeDeleteSObjectCollections ● bulk2GetAllJobs ● bulk2CreateJob ● bulk2GetJob ● bulk2CreateBatch ● bulk2CloseJob ● bulk2AbortJob ● bulk2DeleteJob ● bulk2GetSuccessfulResults ● bulk2GetFailedResults ● bulk2GetUnprocessedRecords ● bulk2CreateQueryJob ● bulk2GetQueryJob ● bulk2GetAllQueryJobs ● bulk2GetQueryJobResults ● bulk2AbortQueryJob ● bulk2DeleteQueryJob ● raw 		
topicName (consumer)	The name of the topic/channel to use.		String

43.3.2. Query Parameters (57 parameters)

Name	Description	Default	Type
apexMethod (common)	APEX method name.		String
apexQueryParams (common)	Query params for APEX method.		Map
apiVersion (common)	Salesforce API version.	53.0	String
backoffIncrement (common)	Backoff interval increment for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	1000	long
batchId (common)	Bulk API Batch ID.		String
contentType (common)	Bulk API content type, one of XML, CSV, ZIP_XML, ZIP_CSV. Enum values: <ul style="list-style-type: none"> • XML • CSV • JSON • ZIP_XML • ZIP_CSV • ZIP_JSON 		ContentType
defaultReplayId (common)	Default replayId setting if no value is found in initialReplayIdMap.	-1	Long
fallBackReplayId (common)	ReplayId to fall back to after an Invalid Replay Id response.	-1	Long
format (common)	Payload format to use for Salesforce API calls, either JSON or XML, defaults to JSON. As of Camel 3.12, this option only applies to the Raw operation. Enum values: <ul style="list-style-type: none"> • JSON • XML 		PayloadFormat

Name	Description	Default	Type
httpClient (common)	Custom Jetty Http Client to use to connect to Salesforce.		SalesforceHttpClient
includeDetails (common)	Include details in Salesforce1 Analytics report, defaults to false.		Boolean
initialReplayIdMap (common)	Replay IDs to start from per channel name.		Map
instanceId (common)	Salesforce1 Analytics report execution instance ID.		String
jobId (common)	Bulk API Job ID.		String
limit (common)	Limit on number of returned records. Applicable to some of the API, check the Salesforce documentation.		Integer
locator (common)	Locator provided by salesforce Bulk 2.0 API for use in getting results for a Query job.		String
maxBackoff (common)	Maximum backoff interval for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	30000	long
maxRecords (common)	The maximum number of records to retrieve per set of results for a Bulk 2.0 Query. The request is still subject to the size limits. If you are working with a very large number of query results, you may experience a timeout before receiving all the data from Salesforce. To prevent a timeout, specify the maximum number of records your client is expecting to receive in the maxRecords parameter. This splits the results into smaller sets with this value as the maximum size.		Integer
notFoundBehaviour (common)	Sets the behaviour of 404 not found status received from Salesforce API. Should the body be set to NULL NotFoundBehaviour#NULL or should a exception be signaled on the exchange NotFoundBehaviour#EXCEPTION - the default. Enum values: <ul style="list-style-type: none"> ● EXCEPTION ● NULL 	EXCEPTION	NotFoundBehaviour

Name	Description	Default	Type
notifyForFields (common)	Notify for fields, options are ALL, REFERENCED, SELECT, WHERE. Enum values: <ul style="list-style-type: none"> ● ALL ● REFERENCED ● SELECT ● WHERE 		NotifyForFieldsEnum
notifyForOperationCreate (common)	Notify for create operation, defaults to false (API version = 29.0).		Boolean
notifyForOperationDelete (common)	Notify for delete operation, defaults to false (API version = 29.0).		Boolean
notifyForOperations (common)	Notify for operations, options are ALL, CREATE, EXTENDED, UPDATE (API version 29.0). Enum values: <ul style="list-style-type: none"> ● ALL ● CREATE ● EXTENDED ● UPDATE 		NotifyForOperationsEnum
notifyForOperationUndelete (common)	Notify for un-delete operation, defaults to false (API version = 29.0).		Boolean
notifyForOperationUpdate (common)	Notify for update operation, defaults to false (API version = 29.0).		Boolean
objectMapper (common)	Custom Jackson ObjectMapper to use when serializing/deserializing Salesforce objects.		ObjectMapper
pkChunking (common)	Use PK Chunking. Only for use in original Bulk API. Bulk 2.0 API performs PK chunking automatically, if necessary.		Boolean

Name	Description	Default	Type
pkChunkingChunkSize (common)	Chunk size for use with PK Chunking. If unspecified, salesforce default is 100,000. Maximum size is 250,000.		Integer
pkChunkingParent (common)	Specifies the parent object when you're enabling PK chunking for queries on sharing objects. The chunks are based on the parent object's records rather than the sharing object's records. For example, when querying on AccountShare, specify Account as the parent object. PK chunking is supported for sharing objects as long as the parent object is supported.		String
pkChunkingStartRow (common)	Specifies the 15-character or 18-character record ID to be used as the lower boundary for the first chunk. Use this parameter to specify a starting ID when restarting a job that failed between batches.		String
queryLocator (common)	Query Locator provided by salesforce for use when a query results in more records than can be retrieved in a single call. Use this value in a subsequent call to retrieve additional records.		String
rawPayload (common)	Use raw payload String for request and response (either JSON or XML depending on format), instead of DTOs, false by default.	false	boolean
reportId (common)	Salesforce1 Analytics report Id.		String
reportMetadata (common)	Salesforce1 Analytics report metadata for filtering.		ReportMetadata
resultId (common)	Bulk API Result ID.		String
sObjectBlobFieldName (common)	SObject blob field name.		String
sObjectClass (common)	Fully qualified SObject class name, usually generated using camel-salesforce-maven-plugin.		String
sObjectFields (common)	SObject fields to retrieve.		String
sObjectId (common)	SObject ID if required by API.		String

Name	Description	Default	Type
sObjectIdName (common)	SObject external ID field name.		String
sObjectIdValue (common)	SObject external ID field value.		String
sObjectName (common)	SObject name if required or supported by API.		String
sObjectQuery (common)	Salesforce SOQL query string.		String
sObjectSearch (common)	Salesforce SOSL search string.		String
updateTopic (common)	Whether to update an existing Push Topic when using the Streaming API, defaults to false.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
replayId (consumer)	The replayId value to use when subscribing.		Long
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
allOrNone (producer)	Composite API option to indicate to rollback all records if any are not successful.	false	boolean
apexUrl (producer)	APEX method URL.		String
compositeMethod (producer)	Composite (raw) method.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
rawHttpHeaders (producer)	Comma separated list of message headers to include as HTTP parameters for Raw operation.		String
rawMethod (producer)	HTTP method to use for the Raw operation.		String
rawPath (producer)	The portion of the endpoint URL after the domain name. E.g., '/services/data/v52.0/subjects/Account/'.		String
rawQueryParameters (producer)	Comma separated list of message headers to include as query parameters for Raw operation. Do not url-encode values as this will be done automatically.		String

43.4. AUTHENTICATING TO SALESFORCE

The component supports three OAuth authentication flows:

- [OAuth 2.0 Username-Password Flow](#)
- [OAuth 2.0 Refresh Token Flow](#)
- [OAuth 2.0 JWT Bearer Token Flow](#)

For each of the flow different set of properties needs to be set:

Table 43.1. Table 1. Properties to set for each authentication flow

Property	Where to find it on Salesforce	Flow
clientId	Connected App, Consumer Key	All flows
clientSecret	Connected App, Consumer Secret	Username-Password, Refresh Token
userName	Salesforce user username	Username-Password, JWT Bearer Token
password	Salesforce user password	Username-Password
refreshToken	From OAuth flow callback	Refresh Token
keystore	Connected App, Digital Certificate	JWT Bearer Token

The component auto determines what flow you're trying to configure, to be remove ambiguity set the **authenticationType** property.

**NOTE**

Using Username-Password Flow in production is not encouraged.

**NOTE**

The certificate used in JWT Bearer Token Flow can be a selfsigned certificate. The KeyStore holding the certificate and the private key must contain only single certificate-private key entry.

43.5. URI FORMAT

When used as a consumer, receiving streaming events, the URI scheme is:

```
salesforce:topic?options
```

When used as a producer, invoking the Salesforce REST APIs, the URI scheme is:

```
salesforce:operationName?options
```

43.6. PASSING IN SALESFORCE HEADERS AND FETCHING SALESFORCE RESPONSE HEADERS

There is support to pass [Salesforce headers](#) via inbound message headers, header names that start with **Sforce** or **x-sfdc** on the Camel message will be passed on in the request, and response headers that start with **Sforce** will be present in the outbound message headers.

For example to fetch API limits you can specify:

```

// in your Camel route set the header before Salesforce endpoint
//...
.setHeader("Sforce-Limit-Info", constant("api-usage"))
.to("salesforce:getGlobalObjects")
.to(myProcessor);

// myProcessor will receive `Sforce-Limit-Info` header on the outbound
// message
class MyProcessor implements Processor {
    public void process(Exchange exchange) throws Exception {
        Message in = exchange.getIn();
        String apiLimits = in.getHeader("Sforce-Limit-Info", String.class);
    }
}

```

In addition, HTTP response status code and text are available as headers **Exchange.HTTP_RESPONSE_CODE** and **Exchange.HTTP_RESPONSE_TEXT**.

43.7. SUPPORTED SALESFORCE APIS

The component supports the following Salesforce APIs

Producer endpoints can use the following APIs. Most of the APIs process one record at a time, the Query API can retrieve multiple Records.

43.7.1. Rest API

You can use the following for **operationName**:

- `getVersions` - Gets supported Salesforce REST API versions
- `getResources` - Gets available Salesforce REST Resource endpoints
- `getGlobalObjects` - Gets metadata for all available SObject types
- `getBasicInfo` - Gets basic metadata for a specific SObject type
- `getDescription` - Gets comprehensive metadata for a specific SObject type
- `getSObject` - Gets an SObject using its Salesforce Id
- `createSObject` - Creates an SObject
- `updateSObject` - Updates an SObject using Id
- `deleteSObject` - Deletes an SObject using Id
- `getSObjectWithId` - Gets an SObject using an external (user defined) id field
- `upsertSObject` - Updates or inserts an SObject using an external id
- `deleteSObjectWithId` - Deletes an SObject using an external id
- `query` - Runs a Salesforce SOQL query

- queryMore - Retrieves more results (in case of large number of results) using result link returned from the 'query' API
- search - Runs a Salesforce SOSL query
- limits - fetching organization API usage limits
- recent - fetching recent items
- approval - submit a record or records (batch) for approval process
- approvals - fetch a list of all approval processes
- composite - submit up to 25 possibly related REST requests and receive individual responses. It's also possible to use "raw" composite without limitation.
- composite-tree - create up to 200 records with parent-child relationships (up to 5 levels) in one go
- composite-batch - submit a composition of requests in batch
- compositeRetrieveSObjectCollections - Retrieve one or more records of the same object type.
- compositeCreateSObjectCollections - Add up to 200 records, returning a list of SaveSObjectResult objects.
- compositeUpdateSObjectCollections - Update up to 200 records, returning a list of SaveSObjectResult objects.
- compositeUpsertSObjectCollections - Create or update (upsert) up to 200 records based on an external ID field. Returns a list of UpsertSObjectResult objects.
- compositeDeleteSObjectCollections - Delete up to 200 records, returning a list of SaveSObjectResult objects.
- queryAll - Runs a SOQL query. It returns the results that are deleted because of a merge (merges up to three records into one of the records, deletes the others, and reparents any related records) or delete. Also returns the information about archived Task and Event records.
- getBlobField - Retrieves the specified blob field from an individual record.
- apexCall - Executes a user defined APEX REST API call.
- raw - Send requests to salesforce and have full, raw control over endpoint, parameters, body, etc.

For example, the following producer endpoint uses the upsertSObject API, with the sObjectIdName parameter specifying 'Name' as the external id field. The request message body should be an SObject DTO generated using the maven plugin. The response message will either be **null** if an existing record was updated, or **CreateSObjectResult** with an id of the new record, or a list of errors while creating the new object.

```
...to("salesforce:upsertSObject?sObjectIdName=Name")...
```

43.7.2. Bulk 2.0 API

The Bulk 2.0 API has a simplified model over the original Bulk API. Use it to quickly load a large amount of

data into salesforce, or query a large amount of data out of salesforce. Data must be provided in CSV format. The minimum API version for Bulk 2.0 is v41.0. The minimum API version for Bulk Queries is v47.0. DTO classes mentioned below are from the **org.apache.camel.component.salesforce.api.dto.bulkv2** package. The following operations are supported:

- **bulk2CreateJob** - Create a bulk job. Supply an instance of **Job** in the message body.
- **bulk2GetJob** - Get an existing Job. **jobId** parameter is required.
- **bulk2CreateBatch** - Add a Batch of CSV records to a job. Supply CSV data in the message body. The first row must contain headers. **jobId** parameter is required.
- **bulk2CloseJob** - Close a job. You must close the job in order for it to be processed or aborted/deleted. **jobId** parameter is required.
- **bulk2AbortJob** - Abort a job. **jobId** parameter is required.
- **bulk2DeleteJob** - Delete a job. **jobId** parameter is required.
- **bulk2GetSuccessfulResults** - Get successful results for a job. Returned message body will contain an InputStream of CSV data. **jobId** parameter is required.
- **bulk2GetFailedResults** - Get failed results for a job. Returned message body will contain an InputStream of CSV data. **jobId** parameter is required.
- **bulk2GetUnprocessedRecords** - Get unprocessed records for a job. Returned message body will contain an InputStream of CSV data. **jobId** parameter is required.
- **bulk2GetAllJobs** - Get all jobs. Response body is an instance of **Jobs**. If the **done** property is false, there are additional pages to fetch, and the **nextRecordsUrl** property contains the value to be set in the **queryLocator** parameter on subsequent calls.
- **bulk2CreateQueryJob** - Create a bulk query job. Supply an instance of **QueryJob** in the message body.
- **bulk2GetQueryJob** - Get a bulk query job. **jobId** parameter is required.
- **bulk2GetQueryJobResults** - Get bulk query job results. **jobId** parameter is required. Accepts **maxRecords** and **locator** parameters. Response message headers include **Sforce-NumberOfRecords** and **Sforce-Locator** headers. The value of **Sforce-Locator** can be passed into subsequent calls via the **locator** parameter.
- **bulk2AbortQueryJob** - Abort a bulk query job. **jobId** parameter is required.
- **bulk2DeleteQueryJob** - Delete a bulk query job. **jobId** parameter is required.
- **bulk2GetAllQueryJobs** - Get all jobs. Response body is an instance of **QueryJobs**. If the **done** property is false, there are additional pages to fetch, and the **nextRecordsUrl** property contains the value to be set in the **queryLocator** parameter on subsequent calls.

43.7.3. Rest Bulk (original) API

Producer endpoints can use the following APIs. All Job data formats, i.e. xml, csv, zip/xml, and zip/csv are supported.

The request and response have to be marshalled/unmarshalled by the route. Usually the request will be

some stream source like a CSV file,
and the response may also be saved to a file to be correlated with the request.

You can use the following for **operationName**:

- **createJob** - Creates a Salesforce Bulk Job. Must supply a **JobInfo** instance in body. PK Chunking is supported via the `pkChunking*` options. See an explanation [here](#).
- **getJob** - Gets a Job using its Salesforce Id
- **closeJob** - Closes a Job
- **abortJob** - Aborts a Job
- **createBatch** - Submits a Batch within a Bulk Job
- **getBatch** - Gets a Batch using Id
- **getAllBatches** - Gets all Batches for a Bulk Job Id
- **getRequest** - Gets Request data (XML/CSV) for a Batch
- **getResults** - Gets the results of the Batch when its complete
- **createBatchQuery** - Creates a Batch from an SOQL query
- **getQueryResultIds** - Gets a list of Result Ids for a Batch Query
- **getQueryResult** - Gets results for a Result Id
- **getRecentReports** - Gets up to 200 of the reports you most recently viewed by sending a GET request to the Report List resource.
- **getReportDescription** - Retrieves the report, report type, and related metadata for a report, either in a tabular or summary or matrix format.
- **executeSyncReport** - Runs a report synchronously with or without changing filters and returns the latest summary data.
- **executeAsyncReport** - Runs an instance of a report asynchronously with or without filters and returns the summary data with or without details.
- **getReportInstances** - Returns a list of instances for a report that you requested to be run asynchronously. Each item in the list is treated as a separate instance of the report.
- **getReportResults**: Contains the results of running a report.

For example, the following producer endpoint uses the `createBatch` API to create a Job Batch. The in message must contain a body that can be converted into an **InputStream** (usually UTF-8 CSV or XML content from a file, etc.) and header fields `'jobId'` for the Job and `'contentType'` for the Job content type, which can be XML, CSV, ZIP_XML or ZIP_CSV. The put message body will contain **BatchInfo** on success, or throw a **SalesforceException** on error.

```
...to("salesforce:createBatch")..
```

43.7.4. Rest Streaming API

Consumer endpoints can use the following syntax for streaming endpoints to receive Salesforce notifications on create/update.

To create and subscribe to a topic

```
from("salesforce:CamelTestTopic?
notifyForFields=ALL&notifyForOperations=ALL&sObjectName=Merchandise__c&updateTopic=true&sO
bjectQuery=SELECT Id, Name FROM Merchandise__c")...
```

To subscribe to an existing topic

```
from("salesforce:CamelTestTopic&sObjectName=Merchandise__c")...
```

43.7.5. Platform events

To emit a platform event use **createSObject** operation. And set the message body can be JSON string or InputStream with key-value data – in that case **sObjectName** needs to be set to the API name of the event, or a class that extends from AbstractDTOBase with the appropriate class name for the event.

For example using a DTO:

```
class Order_Event__e extends AbstractDTOBase {
    @JsonProperty("OrderNumber")
    private String orderNumber;
    // ... other properties and getters/setters
}

from("timer:tick")
    .process(exchange -> {
        final Message in = exchange.getIn();
        String orderNumber = "ORD" + exchange.getProperty(Exchange.TIMER_COUNTER);
        Order_Event__e event = new Order_Event__e();
        event.setOrderNumber(orderNumber);
        in.setBody(event);
    })
    .to("salesforce:createSObject");
```

Or using JSON event data:

```
from("timer:tick")
    .process(exchange -> {
        final Message in = exchange.getIn();
        String orderNumber = "ORD" + exchange.getProperty(Exchange.TIMER_COUNTER);
        in.setBody("{\"OrderNumber\":\"" + orderNumber + "\"}");
    })
    .to("salesforce:createSObject?sObjectName=Order_Event__e");
```

To receive platform events use the consumer endpoint with the API name of the platform event prefixed with **event/** (or **/event/**), e.g.: **salesforce:events/Order_Event__e**. Processor consuming from that endpoint will receive either **org.apache.camel.component.salesforce.api.dto.PlatformEvent** object or **org.cometd.bayeux.Message** in the body depending on the **rawPayload** being **false** or **true** respectively.

For example, in the simplest form to consume one event:

```
PlatformEvent event = consumer.receiveBody("salesforce:event/Order_Event__e",
PlatformEvent.class);
```

43.7.6. Change data capture events

On the one hand, Salesforce could be configured to emit notifications for record changes of select objects. On the other hand, the Camel Salesforce component could react to such notifications, allowing for instance to [synchronize those changes into an external system](#) .

The notifications of interest could be specified in the **from("salesforce:XXX")** clause of a Camel route via the subscription channel, e.g:

```
from("salesforce:data/ChangeEvents?replayId=-1").log("being notified of all change events")
from("salesforce:data/AccountChangeEvent?replayId=-1").log("being notified of change events for
Account records")
from("salesforce:data/Employee__ChangeEvent?replayId=-1").log("being notified of change events
for Employee__c custom object")
```

The received message contains either **java.util.Map<String, Object>** or **org.cometd.bayeux.Message** in the body depending on the **rawPayload** being **false** or **true** respectively. The **CamelSalesforceChangeType** header could be valued to one of **CREATE, UPDATE, DELETE** or **UNDELETE**.

More details about how to use the Camel Salesforce component change data capture capabilities could be found in the [ChangeEventsConsumerIntegrationTest](#).

The [Salesforce developer guide](#) is a good fit to better know the subtleties of implementing a change data capture integration application. The dynamic nature of change event body fields, high level replication steps as well as security considerations could be of interest.

43.8. EXAMPLES

43.8.1. Uploading a document to a ContentWorkspace

Create the ContentVersion in Java, using a Processor instance:

```
public class ContentProcessor implements Processor {
    public void process(Exchange exchange) throws Exception {
        Message message = exchange.getIn();

        ContentVersion cv = new ContentVersion();
        ContentWorkspace cw = getWorkspace(exchange);
        cv.setFirstPublishLocationId(cw.getId());
        cv.setTitle("test document");
        cv.setPathOnClient("test_doc.html");
        byte[] document = message.getBody(byte[].class);
        ObjectMapper mapper = new ObjectMapper();
        String enc = mapper.convertValue(document, String.class);
        cv.setVersionDataUrl(enc);
        message.setBody(cv);
    }

    protected ContentWorkspace getWorkSpace(Exchange exchange) {
```

```

// Look up the content workspace somehow, maybe use enrich() to add it to a
// header that can be extracted here
----
}
}

```

Give the output from the processor to the Salesforce component:

```

from("file:///home/camel/library")
  .to(new ContentProcessor()) // convert bytes from the file into a ContentVersion SObject
  // for the salesforce component
  .to("salesforce:createSObject");

```

43.9. USING SALESFORCE LIMITS API

With **salesforce:limits** operation you can fetch of API limits from Salesforce and then act upon that data received. The result of **salesforce:limits** operation is mapped to **org.apache.camel.component.salesforce.api.dto.Limits** class and can be used in a custom processors or expressions.

For instance, consider that you need to limit the API usage of Salesforce so that 10% of daily API requests is left for other routes. The body of output message contains an instance of **org.apache.camel.component.salesforce.api.dto.Limits** object that can be used in conjunction with Content Based Router and Content Based Router and [Spring Expression Language \(SpEL\)](#) to choose when to perform queries.

Notice how multiplying **1.0** with the integer value held in **body.dailyApiRequests.remaining** makes the expression evaluate as with floating point arithmetic, without it – it would end up making integral division which would result with either **0** (some API limits consumed) or **1** (no API limits consumed).

```

from("direct:querySalesforce")
  .to("salesforce:limits")
  .choice()
  .when(spel("#{1.0 * body.dailyApiRequests.remaining / body.dailyApiRequests.max < 0.1}"))
  .to("salesforce:query?...")
  .otherwise()
  .setBody(constant("Used up Salesforce API limits, leaving 10% for critical routes"))
  .endChoice()

```

43.10. WORKING WITH APPROVALS

All the properties are named exactly the same as in the Salesforce REST API prefixed with **approval.** You can set approval properties by setting **approval.PropertyName** of the Endpoint these will be used as template – meaning that any property not present in either body or header will be taken from the Endpoint configuration. Or you can set the approval template on the Endpoint by assigning **approval** property to a reference onto a bean in the Registry.

You can also provide header values using the same **approval.PropertyName** in the incoming message headers.

And finally body can contain one **ApprovalRequest** or an **Iterable** of **ApprovalRequest** objects to process as a batch.

The important thing to remember is the priority of the values specified in these three mechanisms:

1. value in body takes precedence before any other
2. value in message header takes precedence before template value
3. value in template is set if no other value in header or body was given

For example to send one record for approval using values in headers use:

Given a route:

```
from("direct:example1")//
    .setHeader("approval.ContextId", simple("${body['contextId']}"))
    .setHeader("approval.NextApproverIds", simple("${body['nextApproverIds']}"))
    .to("salesforce:approval?"//
        + "approval.actionType=Submit"//
        + "&approval.comments=this is a test"//
        + "&approval.processDefinitionNameOrId=Test_Account_Process"//
        + "&approval.skipEntryCriteria=true");
```

You could send a record for approval using:

```
final Map<String, String> body = new HashMap<>();
body.put("contextId", accountIds.iterator().next());
body.put("nextApproverIds", userId);

final ApprovalResult result = template.requestBody("direct:example1", body, ApprovalResult.class);
```

43.11. USING SALESFORCE RECENT ITEMS API

To fetch the recent items use **salesforce:recent** operation. This operation returns an **java.util.List** of **org.apache.camel.component.salesforce.api.dto.RecentItem** objects (**List<RecentItem>**) that in turn contain the **Id**, **Name** and **Attributes** (with **type** and **url** properties). You can limit the number of returned items by specifying **limit** parameter set to maximum number of records to return. For example:

```
from("direct:fetchRecentItems")
    to("salesforce:recent")
    .split().body()
    .log("${body.name} at ${body.attributes.url}");
```

43.12. USING SALESFORCE COMPOSITE API TO SUBMIT SUBJECT TREE

To create up to 200 records including parent-child relationships use **salesforce:composite-tree** operation. This requires an instance of **org.apache.camel.component.salesforce.api.dto.composite.SObjectTree** in the input message and returns the same tree of objects in the output message. The **org.apache.camel.component.salesforce.api.dto.AbstractSObjectBase** instances within the tree get updated with the identifier values (**Id** property) or their corresponding **org.apache.camel.component.salesforce.api.dto.composite.SObjectNode** is populated with **errors** on failure.

Note that for some records operation can succeed and for some it can fail – so you need to manually check for errors.

Easiest way to use this functionality is to use the DTOs generated by the **camel-salesforce-maven-plugin**, but you also have the option of customizing the references that identify the each object in the tree, for instance primary keys from your database.

Lets look at an example:

```
Account account = ...
Contact president = ...
Contact marketing = ...

Account anotherAccount = ...
Contact sales = ...
Asset someAsset = ...

// build the tree
SObjectTree request = new SObjectTree();
request.addObject(account).addChildren(president, marketing);
request.addObject(anotherAccount).addChild(sales).addChild(someAsset);

final SObjectTree response = template.requestBody("salesforce:composite-tree", tree,
SObjectTree.class);
final Map<Boolean, List<SObjectNode>> result = response.allNodes()
.collect(Collectors.groupingBy(SObjectNode::hasErrors));

final List<SObjectNode> withErrors = result.get(true);
final List<SObjectNode> succeeded = result.get(false);

final String firstId = succeeded.get(0).getId();
```

43.13. USING SALESFORCE COMPOSITE API TO SUBMIT MULTIPLE REQUESTS IN A BATCH

The Composite API batch operation (**composite-batch**) allows you to accumulate multiple requests in a batch and then submit them in one go, saving the round trip cost of multiple individual requests. Each response is then received in a list of responses with the order preserved, so that the n-th requests response is in the n-th place of the response.



NOTE

The results can vary from API to API so the result of the request is given as a **java.lang.Object**. In most cases the result will be a **java.util.Map** with string keys and values or other **java.util.Map** as value. Requests are made in JSON format and hold some type information (i.e. it is known what values are strings and what values are numbers).

Lets look at an example:

```
final String accountId = ...
final SObjectBatch batch = new SObjectBatch("38.0");

final Account updates = new Account();
updates.setName("NewName");
batch.addUpdate("Account", accountId, updates);

final Account newAccount = new Account();
```

```

newAccount.setName("Account created from Composite batch API");
batch.addCreate(newAccount);

batch.addGet("Account", accountId, "Name", "BillingPostalCode");

batch.addDelete("Account", accountId);

final SObjectBatchResponse response = template.requestBody("salesforce:composite-batch", batch,
SObjectBatchResponse.class);

boolean hasErrors = response.hasErrors(); // if any of the requests has resulted in either 4xx or 5xx
HTTP status
final List<SObjectBatchResult> results = response.getResults(); // results of three operations sent in
batch

final SObjectBatchResult updateResult = results.get(0); // update result
final int updateStatus = updateResult.getStatusCode(); // probably 204
final Object updateResultData = updateResult.getResult(); // probably null

final SObjectBatchResult createResult = results.get(1); // create result
@SuppressWarnings("unchecked")
final Map<String, Object> createData = (Map<String, Object>) createResult.getResult();
final String newAccountId = createData.get("id"); // id of the new account, this is for JSON, for XML it
would be createData.get("Result").get("id")


final SObjectBatchResult retrieveResult = results.get(2); // retrieve result
@SuppressWarnings("unchecked")
final Map<String, Object> retrieveData = (Map<String, Object>) retrieveResult.getResult();
final String accountName = retrieveData.get("Name"); // Name of the retrieved account, this is for
JSON, for XML it would be createData.get("Account").get("Name")
final String accountBillingPostalCode = retrieveData.get("BillingPostalCode"); // Name of the retrieved
account, this is for JSON, for XML it would be createData.get("Account").get("BillingPostalCode")

final SObjectBatchResult deleteResult = results.get(3); // delete result
final int updateStatus = deleteResult.getStatusCode(); // probably 204
final Object updateResultData = deleteResult.getResult(); // probably null

```

43.14. USING SALESFORCE COMPOSITE API TO SUBMIT MULTIPLE CHAINED REQUESTS

The **composite** operation allows submitting up to 25 requests that can be chained together, for instance identifier generated in previous request can be used in subsequent request. Individual requests and responses are linked with the provided *reference*.



NOTE

Composite API supports only JSON payloads.



NOTE

As with the batch API the results can vary from API to API so the result of the request is given as a **java.lang.Object**. In most cases the result will be a **java.util.Map** with string keys and values or other **java.util.Map** as value. Requests are made in JSON format hold some type information (i.e. it is known what values are strings and what values are numbers).

Lets look at an example:

```
SObjectComposite composite = new SObjectComposite("38.0", true);

// first insert operation via an external id
final Account updateAccount = new TestAccount();
updateAccount.setName("Salesforce");
updateAccount.setBillingStreet("Landmark @ 1 Market Street");
updateAccount.setBillingCity("San Francisco");
updateAccount.setBillingState("California");
updateAccount.setIndustry(Account_IndustryEnum.TECHNOLOGY);
composite.addUpdate("Account", "001xx000003DlpcAAG", updateAccount, "UpdatedAccount");

final Contact newContact = new TestContact();
newContact.setLastName("John Doe");
newContact.setPhone("1234567890");
composite.addCreate(newContact, "NewContact");

final AccountContactJunction__c junction = new AccountContactJunction__c();
junction.setAccount__c("001xx000003DlpcAAG");
junction.setContactId__c("@{NewContact.id}");
composite.addCreate(junction, "JunctionRecord");

final SObjectCompositeResponse response = template.requestBody("salesforce:composite",
composite, SObjectCompositeResponse.class);
final List<SObjectCompositeResult> results = response.getCompositeResponse();

final SObjectCompositeResult accountUpdateResult = results.stream().filter(r ->
"UpdatedAccount".equals(r.getReferenceId())).findFirst().get()
final int statusCode = accountUpdateResult.getHttpStatusCode(); // should be 200
final Map<String, ?> accountUpdateBody = accountUpdateResult.getBody();

final SObjectCompositeResult contactCreationResult = results.stream().filter(r ->
"JunctionRecord".equals(r.getReferenceId())).findFirst().get()
```

43.15. USING "RAW" SALESFORCE COMPOSITE

It's possible to directly call Salesforce composite by preparing the Salesforce JSON request in the route thanks to the **rawPayload** option.

For instance, you can have the following route:

```
from("timer:fire?period=2000").setBody(constant("{\n" +
  "\nallOrNone\" : true,\n" +
  "\nrecords\" : [ { \n" +
  "\nattributes\" : {\"type\" : \"FOO\"},\n" +
```

```

"  \"Name\" : \"123456789\", \"n\" +
"  \"FOO\" : \"XXX\", \"n\" +
"  \"ACCOUNT\" : 2100.0 \"n\" +
"  \"ExternalID\" : \"EXTERNAL\" \"n\"
" }} \"n\" +
" }")
.to("salesforce:composite?rawPayload=true")
.log("${body}");

```

The route directly creates the body as JSON and directly submit to salesforce endpoint using **rawPayload=true** option.

With this approach, you have the complete control on the Salesforce request.

POST is the default HTTP method used to send raw Composite requests to salesforce. Use the **compositeMethod** option to override to the other supported value, **GET**, which returns a list of other available composite resources.

43.16. USING RAW OPERATION

Send HTTP requests to salesforce with full, raw control of all aspects of the call. Any serialization or deserialization of request and response bodies must be performed in the route. The **Content-Type** HTTP header will be automatically set based on the **format** option, but this can be overridden with the **rawHttpHeaders** option.

Parameter	Type	Description	Default	Required
request body	String or InputStream	Body of the HTTP request		
rawPath	String	The portion of the endpoint URL after the domain name, e.g., '/services/data/v5.1.0/subjects/Account/'		x
rawMethod	String	The HTTP method		x
rawQueryParameters	String	Comma separated list of message headers to include as query parameters. Do not url-encode values as this will be done automatically.		
rawHttpHeaders	String	Comma separated list of message headers to include as HTTP headers		

43.16.1. Query example

In this example we'll send a query to the REST API. The query must be passed in a URL parameter called "q", so we'll create a message header called q and tell the raw operation to include that message header as a URL parameter:

```
from("direct:queryExample")
  .setHeader("q", "SELECT Id, LastName FROM Contact")
  .to("salesforce:raw?
format=JSON&rawMethod=GET&rawQueryParameters=q&rawPath=/services/data/v51.0/query")
  // deserialize JSON results or handle in some other way
```

43.16.2. SObject example

In this example, we'll pass a Contact the REST API in a **create** operation. Since the **raw** operation does not perform any serialization, we make sure to pass XML in the message body

```
from("direct:createAContact")
  .setBody(constant("<Contact><LastName>TestLast</LastName></Contact>"))
  .to("salesforce:raw?
format=XML&rawMethod=POST&rawPath=/services/data/v51.0/objects/Contact")
```

The response is:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<Result>
  <id>0034x00000RnV6zAAF</id>
  <success>true</success>
</Result>
```

43.17. USING COMPOSITE SUBJECT COLLECTIONS

The SObject Collections API executes actions on multiple records in one request. Use sObject Collections to reduce the number of round-trips between the client and server. The entire request counts as a single call toward your API limits. This resource is available in API version 42.0 and later.

SObject records (aka DTOs) supplied to these operations must be instances of subclasses of **AbstractDescribedObjectBase**. See the Maven Plugin section for information on generating these DTO classes. These operations serialize supplied DTOs to JSON.

43.17.1. compositeRetrieveSObjectCollections

Retrieve one or more records of the same object type.

Parameter	Type	Description	Default	Required
ids	List of String or comma-separated string	A list of one or more IDs of the objects to return. All IDs must belong to the same object type.		x

Parameter	Type	Description	Default	Required
fields	List of String or comma-separated string	A list of fields to include in the response. The field names you specify must be valid, and you must have read-level permissions to each field.		x
sObjectName	String	Type of SObject, e.g. Account		x
sObjectClass	String	Fully-qualified class name of DTO class to use for deserializing the response.		Required if sObjectName parameter does not resolve to a class that exists in the package specified by the package option.

43.17.2. compositeCreateObjectCollections

Add up to 200 records, returning a list of SaveSObjectResult objects. Mixed SObject types is supported.

Parameter	Type	Description	Default	Required
request body	List of SObject	A list of SObjects to create		x
allOrNone	boolean	Indicates whether to roll back the entire request when the creation of any object fails (true) or to continue with the independent creation of other objects in the request.	false	

43.17.3. compositeUpdateObjectCollections

Update up to 200 records, returning a list of SaveSObjectResult objects. Mixed SObject types is supported.

Parameter	Type	Description	Default	Required
request body	List of SObject	A list of SObjects to update		x
allOrNone	boolean	Indicates whether to roll back the entire request when the update of any object fails (true) or to continue with the independent update of other objects in the request.	false	

43.17.4. compositeUpsertObjectCollections

Create or update (upsert) up to 200 records based on an external ID field, returning a list of UpsertSObjectResult objects. Mixed SObject types is not supported.

Parameter	Type	Description	Default	Required
request body	List of SObject	A list of SObjects to upsert		x
allOrNone	boolean	Indicates whether to roll back the entire request when the upsert of any object fails (true) or to continue with the independent upsert of other objects in the request.	false	
sObjectName	String	Type of SObject, e.g. Account		x
sObjectIdName	String	Name of External ID field		x

43.17.5. compositeDeleteObjectCollections

Delete up to 200 records, returning a list of DeleteSObjectResult objects. Mixed SObject types is supported.

Parameter	Type	Description	Default	Required
sObjectIds or request body	List of String or comma-separated string	A list of up to 200 IDs of objects to be deleted.		x

Parameter	Type	Description	Default	Required
allOrNone	boolean	Indicates whether to roll back the entire request when the deletion of any object fails (true) or to continue with the independent deletion of other objects in the request.	false	

43.18. SENDING NULL VALUES TO SALESFORCE

By default, SObject fields with null values are not sent to salesforce. In order to send null values to salesforce, use the **fieldsToNull** property, as follows:

```
accountSObject.getFieldsToNull().add("Site");
```

43.19. GENERATING SOQL QUERY STRINGS

org.apache.camel.component.salesforce.api.utils.QueryHelper contains helper methods to generate SOQL queries. For instance to fetch all custom fields from *Account* SObject you can simply generate the SOQL SELECT by invoking:

```
String allCustomFieldsQuery = QueryHelper.queryToFetchFilteredFieldsOf(new Account(),
    SObjectField::isCustom);
```

43.20. CAMEL SALESFORCE MAVEN PLUGIN

This Maven plugin generates DTOs for the Camel.

For obvious security reasons it is recommended that the `clientId`, `clientSecret`, `userName` and `password` fields be not set in the `pom.xml`. The plugin should be configured for the rest of the properties, and can be executed using the following command:

```
mvn camel-salesforce:generate -DcamelSalesforce.clientId=<clientId> -
DcamelSalesforce.clientSecret=<clientsecret> \
-DcamelSalesforce.userName=<username> -DcamelSalesforce.password=<password>
```

The generated DTOs use Jackson annotations. All Salesforce field types are supported. Date and time fields are mapped to **java.time.ZonedDateTime** by default, and picklist fields are mapped to generated Java Enumerations.

Please refer to [README.md](#) for details on how to generate the DTO.

43.21. SPRING BOOT AUTO-CONFIGURATION

When using salesforce with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
```

```

<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-salesforce-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 91 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.salesforce.all-or-none</code>	Composite API option to indicate to rollback all records if any are not successful.	false	Boolean
<code>camel.component.salesforce.apex-method</code>	APEX method name.		String
<code>camel.component.salesforce.apex-query-params</code>	Query params for APEX method.		Map
<code>camel.component.salesforce.apex-url</code>	APEX method URL.		String
<code>camel.component.salesforce.api-version</code>	Salesforce API version.	53.0	String
<code>camel.component.salesforce.authentication-type</code>	Explicit authentication method to be used, one of USERNAME_PASSWORD, REFRESH_TOKEN or JWT. Salesforce component can auto-determine the authentication method to use from the properties set, set this property to eliminate any ambiguity.		AuthenticationType
<code>camel.component.salesforce.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.salesforce.backoff-increment</code>	Backoff interval increment for Streaming connection restart attempts for failures beyond CometD auto-reconnect. The option is a long type.	1000	Long

Name	Description	Default	Type
<code>camel.component.salesforce.batch-id</code>	Bulk API Batch ID.		String
<code>camel.component.salesforce.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.salesforce.client-id</code>	OAuth Consumer Key of the connected app configured in the Salesforce instance setup. Typically a connected app needs to be configured but one can be provided by installing a package.		String
<code>camel.component.salesforce.client-secret</code>	OAuth Consumer Secret of the connected app configured in the Salesforce instance setup.		String
<code>camel.component.salesforce.composite-method</code>	Composite (raw) method.		String
<code>camel.component.salesforce.config</code>	Global endpoint configuration - use to set values that are common to all endpoints. The option is a <code>org.apache.camel.component.salesforce.SalesforceEndpointConfig</code> type.		SalesforceEndpointConfig
<code>camel.component.salesforce.content-type</code>	Bulk API content type, one of XML, CSV, ZIP_XML, ZIP_CSV.		ContentType
<code>camel.component.salesforce.default-replay-id</code>	Default replayId setting if no value is found in <code>initialReplayIdMap</code> .	-1	Long
<code>camel.component.salesforce.enabled</code>	Whether to enable auto configuration of the salesforce component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.salesforce.fallback-replay-id</code>	ReplayId to fall back to after an Invalid Replay Id response.	-1	Long
<code>camel.component.salesforce.format</code>	Payload format to use for Salesforce API calls, either JSON or XML, defaults to JSON. As of Camel 3.12, this option only applies to the Raw operation.		PayloadFormat
<code>camel.component.salesforce.http-client</code>	Custom Jetty Http Client to use to connect to Salesforce. The option is a <code>org.apache.camel.component.salesforce.SalesforceHttpClient</code> type.		SalesforceHttpClient
<code>camel.component.salesforce.http-client-connection-timeout</code>	Connection timeout used by the HttpClient when connecting to the Salesforce server.	60000	Long
<code>camel.component.salesforce.http-client-idle-timeout</code>	Timeout used by the HttpClient when waiting for response from the Salesforce server.	10000	Long
<code>camel.component.salesforce.http-client-properties</code>	Used to set any properties that can be configured on the underlying HTTP client. Have a look at properties of <code>SalesforceHttpClient</code> and the <code>Jetty HttpClient</code> for all available options.		Map
<code>camel.component.salesforce.http-max-content-length</code>	Max content length of an HTTP response.		Integer
<code>camel.component.salesforce.http-proxy-auth-uri</code>	Used in authentication against the HTTP proxy server, needs to match the URI of the proxy server in order for the <code>httpProxyUsername</code> and <code>httpProxyPassword</code> to be used for authentication.		String
<code>camel.component.salesforce.http-proxy-excluded-addresses</code>	A list of addresses for which HTTP proxy server should not be used.		Set
<code>camel.component.salesforce.http-proxy-host</code>	Hostname of the HTTP proxy server to use.		String

Name	Description	Default	Type
<code>camel.component.salesforce.http-proxy-included-addresses</code>	A list of addresses for which HTTP proxy server should be used.		Set
<code>camel.component.salesforce.http-proxy-password</code>	Password to use to authenticate against the HTTP proxy server.		String
<code>camel.component.salesforce.http-proxy-port</code>	Port number of the HTTP proxy server to use.		Integer
<code>camel.component.salesforce.http-proxy-realm</code>	Realm of the proxy server, used in preemptive Basic/Digest authentication methods against the HTTP proxy server.		String
<code>camel.component.salesforce.http-proxy-secure</code>	If set to false disables the use of TLS when accessing the HTTP proxy.	true	Boolean
<code>camel.component.salesforce.http-proxy-socks4</code>	If set to true the configures the HTTP proxy to use as a SOCKS4 proxy.	false	Boolean
<code>camel.component.salesforce.http-proxy-use-digest-auth</code>	If set to true Digest authentication will be used when authenticating to the HTTP proxy, otherwise Basic authorization method will be used.	false	Boolean
<code>camel.component.salesforce.http-proxy-username</code>	Username to use to authenticate against the HTTP proxy server.		String
<code>camel.component.salesforce.http-request-buffer-size</code>	HTTP request buffer size. May need to be increased for large SOQL queries.	8192	Integer
<code>camel.component.salesforce.include-details</code>	Include details in Salesforce Analytics report, defaults to false.		Boolean

Name	Description	Default	Type
<code>camel.component.salesforce.initial-replay-id-map</code>	Replay IDs to start from per channel name.		Map
<code>camel.component.salesforce.instance-id</code>	Salesforce1 Analytics report execution instance ID.		String
<code>camel.component.salesforce.instance-url</code>	URL of the Salesforce instance used after authentication, by default received from Salesforce on successful authentication.		String
<code>camel.component.salesforce.job-id</code>	Bulk API Job ID.		String
<code>camel.component.salesforce.jwt-audience</code>	Value to use for the Audience claim (aud) when using OAuth JWT flow. If not set, the login URL will be used, which is appropriate in most cases.		String
<code>camel.component.salesforce.keystore</code>	KeyStore parameters to use in OAuth JWT flow. The KeyStore should contain only one entry with private key and certificate. Salesforce does not verify the certificate chain, so this can easily be a selfsigned certificate. Make sure that you upload the certificate to the corresponding connected app. The option is a <code>org.apache.camel.support.jsse.KeyStoreParameters</code> type.		KeyStoreParameters
<code>camel.component.salesforce.lazy-login</code>	If set to true prevents the component from authenticating to Salesforce with the start of the component. You would generally set this to the (default) false and authenticate early and be immediately aware of any authentication issues.	false	Boolean
<code>camel.component.salesforce.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

Name	Description	Default	Type
<code>camel.component.salesforce.limit</code>	Limit on number of returned records. Applicable to some of the API, check the Salesforce documentation.		Integer
<code>camel.component.salesforce.locator</code>	Locator provided by salesforce Bulk 2.0 API for use in getting results for a Query job.		String
<code>camel.component.salesforce.login-config</code>	All authentication configuration in one nested bean, all properties set there can be set directly on the component as well. The option is a <code>org.apache.camel.component.salesforce.SalesforceLoginConfig</code> type.		<code>SalesforceLoginConfig</code>
<code>camel.component.salesforce.login-url</code>	URL of the Salesforce instance used for authentication, by default set to .		String
<code>camel.component.salesforce.long-polling-transport-properties</code>	Used to set any properties that can be configured on the <code>LongPollingTransport</code> used by the <code>BayeuxClient</code> (<code>CometD</code>) used by the streaming api.		Map
<code>camel.component.salesforce.max-backoff</code>	Maximum backoff interval for Streaming connection restart attempts for failures beyond <code>CometD</code> auto-reconnect. The option is a long type.	30000	Long
<code>camel.component.salesforce.max-records</code>	The maximum number of records to retrieve per set of results for a Bulk 2.0 Query. The request is still subject to the size limits. If you are working with a very large number of query results, you may experience a timeout before receiving all the data from Salesforce. To prevent a timeout, specify the maximum number of records your client is expecting to receive in the <code>maxRecords</code> parameter. This splits the results into smaller sets with this value as the maximum size.		Integer
<code>camel.component.salesforce.not-found-behaviour</code>	Sets the behaviour of 404 not found status received from Salesforce API. Should the body be set to <code>NULL</code> <code>NotFoundBehaviour#NULL</code> or should a exception be signaled on the exchange <code>NotFoundBehaviour#EXCEPTION</code> - the default.		<code>NotFoundBehaviour</code>
<code>camel.component.salesforce.notify-for-fields</code>	Notify for fields, options are ALL, REFERENCED, SELECT, WHERE.		<code>NotifyForFieldsEnum</code>

Name	Description	Default	Type
<code>camel.component.salesforce.notify-for-operation-create</code>	Notify for create operation, defaults to false (API version = 29.0).		Boolean
<code>camel.component.salesforce.notify-for-operation-delete</code>	Notify for delete operation, defaults to false (API version = 29.0).		Boolean
<code>camel.component.salesforce.notify-for-operation-undelete</code>	Notify for un-delete operation, defaults to false (API version = 29.0).		Boolean
<code>camel.component.salesforce.notify-for-operation-update</code>	Notify for update operation, defaults to false (API version = 29.0).		Boolean
<code>camel.component.salesforce.notify-for-operations</code>	Notify for operations, options are ALL, CREATE, EXTENDED, UPDATE (API version 29.0).		NotifyForOperationsEnum
<code>camel.component.salesforce.object-mapper</code>	Custom Jackson ObjectMapper to use when serializing/deserializing Salesforce objects. The option is a <code>com.fasterxml.jackson.databind.ObjectMapper</code> type.		ObjectMapper
<code>camel.component.salesforce.packages</code>	In what packages are the generated DTO classes. Typically the classes would be generated using <code>camel-salesforce-maven-plugin</code> . Set it if using the generated DTOs to gain the benefit of using short SObject names in parameters/header values. Multiple packages can be separated by comma.		String
<code>camel.component.salesforce.password</code>	Password used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows. Make sure that you append security token to the end of the password if using one.		String
<code>camel.component.salesforce.pk-chunking</code>	Use PK Chunking. Only for use in original Bulk API. Bulk 2.0 API performs PK chunking automatically, if necessary.		Boolean

Name	Description	Default	Type
<code>camel.component.salesforce.pk-chunking-chunk-size</code>	Chunk size for use with PK Chunking. If unspecified, salesforce default is 100,000. Maximum size is 250,000.		Integer
<code>camel.component.salesforce.pk-chunking-parent</code>	Specifies the parent object when you're enabling PK chunking for queries on sharing objects. The chunks are based on the parent object's records rather than the sharing object's records. For example, when querying on AccountShare, specify Account as the parent object. PK chunking is supported for sharing objects as long as the parent object is supported.		String
<code>camel.component.salesforce.pk-chunking-start-row</code>	Specifies the 15-character or 18-character record ID to be used as the lower boundary for the first chunk. Use this parameter to specify a starting ID when restarting a job that failed between batches.		String
<code>camel.component.salesforce.query-locator</code>	Query Locator provided by salesforce for use when a query results in more records than can be retrieved in a single call. Use this value in a subsequent call to retrieve additional records.		String
<code>camel.component.salesforce.raw-http-headers</code>	Comma separated list of message headers to include as HTTP parameters for Raw operation.		String
<code>camel.component.salesforce.raw-method</code>	HTTP method to use for the Raw operation.		String
<code>camel.component.salesforce.raw-path</code>	The portion of the endpoint URL after the domain name. E.g., <code>'/services/data/v52.0/subjects/Account/'</code> .		String
<code>camel.component.salesforce.raw-payload</code>	Use raw payload String for request and response (either JSON or XML depending on format), instead of DTOs, false by default.	false	Boolean
<code>camel.component.salesforce.raw-query-parameters</code>	Comma separated list of message headers to include as query parameters for Raw operation. Do not url-encode values as this will be done automatically.		String

Name	Description	Default	Type
<code>camel.component.salesforce.refresh-token</code>	Refresh token already obtained in the refresh token OAuth flow. One needs to setup a web application and configure a callback URL to receive the refresh token, or configure using the builtin callback at and then retrieve the refresh_token from the URL at the end of the flow. Note that in development organizations Salesforce allows hosting the callback web application at localhost.		String
<code>camel.component.salesforce.report-id</code>	Salesforce1 Analytics report Id.		String
<code>camel.component.salesforce.report-metadata</code>	Salesforce1 Analytics report metadata for filtering. The option is a <code>org.apache.camel.component.salesforce.api.dto.analytics.reports.ReportMetadata</code> type.		ReportMetadata
<code>camel.component.salesforce.result-id</code>	Bulk API Result ID.		String
<code>camel.component.salesforce.s-object-blob-field-name</code>	SObject blob field name.		String
<code>camel.component.salesforce.s-object-class</code>	Fully qualified SObject class name, usually generated using <code>camel-salesforce-maven-plugin</code> .		String
<code>camel.component.salesforce.s-object-fields</code>	SObject fields to retrieve.		String
<code>camel.component.salesforce.s-object-id</code>	SObject ID if required by API.		String
<code>camel.component.salesforce.s-object-id-name</code>	SObject external ID field name.		String
<code>camel.component.salesforce.s-object-id-value</code>	SObject external ID field value.		String

Name	Description	Default	Type
<code>camel.component.salesforce.s-object-name</code>	SObject name if required or supported by API.		String
<code>camel.component.salesforce.s-object-query</code>	Salesforce SOQL query string.		String
<code>camel.component.salesforce.s-object-search</code>	Salesforce SOSL search string.		String
<code>camel.component.salesforce.ssl-context-parameters</code>	SSL parameters to use, see <code>SSLContextParameters</code> class for all available options. The option is a <code>org.apache.camel.support.jsse.SSLContextParameters</code> type.		<code>SSLContextParameters</code>
<code>camel.component.salesforce.update-topic</code>	Whether to update an existing Push Topic when using the Streaming API, defaults to false.	false	Boolean
<code>camel.component.salesforce.use-global-ssl-context-parameters</code>	Enable usage of global SSL context parameters.	false	Boolean
<code>camel.component.salesforce.username</code>	Username used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows.		String
<code>camel.component.salesforce.worker-pool-max-size</code>	Maximum size of the thread pool used to handle HTTP responses.	20	Integer
<code>camel.component.salesforce.worker-pool-size</code>	Size of the thread pool used to handle HTTP responses.	10	Integer

CHAPTER 44. SCHEDULER

Only consumer is supported


The Scheduler component is used to generate message exchanges when a scheduler fires. This component is similar to the [Timer](#) component, but it offers more functionality in terms of scheduling. Also this component uses JDK **ScheduledExecutorService**. Where as the timer uses a JDK **Timer**.

You can only consume events from this endpoint.

44.1. URI FORMAT

```
scheduler:name[?options]
```

Where **name** is the name of the scheduler, which is created and shared across endpoints. So if you use the same name for all your scheduler endpoints, only one scheduler thread pool and thread will be used - but you can configure the thread pool to allow more concurrent threads.



NOTE

The IN body of the generated exchange is **null**. So **exchange.getIn().getBody()** returns **null**.

44.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

44.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

44.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

44.3. COMPONENT OPTIONS

The Scheduler component supports 3 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
poolSize (scheduler)	Number of core threads in the thread pool used by the scheduling thread pool. Is by default using a single thread.	1	int

44.4. ENDPOINT OPTIONS

The Scheduler endpoint is configured using URI syntax:

```
scheduler:name
```

with the following path and query parameters:

44.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (consumer)	Required The name of the scheduler.		String

44.4.2. Query Parameters (21 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • <code>InOnly</code> • <code>InOut</code> • <code>InOptionalOut</code> 		<code>ExchangePattern</code>
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.		int

Name	Description	Default	Type
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If <code>greedy</code> is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
poolSize (scheduler)	Number of core threads in the thread pool used by the scheduling thread pool. Is by default using a single thread.	1	int
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService

Name	Description	Default	Type
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> ● NANoseconds ● MICROseconds ● MILLIseconds ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean

44.5. MORE INFORMATION

This component is a scheduler [Polling Consumer](#) where you can find more information about the options above, and examples at the [Polling Consumer](#) page.

44.6. EXCHANGE PROPERTIES

When the timer is fired, it adds the following information as properties to the **Exchange**:

Name	Type	Description
Exchange.TIMER_NAME	String	The value of the name option.
Exchange.TIMER_FIRED_TIME	Date	The time when the consumer fired.

44.7. SAMPLE

To set up a route that generates an event every 60 seconds:

```
from("scheduler://foo?delay=60000").to("bean:myBean?method=someMethodName");
```

The above route will generate an event and then invoke the **someMethodName** method on the bean called **myBean** in the Registry such as JNDI or Spring.

And the route in Spring DSL:

```
<route>
  <from uri="scheduler://foo?delay=60000"/>
  <to uri="bean:myBean?method=someMethodName"/>
</route>
```

44.8. FORCING THE SCHEDULER TO TRIGGER IMMEDIATELY WHEN COMPLETED

To let the scheduler trigger as soon as the previous task is complete, you can set the option **greedy=true**. But beware then the scheduler will keep firing all the time. So use this with caution.

44.9. FORCING THE SCHEDULER TO BE IDLE

There can be use cases where you want the scheduler to trigger and be greedy. But sometimes you want "tell the scheduler" that there was no task to poll, so the scheduler can change into idle mode using the backoff options. To do this you would need to set a property on the exchange with the key **Exchange.SCHEDULER_POLLED_MESSAGES** to a boolean value of false. This will cause the consumer to indicate that there was no messages polled.

The consumer will otherwise as by default return 1 message polled to the scheduler, every time the consumer has completed processing the exchange.

44.10. SPRING BOOT AUTO-CONFIGURATION

When using scheduler with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

-

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-scheduler-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.scheduler.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.scheduler.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.scheduler.enabled</code>	Whether to enable auto configuration of the scheduler component. This is enabled by default.		Boolean
<code>camel.component.scheduler.pool-size</code>	Number of core threads in the thread pool used by the scheduling thread pool. Is by default using a single thread.	1	Integer

CHAPTER 45. SEDA

Both producer and consumer are supported

The SEDA component provides asynchronous [SEDA](#) behavior, so that messages are exchanged on a [BlockingQueue](#) and consumers are invoked in a separate thread from the producer.

Note that queues are only visible within a *single* CamelContext. If you want to communicate across **CamelContext** instances (for example, communicating between Web applications), see the component.

This component does not implement any kind of persistence or recovery, if the VM terminates while messages are yet to be processed. If you need persistence, reliability or distributed SEDA, try using either [JMS](#) or ActiveMQ.



NOTE

Synchronous

The [Direct](#) component provides synchronous invocation of any consumers when a producer sends a message exchange.

45.1. URI FORMAT

```
seda:someName[?options]
```

Where **someName** can be any string that uniquely identifies the endpoint within the current CamelContext.

45.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

45.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

45.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

45.3. COMPONENT OPTIONS

The SEDA component supports 10 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Sets the default number of concurrent threads processing exchanges.	1	int
defaultPollTimeout (consumer (advanced))	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	int
defaultBlockWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	boolean
defaultDiscardWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	boolean

Name	Description	Default	Type
defaultOfferTimeout (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, where a configured timeout can be added to the block case. Utilizing the <code>.offer(timeout)</code> method of the underlining java queue.		long
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
defaultQueueFactory (advanced)	Sets the default queue factory.		BlockingQueueFactory
queueSize (advanced)	Sets the default maximum capacity of the SEDA queue (i.e., the number of messages it can hold).	1000	int

45.4. ENDPOINT OPTIONS

The SEDA endpoint is configured using URI syntax:

```
seda:name
```

with the following path and query parameters:

45.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required Name of queue.		String

45.4.2. Query Parameters (18 parameters)

Name	Description	Default	Type
size (common)	The maximum capacity of the SEDA queue (i.e., the number of messages it can hold). Will by default use the <code>defaultSize</code> set on the SEDA component.	1000	int
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Number of concurrent threads processing exchanges.	1	int
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
limitConcurrentConsumers (consumer (advanced))	Whether to limit the number of <code>concurrentConsumers</code> to the maximum of 500. By default, an exception will be thrown if an endpoint is configured with a greater number. You can disable that check by turning this option off.	true	boolean

Name	Description	Default	Type
multipleConsumers (consumer (advanced))	Specifies whether multiple consumers are allowed. If enabled, you can use SEDA for Publish-Subscribe messaging. That is, you can send a message to the SEDA queue and have each consumer receive a copy of the message. When enabled, this option should be specified on every consumer endpoint.	false	boolean
pollTimeout (consumer (advanced))	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	int
purgeWhenStopping (consumer (advanced))	Whether to purge the task queue when stopping the consumer/route. This allows to stop faster, as any pending messages on the queue is discarded.	false	boolean
blockWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	boolean
discardIfNoConsumers (producer)	Whether the producer should discard the message (do not add the message to the queue), when sending to a queue with no active consumers. Only one of the options <code>discardIfNoConsumers</code> and <code>failIfNoConsumers</code> can be enabled at the same time.	false	boolean
discardWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	boolean
failIfNoConsumers (producer)	Whether the producer should fail by throwing an exception, when sending to a queue with no active consumers. Only one of the options <code>discardIfNoConsumers</code> and <code>failIfNoConsumers</code> can be enabled at the same time.	false	boolean

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
offerTimeout (producer)	Offer timeout (in milliseconds) can be added to the block case when queue is full. You can disable timeout by using 0 or a negative value.		long
timeout (producer)	Timeout (in milliseconds) before a SEDA producer will stop waiting for an asynchronous task to complete. You can disable timeout by using 0 or a negative value.	30000	long
waitForTaskToComplete (producer)	Option to specify whether the caller should wait for the async task to complete or not before continuing. The following three options are supported: Always, Never or IfReplyExpected. The first two values are self-explanatory. The last value, IfReplyExpected, will only wait if the message is Request Reply based. The default option is IfReplyExpected. Enum values: <ul style="list-style-type: none"> ● Never ● IfReplyExpected ● Always 	IfReplyExpected	WaitForTaskToComplete
queue (advanced)	Define the queue instance which will be used by the endpoint.		BlockingQueue

45.5. CHOOSING BLOCKINGQUEUE IMPLEMENTATION

By default, the SEDA component always instantiates `LinkedBlockingQueue`, but you can use different implementation, you can reference your own `BlockingQueue` implementation, in this case the `size` option is not used

```
<bean id="arrayQueue" class="java.util.ArrayBlockingQueue">
  <constructor-arg index="0" value="10" ><!-- size -->
  <constructor-arg index="1" value="true" ><!-- fairness -->
```

```

</bean>

<!-- ... and later -->
<from>seda:array?queue=#arrayQueue</from>

```

Or you can reference a `BlockingQueueFactory` implementation, 3 implementations are provided `LinkedBlockingQueueFactory`, `ArrayBlockingQueueFactory` and `PriorityBlockingQueueFactory`:

```

<bean id="priorityQueueFactory"
class="org.apache.camel.component.seda.PriorityBlockingQueueFactory">
  <property name="comparator">
    <bean class="org.apache.camel.demo.MyExchangeComparator" />
  </property>
</bean>

<!-- ... and later -->
<from>seda:priority?queueFactory=#priorityQueueFactory&size=100</from>

```

45.6. USE OF REQUEST REPLY

The [SEDA](#) component supports using Request Reply, where the caller will wait for the Async route to complete. For instance:

```

from("mina:tcp://0.0.0.0:9876?textline=true&sync=true").to("seda:input");

from("seda:input").to("bean:processInput").to("bean:createResponse");

```

In the route above, we have a TCP listener on port 9876 that accepts incoming requests. The request is routed to the `seda:input` queue. As it is a Request Reply message, we wait for the response. When the consumer on the `seda:input` queue is complete, it copies the response to the original message response.

45.7. CONCURRENT CONSUMERS

By default, the SEDA endpoint uses a single consumer thread, but you can configure it to use concurrent consumer threads. So instead of thread pools you can use:

```

from("seda:stageName?concurrentConsumers=5").process(...)

```

As for the difference between the two, note a *thread pool* can increase/shrink dynamically at runtime depending on load, whereas the number of concurrent consumers is always fixed.

45.8. THREAD POOLS

Be aware that adding a thread pool to a SEDA endpoint by doing something like:

```

from("seda:stageName").thread(5).process(...)

```

Can wind up with two **BlockQueues**: one from the SEDA endpoint, and one from the workqueue of the thread pool, which may not be what you want. Instead, you might wish to configure a Direct endpoint with a thread pool, which can process messages both synchronously and asynchronously. For example:

```
from("direct:stageName").thread(5).process(...)
```

You can also directly configure number of threads that process messages on a SEDA endpoint using the **concurrentConsumers** option.

45.9. SAMPLE

In the route below we use the SEDA queue to send the request to this async queue to be able to send a fire-and-forget message for further processing in another thread, and return a constant reply in this thread to the original caller.

We send a Hello World message and expects the reply to be OK.

```
@Test
public void testSendAsync() throws Exception {
    MockEndpoint mock = getMockEndpoint("mock:result");
    mock.expectedBodiesReceived("Hello World");

    // START SNIPPET: e2
    Object out = template.requestBody("direct:start", "Hello World");
    assertEquals("OK", out);
    // END SNIPPET: e2

    assertMockEndpointsSatisfied();
}

@Override
protected RouteBuilder createRouteBuilder() throws Exception {
    return new RouteBuilder() {
        // START SNIPPET: e1
        public void configure() throws Exception {
            from("direct:start")
                // send it to the seda queue that is async
                .to("seda:next")
                // return a constant response
                .transform(constant("OK"));

            from("seda:next").to("mock:result");
        }
        // END SNIPPET: e1
    };
}
```

The "Hello World" message will be consumed from the SEDA queue from another thread for further processing. Since this is from a unit test, it will be sent to a **mock** endpoint where we can do assertions in the unit test.

45.10. USING MULTIPLECONSUMERS

In this example we have defined two consumers.

```
@Test
public void testSameOptionsProducerStillOkay() throws Exception {
    getMockEndpoint("mock:foo").expectedBodiesReceived("Hello World");
```

```

getMockEndpoint("mock:bar").expectedBodiesReceived("Hello World");

template.sendBody("seda:foo", "Hello World");

assertMockEndpointsSatisfied();
}

@Override
protected RouteBuilder createRouteBuilder() throws Exception {
    return new RouteBuilder() {
        @Override
        public void configure() throws Exception {
            from("seda:foo?multipleConsumers=true").routeId("foo").to("mock:foo");
            from("seda:foo?multipleConsumers=true").routeId("bar").to("mock:bar");
        }
    };
}
}

```

Since we have specified **multipleConsumers=true** on the seda foo endpoint we can have those two consumers receive their own copy of the message as a kind of pub-sub style messaging.

As the beans are part of an unit test they simply send the message to a mock endpoint.

45.11. EXTRACTING QUEUE INFORMATION

If needed, information such as queue size, etc. can be obtained without using JMX in this fashion:

```

SedaEndpoint seda = context.getEndpoint("seda:xxx");
int size = seda.getExchanges().size();

```

45.12. SPRING BOOT AUTO-CONFIGURATION

When using seda with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-seda-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 11 options, which are listed below.

Name	Description	Default	Type
		t	

Name	Description	Default	Type
<code>camel.component.seda.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.seda.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.seda.concurrent-consumers</code>	Sets the default number of concurrent threads processing exchanges.	1	Integer
<code>camel.component.seda.default-block-when-full</code>	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	Boolean
<code>camel.component.seda.default-discard-when-full</code>	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	Boolean
<code>camel.component.seda.default-offer-timeout</code>	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, where a configured timeout can be added to the block case. Utilizing the <code>.offer(timeout)</code> method of the underlying java queue.		Long
<code>camel.component.seda.default-poll-timeout</code>	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	Integer

Name	Description	Default	Type
camel.component.seda.default-queue-factory	Sets the default queue factory. The option is a <code>org.apache.camel.component.seda.BlockingQueueFactory<org.apache.camel.Exchange></code> type.		BlockingQueueFactory
camel.component.seda.enabled	Whether to enable auto configuration of the seda component. This is enabled by default.		Boolean
camel.component.seda.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.seda.queue-size	Sets the default maximum capacity of the SEDA queue (i.e., the number of messages it can hold).	1000	Integer


CHAPTER 46. SERVLET

Only consumer is supported

The Servlet component provides HTTP based endpoints for consuming HTTP requests that arrive at a HTTP endpoint that is bound to a published Servlet.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-servlet</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```



NOTE

Stream

Servlet is stream based, which means the input it receives is submitted to Camel as a stream. That means you will only be able to read the content of the stream **once**. If you find a situation where the message body appears to be empty or you need to access the data multiple times (eg: doing multicasting, or redelivery error handling) you should use Stream caching or convert the message body to a **String** which is safe to be read multiple times.

46.1. URI FORMAT

```
servlet://relative_path[?options]
```

46.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

46.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre-configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

46.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a *type safe* way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

46.3. COMPONENT OPTIONS

The Servlet component supports 11 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which means any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
muteException (consumer)	If enabled and an Exchange failed processing on the consumer side the response's body won't contain the exception's stack trace.	false	boolean
serviceName (consumer)	Default name of servlet to use. The default name is <code>CamelServlet</code> .	CamelServlet	String
attachmentMultipartBinding (consumer (advanced))	Whether to automatic bind multipart/form-data as attachments on the Camel Exchange. The options <code>attachmentMultipartBinding=true</code> and <code>disableStreamCache=false</code> cannot work together. Remove <code>disableStreamCache</code> to use <code>AttachmentMultipartBinding</code> . This is turned off by default as this may require servlet specific configuration to enable this when using Servlets.	false	boolean
fileNameExtWhitelist (consumer (advanced))	Whitelist of accepted filename extensions for accepting uploaded files. Multiple extensions can be separated by comma, such as <code>txt,xml</code> .		String

Name	Description	Default	Type
httpRegistry (consumer (advanced))	To use a custom <code>org.apache.camel.component.servlet.HttpRegistry</code> .		HttpRegistry
allowJavaSerializedObject (advanced)	Whether to allow java serialization when a request uses <code>context-type=application/x-java-serialized-object</code> . This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
httpBinding (advanced)	To use a custom <code>HttpBinding</code> to control the mapping between Camel message and <code>HttpClient</code> .		HttpBinding
httpConfiguration (advanced)	To use the shared <code>HttpConfiguration</code> as base configuration.		HttpConfiguration
headerFilterStrategy (filter)	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message.		HeaderFilterStrategy

46.4. ENDPOINT OPTIONS

The Servlet endpoint is configured using URI syntax:

```
servlet:contextPath
```

with the following path and query parameters:

46.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
contextPath (consumer)	Required The context-path to use.		String

46.4.2. Query Parameters (22 parameters)

Name	Description	Default	Type
chunked (consumer)	If this option is false the Servlet will disable the HTTP streaming and set the content-length header on the response.	true	boolean
disableStreamCache (common)	Determines whether or not the raw input stream from Servlet is cached or not (Camel will read the stream into a in memory/overflow to file, Stream caching) cache. By default Camel will cache the Servlet input stream to support reading it multiple times to ensure Camel can retrieve all data from the stream. However you can set this option to true when you for example need to access the raw stream, such as streaming it directly to a file or other persistent store. DefaultHttpBinding will copy the request input stream into a stream cache and put it into message body if this option is false to support reading the stream multiple times. If you use Servlet to bridge/proxy an endpoint then consider enabling this option to improve performance, in case you do not need to read the message payload multiple times. The http producer will by default cache the response body stream. If this option is set to true, then the producers will not cache the response body stream but use the response stream as-is as the message body.	false	boolean
headerFilterStrategy (common)	To use a custom HeaderFilterStrategy to filter header to and from Camel message.		HeaderFilterStrategy
httpBinding (common (advanced))	To use a custom HttpBinding to control the mapping between Camel message and HttpClient.		HttpBinding
async (consumer)	Configure the consumer to work in async mode.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which means any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean

Name	Description	Default	Type
httpMethodRestrict (consumer)	Used to only allow consuming if the HttpMethod matches, such as GET/POST/PUT etc. Multiple methods can be specified separated by comma.		String
matchOnUriPrefix (consumer)	Whether or not the consumer should try to find a target consumer by matching the URI prefix if no exact match is found.	false	boolean
muteException (consumer)	If enabled and an Exchange failed processing on the consumer side the response's body won't contain the exception's stack trace.	false	boolean
responseBufferSize (consumer)	To use a custom buffer size on the javax.servlet.ServletResponse.		Integer
serviceName (consumer)	Name of the servlet to use.	Camel Servlet	String
transferException (consumer)	If enabled and an Exchange failed processing on the consumer side, and if the caused Exception was sent back serialized in the response as an application/x-java-serialized-object content type. On the producer side the exception will be deserialized and thrown as is, instead of the HttpOperationFailedException. The caused exception is required to be serialized. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	boolean
attachmentMultipartBinding (consumer (advanced))	Whether to automatic bind multipart/form-data as attachments on the Camel Exchange. The options attachmentMultipartBinding=true and disableStreamCache=false cannot work together. Remove disableStreamCache to use AttachmentMultipartBinding. This is turned off by default as this may require servlet specific configuration to enable this when using Servlets.	false	boolean
eagerCheckContentAvailable (consumer (advanced))	Whether to eager check whether the HTTP requests has content if the content-length header is 0 or not present. This can be turned on in case HTTP clients do not send streamed data.	false	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler

Name	Description	Default	Type
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
fileNameExtWhitelist (consumer (advanced))	Whitelist of accepted filename extensions for accepting uploaded files. Multiple extensions can be separated by comma, such as txt,xml.		String
mapHttpMessageBody (consumer (advanced))	If this option is true then IN exchange Body of the exchange will be mapped to HTTP body. Setting this to false will avoid the HTTP mapping.	true	boolean
mapHttpMessageFormUrlEncodedBody (consumer (advanced))	If this option is true then IN exchange Form Encoded body of the exchange will be mapped to HTTP. Setting this to false will avoid the HTTP Form Encoded body mapping.	true	boolean
mapHttpMessageHeaders (consumer (advanced))	If this option is true then IN exchange Headers of the exchange will be mapped to HTTP headers. Setting this to false will avoid the HTTP Headers mapping.	true	boolean
optionsEnabled (consumer (advanced))	Specifies whether to enable HTTP OPTIONS for this Servlet consumer. By default OPTIONS is turned off.	false	boolean
traceEnabled (consumer (advanced))	Specifies whether to enable HTTP TRACE for this Servlet consumer. By default TRACE is turned off.	false	boolean

46.5. MESSAGE HEADERS

Camel will apply the same Message Headers as the [HTTP](#) component.

Camel will also populate **all request.parameter** and **request.headers**. For example, if a client request has the URL, <http://myserver/myserver?orderid=123>, the exchange will contain a header named **orderid** with the value 123.

46.6. USAGE

You can consume only **from** endpoints generated by the Servlet component. Therefore, it should be used only as input into your Camel routes. To issue HTTP requests against other HTTP endpoints, use the [HTTP](#) component.

46.7. SPRING BOOT AUTO-CONFIGURATION

When using servlet with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-servlet-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 15 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.servlet.allow-java-serialized-object</code>	Whether to allow java serialization when a request uses context-type=application/x-java-serialized-object. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	Boolean
<code>camel.component.servlet.attachment-multipart-binding</code>	Whether to automatic bind multipart/form-data as attachments on the Camel Exchange. The options <code>attachmentMultipartBinding=true</code> and <code>disableStreamCache=false</code> cannot work together. Remove <code>disableStreamCache</code> to use <code>AttachmentMultipartBinding</code> . This is turned off by default as this may require servlet specific configuration to enable this when using Servlet's.	false	Boolean
<code>camel.component.servlet.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean

Name	Description	Default	Type
<code>camel.component.servlet.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which means any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.servlet.enabled</code>	Whether to enable auto configuration of the servlet component. This is enabled by default.		Boolean
<code>camel.component.servlet.filename-ext-whitelist</code>	Whitelist of accepted filename extensions for accepting uploaded files. Multiple extensions can be separated by comma, such as txt,xml.		String
<code>camel.component.servlet.header-filter-strategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message. The option is a <code>org.apache.camel.spi.HeaderFilterStrategy</code> type.		HeaderFilterStrategy
<code>camel.component.servlet.http-binding</code>	To use a custom <code>HttpBinding</code> to control the mapping between Camel message and <code>HttpClient</code> . The option is a <code>org.apache.camel.http.common.HttpBinding</code> type.		HttpBinding
<code>camel.component.servlet.http-configuration</code>	To use the shared <code>HttpConfiguration</code> as base configuration. The option is a <code>org.apache.camel.http.common.HttpConfiguration</code> type.		HttpConfiguration
<code>camel.component.servlet.http-registry</code>	To use a custom <code>org.apache.camel.component.servlet.HttpRegistry</code> . The option is a <code>org.apache.camel.http.common.HttpRegistry</code> type.		HttpRegistry
<code>camel.component.servlet.mute-exception</code>	If enabled and an Exchange failed processing on the consumer side the response's body won't contain the exception's stack trace.	false	Boolean
<code>camel.component.servlet.servlet-name</code>	Default name of servlet to use. The default name is <code>CamelServlet</code> .	CamelServlet	String

Name	Description	Default	Type
<code>camel.servlet.mapping.context-path</code>	Context path used by the servlet component for automatic mapping.	<code>/camel/*</code>	String
<code>camel.servlet.mapping.enabled</code>	Enables the automatic mapping of the servlet component into the Spring web context.	<code>true</code>	Boolean
<code>camel.servlet.mapping.servlet-name</code>	The name of the Camel servlet.	Camel Servlet	String

CHAPTER 47. SLACK

Both producer and consumer are supported

The Slack component allows you to connect to an instance of [Slack](#) and delivers a message contained in the message body via a pre established [Slack incoming webhook](#).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-slack</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

47.1. URI FORMAT

To send a message to a channel.

```
slack:#channel[?options]
```

To send a direct message to a slackuser.

```
slack:@userID[?options]
```

47.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

47.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

47.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

47.3. COMPONENT OPTIONS

The Slack component supports 5 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
token (token)	The token to use.		String
webhookUrl (webhook)	The incoming webhook URL.		String

47.4. ENDPOINT OPTIONS

The Slack endpoint is configured using URI syntax:

```
slack:channel
```

with the following path and query parameters:

47.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
channel (common)	Required The channel name (syntax #name) or slackuser (syntax userName) to send a message directly to an user.		String

47.4.2. Query Parameters (29 parameters)

Name	Description	Default	Type
token (common)	The token to use.		String
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
conversationType (consumer)	Type of conversation. Enum values: <ul style="list-style-type: none"> ● PUBLIC_CHANNEL ● PRIVATE_CHANNEL ● MPIM ● IM 	PUBLIC_CHANNEL	ConversationType
maxResults (consumer)	The Max Result for the poll.	10	String
naturalOrder (consumer)	Create exchanges in natural order (oldest to newest) or not.	false	boolean

Name	Description	Default	Type
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
serverUrl (consumer)	The Server URL of the Slack instance.		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • InOnly • InOut • InOptionalOut 		ExchangePattern
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
iconEmoji (producer)	Deprecated Use a Slack emoji as an avatar.		String
iconUrl (producer)	Deprecated The avatar that the component will use when sending message to a channel or user.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
username (producer)	Deprecated This is the username that the bot will have when sending messages to a channel or user.		String
webhookUrl (producer)	The incoming webhook URL.		String
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long

Name	Description	Default	Type
runLoggingLevel (scheduler)	<p>The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	<p>Time unit for initialDelay and delay options.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● NANOSECONDS ● MICROSECONDS ● MILLISECONDS ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit

Name	Description	Default	Type
<code>useFixedDelay</code> (scheduler)	Controls if fixed delay or fixed rate is used. See <code>ScheduledExecutorService</code> in JDK for details.	true	boolean

47.5. CONFIGURING IN SPRINT XML

The Slack component with XML must be configured as a Spring or Blueprint bean that contains the incoming webhook url or the app token for the integration as a parameter.

```
<bean id="slack" class="org.apache.camel.component.slack.SlackComponent">
  <property name="webhookUrl"
value="https://hooks.slack.com/services/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"/>
  <property name="token" value="xoxb-12345678901-1234567890123-
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx"/>
</bean>
```

For Java you can configure this using Java code.

47.6. EXAMPLE

A CamelContext with Blueprint could be as:

```
<?xml version="1.0" encoding="UTF-8"?>
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0" default-activation="lazy">

  <bean id="slack" class="org.apache.camel.component.slack.SlackComponent">
    <property name="webhookUrl"
value="https://hooks.slack.com/services/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"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/blueprint">
    <route>
      <from uri="direct:test"/>
      <to uri="slack:#channel?iconEmoji=:camel:&username=CamelTest"/>
    </route>
  </camelContext>

</blueprint>
```

47.7. PRODUCER

You can now use a token to send a message instead of WebhookUrl.

```
from("direct:test")
  .to("slack:#random?token=RAW(<YOUR_TOKEN>);");
```

You can now use the Slack API model to create blocks. You can read more about it here <https://api.slack.com/block-kit>.


```

public void testSlackAPIModelMessage() {
    Message message = new Message();
    message.setBlocks(Collections.singletonList(SectionBlock
        .builder()
        .text(MarkdownTextObject
            .builder()
            .text("**Hello from Camel!**")
            .build())
        .build()));

    template.sendBody(test, message);
}

```

47.8. CONSUMER

You can use also a consumer for messages in channel.

```

from("slack://general?token=RAW(<YOUR_TOKEN>)&maxResults=1")
    .to("mock:result");

```

In this way you'll get the last message from general channel. The consumer will take track of the timestamp of the last message consumed and in the next poll it will check from that timestamp.

You'll need to create a Slack app and use it on your workspace.


Use the 'Bot User OAuth Access Token' as token for the consumer endpoint.



NOTE

Add the corresponding history (**channels:history** or **groups:history** or **mpim:history** or **im:history**) and read (**channels:read** or **groups:read** or **mpim:read** or **im:read**) user token scope to your app to grant it permission to view messages in the corresponding channel. You will need to use the `conversationType` option to set it up too (**PUBLIC_CHANNEL**, **PRIVATE_CHANNEL**, **MPIM**, **IM**)

The `naturalOrder` option allows consuming messages from the oldest to the newest. Originally you would get the newest first and consume backward (message 3 ⇒ message 2 ⇒ message 1)



NOTE

You can use the `conversationType` option to read history and messages from a channel that is not only public (**PUBLIC_CHANNEL**, **PRIVATE_CHANNEL**, **MPIM**, **IM**)

47.9. SPRING BOOT AUTO-CONFIGURATION

When using slack with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-slack-starter</artifactId>

```

```

<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 6 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.slack.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.slack.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.slack.enabled</code>	Whether to enable auto configuration of the slack component. This is enabled by default.		Boolean
<code>camel.component.slack.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.slack.token</code>	The token to use.		String
<code>camel.component.slack.webhook-url</code>	The incoming webhook URL.		String

CHAPTER 48. SQL

Both producer and consumer are supported

The SQL component allows you to work with databases using JDBC queries. The difference between this component and [JDBC](#) component is that in case of SQL the query is a property of the endpoint and it uses message payload as parameters passed to the query.

This component uses **spring-jdbc** behind the scenes for the actual SQL handling.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-sql</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

The SQL component also supports:

- a JDBC based repository for the Idempotent Consumer EIP pattern. See further below.
- a JDBC based repository for the Aggregator EIP pattern. See further below.

48.1. URI FORMAT



NOTE

This component can be used as a [Transactional Client](#).

The SQL component uses the following endpoint URI notation:

```
sql:select * from table where id=# order by name[?options]
```

You can use named parameters by using: `#name_of_the_parameter` style as shown:

```
sql:select * from table where id=:#myId order by name[?options]
```

When using named parameters, Camel will lookup the names from, in the given precedence:

1. from message body if its a **java.util.Map**
2. from message headers

If a named parameter cannot be resolved, then an exception is thrown.

You can use Simple expressions as parameters as shown:

```
sql:select * from table where id=:${exchangeProperty.myId} order by name[?options]
```

Notice that the standard `?` symbol that denotes the parameters to an SQL query is substituted with the `#` symbol, because the `?` symbol is used to specify options for the endpoint. The `?` symbol replacement can be configured on endpoint basis.

You can externalize your SQL queries to files in the classpath or file system as shown:

```
sql:classpath:sql/myquery.sql[?options]
```

And the `myquery.sql` file is in the classpath and is just a plain text

```
-- this is a comment
select *
from table
where
  id = :#{exchangeProperty.myId}
order by
  name
```

In the file you can use multilines and format the SQL as you wish. And also use comments such as the dash line.

48.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

48.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

48.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

48.3. COMPONENT OPTIONS

The SQL component supports 5 options, which are listed below.

Name	Description	Default	Type
dataSource (common)	Autowired Sets the DataSource to use to communicate with the database.		DataSource
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
usePlaceholder (advanced)	Sets whether to use placeholder and replace all placeholder characters with sign in the SQL queries. This option is default true.	true	boolean

48.4. ENDPOINT OPTIONS

The SQL endpoint is configured using URI syntax:

```
sql:query
```

with the following path and query parameters:

48.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
query (common)	Required Sets the SQL query to perform. You can externalize the query by using file: or classpath: as prefix and specify the location of the file.		String

48.4.2. Query Parameters (45 parameters)

Name	Description	Default	Type
allowNamedParameters (common)	Whether to allow using named parameters in the queries.	true	boolean
dataSource (common)	Autowired Sets the DataSource to use to communicate with the database at endpoint level.		DataSource
outputClass (common)	Specify the full package and class name to use as conversion when outputType=SelectOne.		String
outputHeader (common)	Store the query result in a header instead of the message body. By default, outputHeader == null and the query result is stored in the message body, any existing content in the message body is discarded. If outputHeader is set, the value is used as the name of the header to store the query result and the original message body is preserved.		String

Name	Description	Default	Type
outputType (common)	<p>Make the output of consumer or producer to SelectList as List of Map, or SelectOne as single Java object in the following way: a) If the query has only single column, then that JDBC Column object is returned. (such as SELECT COUNT() FROM PROJECT will return a Long object. b) If the query has more than one column, then it will return a Map of that result. c) If the outputClass is set, then it will convert the query result into an Java bean object by calling all the setters that match the column names. It will assume your class has a default constructor to create instance with. d) If the query resulted in more than one rows, it throws an non-unique result exception. StreamList streams the result of the query using an Iterator. This can be used with the Splitter EIP in streaming mode to process the ResultSet in streaming fashion.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● SelectOne ● SelectList ● StreamList 	Select List	SqlOutputType
separator (common)	The separator to use when parameter values is taken from message body (if the body is a String type), to be inserted at # placeholders. Notice if you use named parameters, then a Map type is used instead. The default value is comma.	,	char
breakBatchOnConsumeFail (consumer)	Sets whether to break batch if onConsume failed.	false	boolean
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
expectedUpdateCount (consumer)	Sets an expected update count to validate when using onConsume.	-1	int
maxMessagesPerPoll (consumer)	Sets the maximum number of messages to poll.		int

Name	Description	Default	Type
onConsume (consumer)	After processing each row then this query can be executed, if the Exchange was processed successfully, for example to mark the row as processed. The query can have parameter.		String
onConsumeBatchComplete (consumer)	After processing the entire batch, this query can be executed to bulk update rows etc. The query cannot have parameters.		String
onConsumeFailed (consumer)	After processing each row then this query can be executed, if the Exchange failed, for example to mark the row as failed. The query can have parameter.		String
routeEmptyResultSet (consumer)	Sets whether empty resultset should be allowed to be sent to the next hop. Defaults to false. So the empty resultset will be filtered out.	false	boolean
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
transacted (consumer)	Enables or disables transaction. If enabled then if processing an exchange failed then the consumer breaks out processing any further exchanges to cause a rollback eager.	false	boolean
useIterator (consumer)	Sets how resultset should be delivered to route. Indicates delivery as either a list or individual object. defaults to true.	true	boolean
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option errorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

Name	Description	Default	Type
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		PollingConsumerPollStrategy
processingStrategy (consumer (advanced))	Allows to plugin to use a custom <code>org.apache.camel.component.sql.SqlProcessingStrategy</code> to execute queries when the consumer has processed the rows/batch.		SqlProcessingStrategy
batch (producer)	Enables or disables batch mode.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
noop (producer)	If set, will ignore the results of the SQL query and use the existing IN message as the OUT message for the continuation of processing.	false	boolean
useMessageBodyForSql (producer)	Whether to use the message body as the SQL and then headers for parameters. If this option is enabled then the SQL in the uri is not used. Note that query parameters in the message body are represented by a question mark instead of a # symbol.	false	boolean
alwaysPopulateStatement (advanced)	If enabled then the <code>populateStatement</code> method from <code>org.apache.camel.component.sql.SqlPrepareStatementStrategy</code> is always invoked, also if there is no expected parameters to be prepared. When this is false then the <code>populateStatement</code> is only invoked if there is 1 or more expected parameters to be set; for example this avoids reading the message body/headers for SQL queries with no parameters.	false	boolean

Name	Description	Default	Type
parametersCount (advanced)	If set greater than zero, then Camel will use this count value of parameters to replace instead of querying via JDBC metadata API. This is useful if the JDBC vendor could not return correct parameters count, then user may override instead.		int
placeholder (advanced)	Specifies a character that will be replaced to in SQL query. Notice, that it is simple String.replaceAll() operation and no SQL parsing is involved (quoted strings will also change).	#	String
prepareStatementStrategy (advanced)	Allows to plugin to use a custom org.apache.camel.component.sql.SqlPreparedStatementStrategy to control preparation of the query and prepared statement.		SqlPreparedStatementStrategy
templateOptions (advanced)	Configures the Spring JdbcTemplate with the key/values from the Map.		Map
usePlaceholder (advanced)	Sets whether to use placeholder and replace all placeholder characters with sign in the SQL queries.	true	boolean
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long

Name	Description	Default	Type
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean

Name	Description	Default	Type
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> • NANOSECONDS • MICROSECONDS • MILLISECONDS • SECONDS • MINUTES • HOURS • DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean

48.5. TREATMENT OF THE MESSAGE BODY

The SQL component tries to convert the message body to an object of **java.util.Iterator** type and then uses this iterator to fill the query parameters (where each query parameter is represented by a # symbol (or configured placeholder) in the endpoint URI). If the message body is not an array or collection, the conversion results in an iterator that iterates over only one object, which is the body itself.

For example, if the message body is an instance of **java.util.List**, the first item in the list is substituted into the first occurrence of # in the SQL query, the second item in the list is substituted into the second occurrence of #, and so on.

If **batch** is set to **true**, then the interpretation of the inbound message body changes slightly – instead of an iterator of parameters, the component expects an iterator that contains the parameter iterators; the size of the outer iterator determines the batch size.

You can use the option **useMessageBodyForSql** that allows to use the message body as the SQL statement, and then the SQL parameters must be provided in a header with the key **SqlConstants.SQL_PARAMETERS**. This allows the SQL component to work more dynamically as the SQL query is from the message body. Use templating (such as [Velocity](#), [Freemarker](#)) for conditional processing, e.g. to include or exclude **where** clauses depending on the presence of query parameters.

48.6. RESULT OF THE QUERY

For **select** operations, the result is an instance of **List<Map<String, Object>>** type, as returned by the [JdbcTemplate.queryForList\(\)](#) method. For **update** operations, a **NULL** body is returned as the **update** operation is only set as a header and never as a body.



NOTE

See [Header Values](#) for more information on the **update** operation.

By default, the result is placed in the message body. If the `outputHeader` parameter is set, the result is placed in the header. This is an alternative to using a full message enrichment pattern to add headers, it provides a concise syntax for querying a sequence or some other small value into a header. It is convenient to use `outputHeader` and `outputType` together:

```
from("jms:order.inbox")
  .to("sql:select order_seq.nextval from dual?outputHeader=OrderId&outputType=SelectOne")
  .to("jms:order.booking");
```

48.7. USING STREAMLIST

The producer supports `outputType=StreamList` that uses an iterator to stream the output of the query. This allows to process the data in a streaming fashion which for example can be used by the Splitter EIP to process each row one at a time, and load data from the database as needed.

```
from("direct:withSplitModel")
  .to("sql:select * from projects order by id?
outputType=StreamList&outputClass=org.apache.camel.component.sql.ProjectModel")
  .to("log:stream")
  .split(body()).streaming()
  .to("log:row")
  .to("mock:result")
  .end();
```

48.8. HEADER VALUES

When performing **update** operations, the SQL Component stores the update count in the following message headers:

Header	Description
CamelSqlUpdateCount	The number of rows updated for update operations, returned as an Integer object. This header is not provided when using <code>outputType=StreamList</code> .
CamelSqlRowCount	The number of rows returned for select operations, returned as an Integer object. This header is not provided when using <code>outputType=StreamList</code> .
CamelSqlQuery	Query to execute. This query takes precedence over the query specified in the endpoint URI. Note that query parameters in the header <i>are</i> represented by a ? instead of a # symbol

When performing **insert** operations, the SQL Component stores the rows with the generated keys and number of these rows in the following message headers:

Header	Description
CamelSqlGeneratedKeysRowCount	The number of rows in the header that contains generated keys.

Header	Description
CamelSqlGeneratedKey Rows	Rows that contains the generated keys (a list of maps of keys).

48.9. GENERATED KEYS

If you insert data using SQL INSERT, then the RDBMS may support auto generated keys. You can instruct the SQL producer to return the generated keys in headers.

To do that set the header **CamelSqlRetrieveGeneratedKeys=true**. Then the generated keys will be provided as headers with the keys listed in the table above.

To specify which generated columns should be retrieved, set the header **CamelSqlGeneratedColumns** to a **String[]** or **int[]**, indicating the column names or indexes, respectively. Some databases requires this, such as Oracle. It may also be necessary to use the **parametersCount** option if the driver cannot correctly determine the number of parameters.

You can see more details in this [unit test](#).

48.10. DATASOURCE

You can set a reference to a **DataSource** in the URI directly:

```
sql:select * from table where id=# order by name?dataSource=#myDS
```

48.11. USING NAMED PARAMETERS

In the given route below, we want to get all the projects from the projects table. Notice the SQL query has 2 named parameters, `:#lic` and `:#min`.

Camel will then lookup for these parameters from the message body or message headers. Notice in the example above we set two headers with constant value for the named parameters:

```
from("direct:projects")
  .setHeader("lic", constant("ASF"))
  .setHeader("min", constant(123))
  .to("sql:select * from projects where license = :#lic and id > :#min order by id")
```

Though if the message body is a **java.util.Map** then the named parameters will be taken from the body.

```
from("direct:projects")
  .to("sql:select * from projects where license = :#lic and id > :#min order by id")
```

48.12. USING EXPRESSION PARAMETERS IN PRODUCERS

In the given route below, we want to get all the project from the database. It uses the body of the exchange for defining the license and uses the value of a property as the second parameter.

```

from("direct:projects")
  .setBody(constant("ASF"))
  .setProperty("min", constant(123))
  .to("sql:select * from projects where license = :#{body} and id > :#{exchangeProperty.min} order
by id")

```

48.12.1. Using expression parameters in consumers

When using the SQL component as consumer, you can now also use expression parameters (simple language) to build dynamic query parameters, such as calling a method on a bean to retrieve an id, date or something.

For example in the sample below we call the `nextId` method on the bean `myIdGenerator`:

```

from("sql:select * from projects where id = :#{bean:myIdGenerator.nextId}")
  .to("mock:result");

```

And the bean has the following method:

```

public static class MyIdGenerator {

    private int id = 1;

    public int nextId() {
        return id++;
    }
}

```

Notice that there is no existing **Exchange** with message body and headers, so the simple expression you can use in the consumer are most useable for calling bean methods as in this example.

48.13. USING IN QUERIES WITH DYNAMIC VALUES

The SQL producer allows to use SQL queries with IN statements where the IN values is dynamic computed. For example from the message body or a header etc.

To use IN you need to:

- prefix the parameter name with **in**:
- add **()** around the parameter

An example explains this better. The following query is used:

```

-- this is a comment
select *
from projects
where project in (:#in:names)
order by id

```

In the following route:

```

from("direct:query")
  .to("sql:classpath:sql/selectProjectsIn.sql")

```

```
.to("log:query")
.to("mock:query");
```

Then the IN query can use a header with the key names with the dynamic values such as:

```
// use an array
template.requestBodyAndHeader("direct:query", "Hi there!", "names", new String[]{"Camel", "AMQ"});

// use a list
List<String> names = new ArrayList<String>();
names.add("Camel");
names.add("AMQ");

template.requestBodyAndHeader("direct:query", "Hi there!", "names", names);

// use a string separated values with comma
template.requestBodyAndHeader("direct:query", "Hi there!", "names", "Camel,AMQ");
```

The query can also be specified in the endpoint instead of being externalized (notice that externalizing makes maintaining the SQL queries easier)

```
from("direct:query")
.to("sql:select * from projects where project in (:#in:names) order by id")
.to("log:query")
.to("mock:query");
```

48.14. USING THE JDBC BASED IDEMPOTENT REPOSITORY

In this section we will use the JDBC based idempotent repository.



NOTE

Abstract class

There is an abstract class

org.apache.camel.processor.idempotent.jdbc.AbstractJdbcMessageIdRepository
you can extend to build custom JDBC idempotent repository.

First we have to create the database table which will be used by the idempotent repository. We use the following schema:

```
CREATE TABLE CAMEL_MESSAGEPROCESSED ( processorName VARCHAR(255),
messageld VARCHAR(100) )
```

We added the createdAt column:

```
CREATE TABLE CAMEL_MESSAGEPROCESSED ( processorName VARCHAR(255),
messageld VARCHAR(100), createdAt TIMESTAMP )
```



NOTE

The SQL Server **TIMESTAMP** type is a fixed-length binary-string type. It does not map to any of the JDBC time types: **DATE**, **TIME**, or **TIMESTAMP**.

When working with concurrent consumers it is crucial to create a unique constraint on the columns `processorName` and `messageId`. Because the syntax for this constraint differs from database to database, we do not show it here.

48.14.1. Customize the JDBC idempotency repository

You have a few options to tune the `org.apache.camel.processor.idempotent.jdbc.JdbcMessageIdRepository` for your needs:

Parameter	Default Value	Description
<code>createTableIfNotExists</code>	<code>true</code>	Defines whether or not Camel should try to create the table if it doesn't exist.
<code>tableName</code>	<code>CAMEL_MESSAGEPROCESSED</code>	To use a custom table name instead of the default name: <code>CAMEL_MESSAGEPROCESSED</code> .
<code>tableExistsString</code>	<code>SELECT 1 FROM CAMEL_MESSAGEPROCESSED WHERE 1 = 0</code>	This query is used to figure out whether the table already exists or not. It must throw an exception to indicate the table doesn't exist.
<code>createString</code>	<code>CREATE TABLE CAMEL_MESSAGEPROCESSED (processorName VARCHAR(255), messageId VARCHAR(100), createdAt TIMESTAMP)</code>	The statement which is used to create the table.
<code>queryString</code>	<code>SELECT COUNT(*) FROM CAMEL_MESSAGEPROCESSED WHERE processorName = ? AND messageId = ?</code>	The query which is used to figure out whether the message already exists in the repository (the result is not equals to '0'). It takes two parameters. This first one is the processor name (String) and the second one is the message id (String).
<code>insertString</code>	<code>INSERT INTO CAMEL_MESSAGEPROCESSED (processorName, messageId, createdAt) VALUES (?, ?, ?)</code>	The statement which is used to add the entry into the table. It takes three parameter. The first one is the processor name (String), the second one is the message id (String) and the third one is the timestamp (java.sql.Timestamp) when this entry was added to the repository.
<code>deleteString</code>	<code>DELETE FROM CAMEL_MESSAGEPROCESSED WHERE processorName = ? AND messageId = ?</code>	The statement which is used to delete the entry from the database. It takes two parameter. This first one is the processor name (String) and the second one is the message id (String).

The option **tableName** can be used to use the default SQL queries but with a different table name. However if you want to customize the SQL queries then you can configure each of them individually.

48.14.2. Orphan Lock aware Jdbc IdempotentRepository

One of the limitations of **org.apache.camel.processor.idempotent.jdbc.JdbcMessageIdRepository** is that it does not handle orphan locks resulting from JVM crash or non graceful shutdown. This can result in unprocessed files/messages if this implementation is used with camel-file, camel-ftp etc. if you need to address orphan locks processing then use **org.apache.camel.processor.idempotent.jdbc.JdbcOrphanLockAwareIdempotentRepository**. This repository keeps track of the locks held by an instance of the application. For each lock held, the application will send keep alive signals to the lock repository resulting in updating the createdAt column with the current Timestamp. When an application instance tries to acquire a lock if the, then there are three possibilities exist :

- lock entry does not exist then the lock is provided using the base implementation of **JdbcMessageIdRepository**.
- lock already exists and the `createdAt < System.currentTimeMillis() - lockMaxAgeMillis`. In this case it is assumed that an active instance has the lock and the lock is not provided to the new instance requesting the lock
- lock already exists and the `createdAt >= System.currentTimeMillis() - lockMaxAgeMillis`. In this case it is assumed that there is no active instance which has the lock and the lock is provided to the requesting instance. The reason behind is that if the original instance which had the lock, if it was still running, it would have updated the Timestamp on createdAt using its keepAlive mechanism

This repository has two additional configuration parameters

Parameter	Description
lockMaxAgeMillis	This refers to the duration after which the lock is considered orphaned i.e. if the <code>currentTimestamp - createdAt >= lockMaxAgeMillis</code> then lock is orphaned.
lockKeepAliveIntervalMillis	The frequency at which keep alive updates are done to createdAt Timestamp column.

48.14.3. Caching Jdbc IdempotentRepository

Some SQL implementations are not fast on a per query basis. The **JdbcMessageIdRepository** implementation does its idempotent checks individually within SQL transactions. Checking a mere 100 keys can take minutes. The **JdbcCachedMessageIdRepository** preloads an in-memory cache on start with the entire list of keys. This cache is then checked first before passing through to the original implementation.

As with all cache implementations, there are considerations that should be made with regard to stale data and your specific usage.

48.15. USING THE JDBC BASED AGGREGATION REPOSITORY

JdbcAggregationRepository is an **AggregationRepository** which on the fly persists the aggregated messages. This ensures that you will not loose messages, as the default aggregator will use an in memory only **AggregationRepository**. The **JdbcAggregationRepository** allows together with Camel to provide persistent support for the Aggregator.

Only when an Exchange has been successfully processed it will be marked as complete which happens when the **confirm** method is invoked on the **AggregationRepository**. This means if the same Exchange fails again it will be kept retried until it success.

You can use option **maximumRedeliveries** to limit the maximum number of redelivery attempts for a given recovered Exchange. You must also set the **deadLetterUri** option so Camel knows where to send the Exchange when the **maximumRedeliveries** was hit.

You can see some examples in the unit tests of camel-sql, for example **JdbcAggregateRecoverDeadLetterChannelTest.java**

48.15.1. Database

To be operational, each aggregator uses two table: the aggregation and completed one. By convention the completed has the same name as the aggregation one suffixed with "**_COMPLETED**". The name must be configured in the Spring bean with the **RepositoryName** property. In the following example aggregation will be used.

The table structure definition of both table are identical: in both case a String value is used as key (**id**) whereas a Blob contains the exchange serialized in byte array.

However one difference should be remembered: the **id** field does not have the same content depending on the table.

In the aggregation table **id** holds the correlation Id used by the component to aggregate the messages. In the completed table, **id** holds the id of the exchange stored in corresponding the blob field.

Here is the SQL query used to create the tables, just replace "**aggregation**" with your aggregator repository name.

```
CREATE TABLE aggregation (
  id varchar(255) NOT NULL,
  exchange blob NOT NULL,
  version BIGINT NOT NULL,
  constraint aggregation_pk PRIMARY KEY (id)
);
CREATE TABLE aggregation_completed (
  id varchar(255) NOT NULL,
  exchange blob NOT NULL,
  version BIGINT NOT NULL,
  constraint aggregation_completed_pk PRIMARY KEY (id)
);
```

48.16. STORING BODY AND HEADERS AS TEXT

You can configure the **JdbcAggregationRepository** to store message body and select(ed) headers as String in separate columns. For example to store the body, and the following two headers **companyName** and **accountName** use the following SQL:

```
CREATE TABLE aggregationRepo3 (
  id varchar(255) NOT NULL,
  exchange blob NOT NULL,
  version BIGINT NOT NULL,
  body varchar(1000),
  companyName varchar(1000),
  accountName varchar(1000),
  constraint aggregationRepo3_pk PRIMARY KEY (id)
```

```
);
CREATE TABLE aggregationRepo3_completed (
  id varchar(255) NOT NULL,
  exchange blob NOT NULL,
  version BIGINT NOT NULL,
  body varchar(1000),
  companyName varchar(1000),
  accountName varchar(1000),
  constraint aggregationRepo3_completed_pk PRIMARY KEY (id)
);
```

And then configure the repository to enable this behavior as shown below:

```
<bean id="repo3"
  class="org.apache.camel.processor.aggregate.jdbc.JdbcAggregationRepository">
  <property name="repositoryName" value="aggregationRepo3"/>
  <property name="transactionManager" ref="txManager3"/>
  <property name="dataSource" ref="dataSource3"/>
  <!-- configure to store the message body and following headers as text in the repo -->
  <property name="storeBodyAsText" value="true"/>
  <property name="headersToStoreAsText">
    <list>
      <value>companyName</value>
      <value>accountName</value>
    </list>
  </property>
</bean>
```

48.16.1. Codec (Serialization)

Since they can contain any type of payload, Exchanges are not serializable by design. It is converted into a byte array to be stored in a database BLOB field. All those conversions are handled by the **JdbcCodec** class. One detail of the code requires your attention: the **ClassLoaderAwareObjectInputStream**.

The **ClassLoaderAwareObjectInputStream** has been reused from the [Apache ActiveMQ](#) project. It wraps an **ObjectInputStream** and use it with the **ContextClassLoader** rather than the **currentThread** one. The benefit is to be able to load classes exposed by other bundles. This allows the exchange body and headers to have custom types object references.

48.16.2. Transaction

A Spring **PlatformTransactionManager** is required to orchestrate transaction.

48.16.2.1. Service (Start/Stop)

The **start** method verify the connection of the database and the presence of the required tables. If anything is wrong it will fail during starting.

48.16.3. Aggregator configuration

Depending on the targeted environment, the aggregator might need some configuration. As you already know, each aggregator should have its own repository (with the corresponding pair of table created in the database) and a data source. If the default lobHandler is not adapted to your database system, it can be injected with the **lobHandler** property.

Here is the declaration for Oracle:

```
<bean id="lobHandler" class="org.springframework.jdbc.support.lob.OracleLobHandler">
  <property name="nativeJdbcExtractor" ref="nativeJdbcExtractor"/>
</bean>
<bean id="nativeJdbcExtractor"
  class="org.springframework.jdbc.support.nativejdbc.CommonsDbcpNativeJdbcExtractor"/>
<bean id="repo"
  class="org.apache.camel.processor.aggregate.jdbc.JdbcAggregationRepository">
  <property name="transactionManager" ref="transactionManager"/>
  <property name="repositoryName" value="aggregation"/>
  <property name="dataSource" ref="dataSource"/>
  <!-- Only with Oracle, else use default -->
  <property name="lobHandler" ref="lobHandler"/>
</bean>
```

48.16.4. Optimistic locking

You can turn on **optimisticLocking** and use this JDBC based aggregation repository in a clustered environment where multiple Camel applications shared the same database for the aggregation repository. If there is a race condition there JDBC driver will throw a vendor specific exception which the **JdbcAggregationRepository** can react upon. To know which caused exceptions from the JDBC driver is regarded as an optimistic locking error we need a mapper to do this. Therefore there is a **org.apache.camel.processor.aggregate.jdbc.JdbcOptimisticLockingExceptionMapper** allows you to implement your custom logic if needed. There is a default implementation **org.apache.camel.processor.aggregate.jdbc.DefaultJdbcOptimisticLockingExceptionMapper** which works as follows:

The following check is done:

- If the caused exception is an **SQLException** then the SQLState is checked if starts with 23.
- If the caused exception is a **DataIntegrityViolationException**
- If the caused exception class name has "ConstraintViolation" in its name.
- Optional checking for FQN class name matches if any class names has been configured.

You can in addition add FQN classnames, and if any of the caused exception (or any nested) equals any of the FQN class names, then its an optimistic locking error.

Here is an example, where we define 2 extra FQN class names from the JDBC vendor.

```
<bean id="repo"
  class="org.apache.camel.processor.aggregate.jdbc.JdbcAggregationRepository">
  <property name="transactionManager" ref="transactionManager"/>
  <property name="repositoryName" value="aggregation"/>
  <property name="dataSource" ref="dataSource"/>
  <property name="jdbcOptimisticLockingExceptionMapper" ref="myExceptionMapper"/>
</bean>
<!-- use the default mapper with extraFQN class names from our JDBC driver -->
<bean id="myExceptionMapper"
  class="org.apache.camel.processor.aggregate.jdbc.DefaultJdbcOptimisticLockingExceptionMapper">
  <property name="classNames">
    <util:set>
```

```

    <value>com.foo.sql.MyViolationExceptoion</value>
    <value>com.foo.sql.MyOtherViolationExceptoion</value>
  </util:set>
</property>
</bean>

```

48.16.5. Propagation behavior

JdbcAggregationRepository uses two distinct *transaction templates* from Spring-TX. One is read-only and one is used for read-write operations.

However, when using **JdbcAggregationRepository** within a route that itself uses `<transacted />` and there's common **PlatformTransactionManager** used, there may be a need to configure *propagation behavior* used by transaction templates inside **JdbcAggregationRepository**.

Here's a way to do it:

```

<bean id="repo"
class="org.apache.camel.processor.aggregate.jdbc.JdbcAggregationRepository">
  <property name="propagationBehaviorName" value="PROPAGATION_NESTED" />
</bean>

```

Propagation is specified by constants of **org.springframework.transaction.TransactionDefinition** interface, so **propagationBehaviorName** is convenient setter that allows to use names of the constants.

48.16.6. PostgreSQL case

There's special database that may cause problems with optimistic locking used by **JdbcAggregationRepository**. PostgreSQL marks connection as invalid in case of data integrity violation exception (the one with `SQLState 23505`). This makes the connection effectively unusable within nested transaction. Details can be found in the [document](#).

org.apache.camel.processor.aggregate.jdbc.PostgresAggregationRepository extends **JdbcAggregationRepository** and uses special **INSERT .. ON CONFLICT ..** statement to provide optimistic locking behavior.

This statement is (with default aggregation table definition):

```

INSERT INTO aggregation (id, exchange) values (?, ?) ON CONFLICT DO NOTHING

```

Details can be found in [PostgreSQL documentation](#).

When this clause is used, **java.sql.PreparedStatement.executeUpdate()** call returns **0** instead of throwing `SQLException` with `SQLState=23505`. Further handling is exactly the same as with generic **JdbcAggregationRepository**, but without marking PostgreSQL connection as invalid.

48.17. CAMEL SQL STARTER

A starter module is available to spring-boot users. When using the starter, the **DataSource** can be directly configured using spring-boot properties.

```

# Example for a mysql datasource
spring.datasource.url=jdbc:mysql://localhost/test

```

```
spring.datasource.username=dbuser
spring.datasource.password=dbpass
spring.datasource.driver-class-name=com.mysql.jdbc.Driver
```

To use this feature, add the following dependencies to your spring boot pom.xml file:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-sql-starter</artifactId>
  <version>${camel.version}</version> <!-- use the same version as your Camel core version -->
</dependency>

<dependency>
  <groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-jdbc</artifactId>
  <version>${spring-boot-version}</version>
</dependency>
```

You should also include the specific database driver, if needed.

48.18. SPRING BOOT AUTO-CONFIGURATION

When using sql with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-sql-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 8 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.sql-stored.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.sql-stored.enabled</code>	Whether to enable auto configuration of the sql-stored component. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.component.sql-stored.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.sql.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.sql.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.sql.enabled</code>	Whether to enable auto configuration of the sql component. This is enabled by default.		Boolean
<code>camel.component.sql.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.sql.use-placeholder</code>	Sets whether to use placeholder and replace all placeholder characters with sign in the SQL queries. This option is default true.	true	Boolean

CHAPTER 49. STUB

Both producer and consumer are supported

The Stub component provides a simple way to stub out any physical endpoints while in development or testing, allowing you for example to run a route without needing to actually connect to a specific [SMTP](#) or [HTTP](#) endpoint. Just add **stub:** in front of any endpoint URI to stub out the endpoint.

Internally the Stub component creates [VM](#) endpoints. The main difference between Stub and [VM](#) is that VM will validate the URI and parameters you give it, so putting `vm:` in front of a typical URI with query arguments will usually fail. Stub won't though, as it basically ignores all query parameters to let you quickly stub out one or more endpoints in your route temporarily.

49.1. URI FORMAT

```
stub:someUri
```

Where **someUri** can be any URI with any query parameters.

49.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

49.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

49.2.1.1. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

49.3. COMPONENT OPTIONS

The Stub component supports 10 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Sets the default number of concurrent threads processing exchanges.	1	int
defaultPollTimeout (consumer (advanced))	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	int
defaultBlockWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	boolean
defaultDiscardWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	boolean
defaultOfferTimeout (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, where a configured timeout can be added to the block case. Utilizing the <code>.offer(timeout)</code> method of the underlining java queue.		long

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
defaultQueueFactory (advanced)	Sets the default queue factory.		BlockingQueueFactory
queueSize (advanced)	Sets the default maximum capacity of the SEDA queue (i.e., the number of messages it can hold).	1000	int

49.4. ENDPOINT OPTIONS

The Stub endpoint is configured using URI syntax:

```
stub:name
```

with the following path and query parameters:

49.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
name (common)	Required Name of queue.		String

49.4.2. Query Parameters (18 parameters)

Name	Description	Default	Type
size (common)	The maximum capacity of the SEDA queue (i.e., the number of messages it can hold). Will by default use the <code>defaultSize</code> set on the SEDA component.	1000	int
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
concurrentConsumers (consumer)	Number of concurrent threads processing exchanges.	1	int
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern
limitConcurrentConsumers (consumer (advanced))	Whether to limit the number of <code>concurrentConsumers</code> to the maximum of 500. By default, an exception will be thrown if an endpoint is configured with a greater number. You can disable that check by turning this option off.	true	boolean
multipleConsumers (consumer (advanced))	Specifies whether multiple consumers are allowed. If enabled, you can use SEDA for Publish-Subscribe messaging. That is, you can send a message to the SEDA queue and have each consumer receive a copy of the message. When enabled, this option should be specified on every consumer endpoint.	false	boolean

Name	Description	Default	Type
pollTimeout (consumer (advanced))	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	int
purgeWhenStopping (consumer (advanced))	Whether to purge the task queue when stopping the consumer/route. This allows to stop faster, as any pending messages on the queue is discarded.	false	boolean
blockWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	boolean
discardIfNoConsumers (producer)	Whether the producer should discard the message (do not add the message to the queue), when sending to a queue with no active consumers. Only one of the options <code>discardIfNoConsumers</code> and <code>failIfNoConsumers</code> can be enabled at the same time.	false	boolean
discardWhenFull (producer)	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	boolean
failIfNoConsumers (producer)	Whether the producer should fail by throwing an exception, when sending to a queue with no active consumers. Only one of the options <code>discardIfNoConsumers</code> and <code>failIfNoConsumers</code> can be enabled at the same time.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean

Name	Description	Default	Type
offerTimeout (producer)	Offer timeout (in milliseconds) can be added to the block case when queue is full. You can disable timeout by using 0 or a negative value.		long
timeout (producer)	Timeout (in milliseconds) before a SEDA producer will stop waiting for an asynchronous task to complete. You can disable timeout by using 0 or a negative value.	30000	long
waitForTaskToComplete (producer)	Option to specify whether the caller should wait for the async task to complete or not before continuing. The following three options are supported: Always, Never or IfReplyExpected. The first two values are self-explanatory. The last value, IfReplyExpected, will only wait if the message is Request Reply based. The default option is IfReplyExpected. Enum values: <ul style="list-style-type: none"> • Never • IfReplyExpected • Always 	IfReplyExpected	WaitForTaskToComplete
queue (advanced)	Define the queue instance which will be used by the endpoint.		BlockingQueue

49.5. EXAMPLES

Here are a few samples of stubbing endpoint uris

```
stub:smtp://somehost.foo.com?user=whatnot&something=else
stub:http://somehost.bar.com/something
```

49.6. SPRING BOOT AUTO-CONFIGURATION

When using stub with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-stub-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 11 options, which are listed below.

Name	Description	Default	Type
camel.component.stub.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.stub.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.stub.concurrent-consumers	Sets the default number of concurrent threads processing exchanges.	1	Integer
camel.component.stub.default-block-when-full	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will instead block and wait until the message can be accepted.	false	Boolean
camel.component.stub.default-discard-when-full	Whether a thread that sends messages to a full SEDA queue will be discarded. By default, an exception will be thrown stating that the queue is full. By enabling this option, the calling thread will give up sending and continue, meaning that the message was not sent to the SEDA queue.	false	Boolean
camel.component.stub.default-offer-timeout	Whether a thread that sends messages to a full SEDA queue will block until the queue's capacity is no longer exhausted. By default, an exception will be thrown stating that the queue is full. By enabling this option, where a configured timeout can be added to the block case. Utilizing the .offer(timeout) method of the underlining java queue.		Long
camel.component.stub.default-poll-timeout	The timeout (in milliseconds) used when polling. When a timeout occurs, the consumer can check whether it is allowed to continue running. Setting a lower value allows the consumer to react more quickly upon shutdown.	1000	Integer

Name	Description	Default	Type
camel.component.stub.default-queue-factory	Sets the default queue factory. The option is a <code>org.apache.camel.component.seda.BlockingQueueFactory<org.apache.camel.Exchange></code> type.		BlockingQueueFactory
camel.component.stub.enabled	Whether to enable auto configuration of the stub component. This is enabled by default.		Boolean
camel.component.stub.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.stub.queue-size	Sets the default maximum capacity of the SEDA queue (i.e., the number of messages it can hold).	1000	Integer

CHAPTER 50. TELEGRAM

Both producer and consumer are supported

The Telegram component provides access to the [Telegram Bot API](#). It allows a Camel-based application to send and receive messages by acting as a Bot, participating in direct conversations with normal users, private and public groups or channels.

A Telegram Bot must be created before using this component, following the instructions at the [Telegram Bot developers](#) home. When a new Bot is created, the [BotFather](#) provides an **authorization token** corresponding to the Bot. The authorization token is a mandatory parameter for the camel-telegram endpoint.



NOTE

In order to allow the Bot to receive all messages exchanged within a group or channel (not just the ones starting with a '/' character), ask the BotFather to **disable the privacy mode**, using the `/setprivacy` command.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-telegram</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

50.1. URI FORMAT

```
telegram:type[?options]
```

50.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

50.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

50.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

50.3. COMPONENT OPTIONS

The Telegram component supports 7 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
baseUri (advanced)	Can be used to set an alternative base URI, e.g. when you want to test the component against a mock Telegram API.		String
client (advanced)	To use a custom AsyncHttpClient.		AsyncHttpClient
clientConfig (advanced)	To configure the AsyncHttpClient to use a custom com.ning.http.client.AsyncHttpClientConfig instance.		AsyncHttpClientConfig
authorizationToken (security)	The default Telegram authorization token to be used when the information is not provided in the endpoints.		String

50.4. ENDPOINT OPTIONS

The Telegram endpoint is configured using URI syntax:

```
telegram:type
```

with the following path and query parameters:

50.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
type (common)	Required The endpoint type. Currently, only the 'bots' type is supported. Enum values: <ul style="list-style-type: none">• bots		String

50.4.2. Query Parameters (30 parameters)

Name	Description	Default	Type
------	-------------	---------	------

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
limit (consumer)	Limit on the number of updates that can be received in a single polling request.	100	Integer
sendEmptyMessageWhenIdle (consumer)	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	boolean
timeout (consumer)	Timeout in seconds for long polling. Put 0 for short polling or a bigger number for long polling. Long polling produces shorter response time.	30	Integer
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> • <code>InOnly</code> • <code>InOut</code> • <code>InOptionalOut</code> 		<code>ExchangePattern</code>
pollStrategy (consumer (advanced))	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.		<code>PollingConsumerPollStrategy</code>

Name	Description	Default	Type
chatId (producer)	The identifier of the chat that will receive the produced messages. Chat ids can be first obtained from incoming messages (eg. when a telegram user starts a conversation with a bot, its client sends automatically a '/start' message containing the chat id). It is an optional parameter, as the chat id can be set dynamically for each outgoing message (using body or headers).		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
baseUri (advanced)	Can be used to set an alternative base URI, e.g. when you want to test the component against a mock Telegram API.		String
bufferSize (advanced)	The initial in-memory buffer size used when transferring data between Camel and AHC Client.	4096	int
clientConfig (advanced)	To configure the AsyncHttpClient to use a custom com.ning.http.client.AsyncHttpClientConfig instance.		AsyncHttpClientConfig
proxyHost (proxy)	HTTP proxy host which could be used when sending out the message.		String
proxyPort (proxy)	HTTP proxy port which could be used when sending out the message.		Integer
proxyType (proxy)	HTTP proxy type which could be used when sending out the message. Enum values: <ul style="list-style-type: none"> ● HTTP ● SOCKS4 ● SOCKS5 	HTTP	TelegramProxyType

Name	Description	Default	Type
backoffErrorThreshold (scheduler)	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.		int
backoffIdleThreshold (scheduler)	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.		int
backoffMultiplier (scheduler)	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.		int
delay (scheduler)	Milliseconds before the next poll.	500	long
greedy (scheduler)	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	boolean
initialDelay (scheduler)	Milliseconds before the first poll starts.	1000	long
repeatCount (scheduler)	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0	long
runLoggingLevel (scheduler)	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. Enum values: <ul style="list-style-type: none"> ● TRACE ● DEBUG ● INFO ● WARN ● ERROR ● OFF 	TRACE	LoggingLevel

Name	Description	Default	Type
scheduledExecutorService (scheduler)	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.		ScheduledExecutorService
scheduler (scheduler)	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler.	none	Object
schedulerProperties (scheduler)	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.		Map
startScheduler (scheduler)	Whether the scheduler should be auto started.	true	boolean
timeUnit (scheduler)	Time unit for initialDelay and delay options. Enum values: <ul style="list-style-type: none"> ● NANoseconds ● MICROseconds ● MILLIseconds ● SECONDS ● MINUTES ● HOURS ● DAYS 	MILLIS ECON DS	TimeUnit
useFixedDelay (scheduler)	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	boolean
authorizationToken (security)	Required The authorization token for using the bot (ask the BotFather).		String

50.4.3. Message Headers

Name	Description
CamelTelegramChatId	This header is used by the producer endpoint in order to resolve the chat id that will receive the message. The recipient chat id can be placed (in order of priority) in message body, in the CamelTelegramChatId header or in the endpoint configuration (chatId option). This header is also present in all incoming messages.

Name	Description
CamelTelegramMediaType	This header is used to identify the media type when the outgoing message is composed of pure binary data. Possible values are strings or enum values belonging to the org.apache.camel.component.telegram.TelegramMediaType enumeration.
CamelTelegramMediaTitleCaption	This header is used to provide a caption or title for outgoing binary messages.
CamelTelegramParseMode	This header is used to format text messages using HTML or Markdown (see org.apache.camel.component.telegram.TelegramParseMode).

50.5. USAGE

The Telegram component supports both consumer and producer endpoints. It can also be used in **reactive chat-bot mode** (to consume, then produce messages).

50.6. PRODUCER EXAMPLE

The following is a basic example of how to send a message to a Telegram chat through the Telegram Bot API.

in Java DSL

```
from("direct:start").to("telegram:bots?
authorizationToken=123456789:insertYourAuthorizationTokenHere");
```

or in Spring XML

```
<route>
  <from uri="direct:start"/>
  <to uri="telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere"/>
</route>
```

The code **123456789:insertYourAuthorizationTokenHere** is the **authorization token** corresponding to the Bot.

When using the producer endpoint without specifying the **chat id** option, the target chat will be identified using information contained in the body or headers of the message. The following message bodies are allowed for a producer endpoint (messages of type **OutgoingXXXMessage** belong to the package **org.apache.camel.component.telegram.model**)

Java Type	Description
OutgoingTextMessage	To send a text message to a chat
OutgoingPhotoMessage	To send a photo (JPG, PNG) to a chat

Java Type	Description
OutgoingAudioMessage	To send a mp3 audio to a chat
OutgoingVideoMessage	To send a mp4 video to a chat
OutgoingDocumentMessage	To send a file to a chat (any media type)
OutgoingStickerMessage	To send a sticker to a chat (WEBP)
OutgoingAnswerInlineQuery	To send answers to an inline query
EditMessageTextMessage	To edit text and game messages (editMessageText)
EditMessageCaptionMessage	To edit captions of messages (editMessageCaption)
EditMessageMediaMessage	To edit animation, audio, document, photo, or video messages. (editMessageMedia)
EditMessageReplyMarkupMessage	To edit only the reply markup of message. (editMessageReplyMarkup)
EditMessageDelete	To delete a message, including service messages. (deleteMessage)
SendLocationMessage	To send a location (setSendLocation)
EditMessageLiveLocationMessage	To send changes to a live location (editMessageLiveLocation)
StopMessageLiveLocationMessage	To stop updating a live location message sent by the bot or via the bot (for inline bots) before live_period expires (stopMessageLiveLocation)
SendVenueMessage	To send information about a venue (sendVenue)
byte[]	To send any media type supported. It requires the CamelTelegramMediaType header to be set to the appropriate media type
String	To send a text message to a chat. It gets converted automatically into a OutgoingTextMessage

50.7. CONSUMER EXAMPLE

The following is a basic example of how to receive all messages that telegram users are sending to the configured Bot. In Java DSL

```
from("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere")
.bean(ProcessorBean.class)
```

or in Spring XML

```
<route>
  <from uri="telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere"/>
  <bean ref="myBean" />
</route>

<bean id="myBean" class="com.example.MyBean"/>
```

The **MyBean** is a simple bean that will receive the messages

```
public class MyBean {

    public void process(String message) {
        // or Exchange, or org.apache.camel.component.telegram.model.IncomingMessage (or both)

        // do process
    }
}
```

Supported types for incoming messages are

Java Type	Description
IncomingMessage	The full object representation of an incoming message
String	The content of the message, for text messages only

50.8. REACTIVE CHAT-BOT EXAMPLE

The reactive chat-bot mode is a simple way of using the Camel component to build a simple chat bot that replies directly to chat messages received from the Telegram users.

The following is a basic configuration of the chat-bot in Java DSL

```
from("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere")
.bean(ChatBotLogic.class)
.to("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere");
```

or in Spring XML

```
<route>
  <from uri="telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere"/>
  <bean ref="chatBotLogic" />
  <to uri="telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere"/>
</route>

<bean id="chatBotLogic" class="com.example.ChatBotLogic"/>
```

The **ChatBotLogic** is a simple bean that implements a generic String-to-String method.

```
public class ChatBotLogic {

    public String chatBotProcess(String message) {
        if( "do-not-reply".equals(message) ) {
            return null; // no response in the chat
        }

        return "echo from the bot: " + message; // echoes the message
    }
}
```

Every non-null string returned by the **chatBotProcess** method is automatically routed to the chat that originated the request (as the **CamelTelegramChatId** header is used to route the message).

50.9. GETTING THE CHAT ID

If you want to push messages to a specific Telegram chat when an event occurs, you need to retrieve the corresponding chat ID. The chat ID is not currently shown in the telegram client, but you can obtain it using a simple route.

First, add the bot to the chat where you want to push messages, then run a route like the following one.

```
from("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere")
.to("log:INFO?showHeaders=true");
```

Any message received by the bot will be dumped to your log together with information about the chat (**CamelTelegramChatId** header).

Once you get the chat ID, you can use the following sample route to push message to it.

```
from("timer:tick")
.setBody().constant("Hello")
to("telegram:bots?
authorizationToken=123456789:insertYourAuthorizationTokenHere&chatId=123456")
```

Note that the corresponding URI parameter is simply **chatId**.

50.10. CUSTOMIZING KEYBOARD

You can customize the user keyboard instead of asking him to write an option. **OutgoingTextMessage** has the property **ReplyMarkup** which can be used for such thing.

```
from("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere")
.process(exchange -> {

    OutgoingTextMessage msg = new OutgoingTextMessage();
    msg.setText("Choose one option!");

    InlineKeyboardButton buttonOptionOne1 = InlineKeyboardButton.builder()
        .text("Option One - I").build();
```

```

InlineKeyboardButton buttonOptionOneI = InlineKeyboardButton.builder()
    .text("Option One - I").build();

InlineKeyboardButton buttonOptionTwoI = InlineKeyboardButton.builder()
    .text("Option Two - I").build();

ReplyKeyboardMarkup replyMarkup = ReplyKeyboardMarkup.builder()
    .keyboard()
    .addRow(Arrays.asList(buttonOptionOneI, buttonOptionOneI))
    .addRow(Arrays.asList(buttonOptionTwoI))
    .close()
    .oneTimeKeyboard(true)
    .build();

msg.setReplyMarkup(replyMarkup);

exchange.getIn().setBody(msg);
})
.to("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere");

```

If you want to disable it the next message must have the property **removeKeyboard** set on **ReplyKeyboardMarkup** object.

```

from("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere")
    .process(exchange -> {

        OutgoingTextMessage msg = new OutgoingTextMessage();
        msg.setText("Your answer was accepted!");

        ReplyKeyboardMarkup replyMarkup = ReplyKeyboardMarkup.builder()
            .removeKeyboard(true)
            .build();

        msg.setReplyKeyboardMarkup(replyMarkup);

        exchange.getIn().setBody(msg);
    })
    .to("telegram:bots?authorizationToken=123456789:insertYourAuthorizationTokenHere");

```

50.11. WEBHOOK MODE

The Telegram component supports usage in the **webhook mode** using the **camel-webhook** component.

In order to enable webhook mode, users need first to add a REST implementation to their application. Maven users, for example, can add **netty-http** to their **pom.xml** file:

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-netty-http</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>

```

Once done, you need to prepend the webhook URI to the telegram URI you want to use.

In Java DSL:

```
from("webhook:telegram:bots?
authorizationToken=123456789:insertYourAuthorizationTokenHere").to("log:info");
```

Some endpoints will be exposed by your application and Telegram will be configured to send messages to them. You need to ensure that your server is exposed to the internet and to pass the right value of the `camel.component.webhook.configuration.webhook-external-url` property.

Refer to the `camel-webhook` component documentation for instructions on how to set it.

50.12. SPRING BOOT AUTO-CONFIGURATION

When using telegram with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-telegram-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 8 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.telegram.authorization-token</code>	The default Telegram authorization token to be used when the information is not provided in the endpoints.		String
<code>camel.component.telegram.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.telegram.base-uri</code>	Can be used to set an alternative base URI, e.g. when you want to test the component against a mock Telegram API.		String

Name	Description	Default	Type
<code>camel.component.telegram.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.telegram.client</code>	To use a custom <code>AsyncHttpClient</code> . The option is a <code>org.apache.camel.async-http-client.AsyncHttpClient</code> type.		<code>AsyncHttpClient</code>
<code>camel.component.telegram.client-config</code>	To configure the <code>AsyncHttpClient</code> to use a custom <code>com.ning.http.client.AsyncHttpClientConfig</code> instance. The option is a <code>org.apache.camel.async-http-client.AsyncHttpClientConfig</code> type.		<code>AsyncHttpClientConfig</code>
<code>camel.component.telegram.enabled</code>	Whether to enable auto configuration of the telegram component. This is enabled by default.		Boolean
<code>camel.component.telegram.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean

CHAPTER 51. TIMER

Only consumer is supported

The Timer component is used to generate message exchanges when a timer fires. You can only consume events from this endpoint.

51.1. URI FORMAT


```
timer:name[?options]
```

Where **name** is the name of the **Timer** object, which is created and shared across endpoints. So if you use the same name for all your timer endpoints, only one **Timer** object and thread will be used.



NOTE

The IN body of the generated exchange is **null**. So `exchange.getIn().getBody()` returns **null**.



NOTE

Advanced Scheduler

See also the [Quartz](#) component that supports much more advanced scheduling.

51.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

51.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

51.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

51.3. COMPONENT OPTIONS

The Timer component supports 2 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

51.4. ENDPOINT OPTIONS

The Timer endpoint is configured using URI syntax:

```
timer:timerName
```

with the following path and query parameters:

51.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
timerName (consumer)	Required The name of the timer.		String

51.4.2. Query Parameters (13 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
delay (consumer)	Delay before first event is triggered.	1000	long
fixedRate (consumer)	Events take place at approximately regular intervals, separated by the specified period.	false	boolean
includeMetadata (consumer)	Whether to include metadata in the exchange such as fired time, timer name, timer count etc. This information is default included.	true	boolean
period (consumer)	If greater than 0, generate periodic events every period.	1000	long
repeatCount (consumer)	Specifies a maximum limit of number of fires. So if you set it to 1, the timer will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.		long
exceptionHandler (consumer (advanced))	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		<code>ExceptionHandler</code>
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● <code>InOnly</code> ● <code>InOut</code> ● <code>InOptionalOut</code> 		<code>ExchangePattern</code>
daemon (advanced)	Specifies whether or not the thread associated with the timer endpoint runs as a daemon. The default value is true.	true	boolean

Name	Description	Default	Type
pattern (advanced)	Allows you to specify a custom Date pattern to use for setting the time option using URI syntax.		String
synchronous (advanced)	Sets whether synchronous processing should be strictly used.	false	boolean
time (advanced)	A java.util.Date the first event should be generated. If using the URI, the pattern expected is: yyyy-MM-dd HH:mm:ss or yyyy-MM-dd'T'HH:mm:ss.		Date
timer (advanced)	To use a custom Timer.		Timer

51.5. EXCHANGE PROPERTIES

When the timer is fired, it adds the following information as properties to the **Exchange**:

Name	Type	Description
Exchange.TIMER_NAME	String	The value of the name option.
Exchange.TIMER_TIME	Date	The value of the time option.
Exchange.TIMER_PERIOD	long	The value of the period option.
Exchange.TIMER_FIRED_TIME	Date	The time when the consumer fired.
Exchange.TIMER_COUNTER	Long	The current fire counter. Starts from 1.

51.6. SAMPLE

To set up a route that generates an event every 60 seconds:

```
from("timer://foo?fixedRate=true&period=60000").to("bean:myBean?method=someMethodName");
```

The above route will generate an event and then invoke the **someMethodName** method on the bean called **myBean** in the Registry.

And the route in Spring DSL:

```
<route>
  <from uri="timer://foo?fixedRate=true&period=60000"/>
  <to uri="bean:myBean?method=someMethodName"/>
</route>
```

```
</route>
```

51.7. FIRING AS SOON AS POSSIBLE

Since Camel 2.17

You may want to fire messages in a Camel route as soon as possible you can use a negative delay:

```
<route>
  <from uri="timer://foo?delay=-1"/>
  <to uri="bean:myBean?method=someMethodName"/>
</route>
```

In this way the timer will fire messages immediately.

You can also specify a repeatCount parameter in conjunction with a negative delay to stop firing messages after a fixed number has been reached.

If you don't specify a repeatCount then the timer will continue firing messages until the route will be stopped.

51.8. FIRING ONLY ONCE

You may want to fire a message in a Camel route only once, such as when starting the route. To do that you use the repeatCount option as shown:

```
<route>
  <from uri="timer://foo?repeatCount=1"/>
  <to uri="bean:myBean?method=someMethodName"/>
</route>
```

51.9. SPRING BOOT AUTO-CONFIGURATION

When using timer with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-timer-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
		t	

Name	Description	Default	Type
camel.component.timer.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.timer.bridge-error-handler	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
camel.component.timer.enabled	Whether to enable auto configuration of the timer component. This is enabled by default.		Boolean

CHAPTER 52. VALIDATOR

Only producer is supported

The Validation component performs XML validation of the message body using the JAXP Validation API and based on any of the supported XML schema languages, which defaults to [XML Schema](#)

Note that the component also supports the following useful schema languages:

- [RelaxNG Compact Syntax](#)
- [RelaxNG XML Syntax](#)

The [MSV](#) component also supports [RelaxNG XML Syntax](#).

52.1. URI FORMAT

```
validator:someLocalOrRemoteResource
```

Where **someLocalOrRemoteResource** is some URL to a local resource on the classpath or a full URL to a remote resource or resource on the file system which contains the XSD to validate against. For example:

- **msv:org/foo/bar.xsd**
- **msv:file:../foo/bar.xsd**
- **msv:http://acme.com/cheese.xsd**
- **validator:com/mypackage/myschema.xsd**

The Validation component is provided directly in the camel-core.

52.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

52.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

52.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

52.3. COMPONENT OPTIONS

The Validator component supports 3 options, which are listed below.

Name	Description	Default	Type
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
resourceResolverFactory (advanced)	To use a custom LSResourceResolver which depends on a dynamic endpoint resource URI.		ValidatorResourceResolverFactory

52.4. ENDPOINT OPTIONS

The Validator endpoint is configured using URI syntax:

```
validator:resourceUri
```

with the following path and query parameters:

52.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
resourceUri (producer)	Required URL to a local resource on the classpath, or a reference to lookup a bean in the Registry, or a full URL to a remote resource or resource on the file system which contains the XSD to validate against.		String

52.4.2. Query Parameters (10 parameters)

Name	Description	Default	Type
failOnNullBody (producer)	Whether to fail if no body exists.	true	boolean
failOnNullHeader (producer)	Whether to fail if no header exists when validating against a header.	true	boolean
headerName (producer)	To validate against a header instead of the message body.		String
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
errorHandler (advanced)	To use a custom <code>org.apache.camel.processor.validation.ValidatorErrorHandler</code> . The default error handler captures the errors and throws an exception.		ValidatorErrorHandler
resourceResolver (advanced)	To use a custom <code>LSResourceResolver</code> . Do not use together with <code>resourceResolverFactory</code> .		LSResourceResolver

Name	Description	Default	Type
resourceResolverFactory (advanced)	To use a custom LSResourceResolver which depends on a dynamic endpoint resource URI. The default resource resolver factory returns a resource resolver which can read files from the class path and file system. Do not use together with resourceResolver.		ValidatorResourceResolverFactory
schemaFactory (advanced)	To use a custom javax.xml.validation.SchemaFactory.		SchemaFactory
schemaLanguage (advanced)	Configures the W3C XML Schema Namespace URI.	http://www.w3.org/2001/XMLSchema	String
useSharedSchema (advanced)	Whether the Schema instance should be shared or not. This option is introduced to work around a JDK 1.6.x bug. Xerces should not have this issue.	true	boolean

52.5. EXAMPLE

The following [example](#) shows how to configure a route from endpoint **direct:start** which then goes to one of two endpoints, either **mock:valid** or **mock:invalid** based on whether or not the XML matches the given schema (which is supplied on the classpath).

52.6. ADVANCED: JMX METHOD CLEARCACHEDSCHEMA

You can force that the cached schema in the validator endpoint is cleared and reread with the next process call with the JMX operation **clearCachedSchema**. You can also use this method to programmatically clear the cache. This method is available on the **ValidatorEndpoint** class.

52.7. SPRING BOOT AUTO-CONFIGURATION

When using validator with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-validator-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
camel.component.validator.autowired-enabled	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
camel.component.validator.enabled	Whether to enable auto configuration of the validator component. This is enabled by default.		Boolean
camel.component.validator.lazy-start-producer	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
camel.component.validator.resource-resolver-factory	To use a custom LSResourceResolver which depends on a dynamic endpoint resource URI. The option is a <code>org.apache.camel.component.validator.ValidatorResourceResolverFactory</code> type.		ValidatorResourceResolverFactory

CHAPTER 53. WEBHOOK

Only consumer is supported

The Webhook meta component allows other Camel components to configure webhooks on a remote webhook provider and listening for them.

The following components currently provide webhook endpoints:

- **Telegram**

Maven users can add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-webhook</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

Typically, other components that support webhook will bring this dependency transitively.

53.1. URI FORMAT

```
webhook:endpoint[?options]
```

53.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

53.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

53.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

53.3. COMPONENT OPTIONS

The Webhook component supports 8 options, which are listed below.

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
webhookAutoRegister (consumer)	Automatically register the webhook at startup and unregister it on shutdown.	true	boolean
webhookBasePath (consumer)	The first (base) path element where the webhook will be exposed. It's a good practice to set it to a random string, so that it cannot be guessed by unauthorized parties.		String
webhookComponentName (consumer)	The Camel Rest component to use for the REST transport, such as <code>netty-http</code> .		String
webhookExternalUrl (consumer)	The URL of the current service as seen by the webhook provider.		String
webhookPath (consumer)	The path where the webhook endpoint will be exposed (relative to <code>basePath</code> , if any).		String
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean

Name	Description	Default	Type
configuration (advanced)	Set the default configuration for the webhook meta-component.		WebhookConfiguration

53.4. ENDPOINT OPTIONS

The Webhook endpoint is configured using URI syntax:

```
webhook:endpointUri
```

with the following path and query parameters:

53.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
endpointUri (consumer)	Required The delegate uri. Must belong to a component that supports webhooks.		String

53.4.2. Query Parameters (8 parameters)

Name	Description	Default	Type
bridgeErrorHandler (consumer)	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	boolean
webhookAutoRegister (consumer)	Automatically register the webhook at startup and unregister it on shutdown.	true	boolean
webhookBasePath (consumer)	The first (base) path element where the webhook will be exposed. It's a good practice to set it to a random string, so that it cannot be guessed by unauthorized parties.		String

Name	Description	Default	Type
webhookComponentName (consumer)	The Camel Rest component to use for the REST transport, such as netty-http.		String
webhookExternalUrl (consumer)	The URL of the current service as seen by the webhook provider.		String
webhookPath (consumer)	The path where the webhook endpoint will be exposed (relative to basePath, if any).		String
exceptionHandler (consumer (advanced))	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.		ExceptionHandler
exchangePattern (consumer (advanced))	Sets the exchange pattern when the consumer creates an exchange. Enum values: <ul style="list-style-type: none"> ● InOnly ● InOut ● InOptionalOut 		ExchangePattern

53.5. EXAMPLES

Examples of webhook component are provided in the documentation of the delegate components that support it.

53.6. SPRING BOOT AUTO-CONFIGURATION

When using webhook with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-webhook-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 9 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.webhook.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.webhook.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.webhook.configuration</code>	Set the default configuration for the webhook meta-component. The option is a <code>org.apache.camel.component.webhook.WebhookConfiguration</code> type.		WebhookConfiguration
<code>camel.component.webhook.enabled</code>	Whether to enable auto configuration of the webhook component. This is enabled by default.		Boolean
<code>camel.component.webhook.webhook-auto-register</code>	Automatically register the webhook at startup and unregister it on shutdown.	true	Boolean
<code>camel.component.webhook.webhook-base-path</code>	The first (base) path element where the webhook will be exposed. It's a good practice to set it to a random string, so that it cannot be guessed by unauthorized parties.		String
<code>camel.component.webhook.webhook-component-name</code>	The Camel Rest component to use for the REST transport, such as <code>netty-http</code> .		String
<code>camel.component.webhook.webhook-external-url</code>	The URL of the current service as seen by the webhook provider.		String
<code>camel.component.webhook.webhook-path</code>	The path where the webhook endpoint will be exposed (relative to <code>basePath</code> , if any).		String

CHAPTER 54. XSLT

Only producer is supported

The XSLT component allows you to process a message using an [XSLT](#) template. This can be ideal when using Templating to generate response for requests.

54.1. URI FORMAT

```
xslt:templateName[?options]
```

The URI format contains **templateName**, which can be one of the following:

- the classpath-local URI of the template to invoke
- the complete URL of the remote template.

You can append query options to the URI in the following format:

```
?option=value&option=value&...
```

Table 54.1. Table 1. Example URIs

URI	Description
xslt:com/acme/mytransform.xml	Refers to the file com/acme/mytransform.xml on the classpath
xslt:file:///foo/bar.xml	Refers to the file /foo/bar.xml
xslt:http://acme.com/cheese/foo.xml	Refers to the remote http resource

54.2. CONFIGURING OPTIONS

Camel components are configured on two separate levels:

- component level
- endpoint level

54.2.1. Configuring Component Options

The component level is the highest level which holds general and common configurations that are inherited by the endpoints. For example a component may have security settings, credentials for authentication, urls for network connection and so forth.

Some components only have a few options, and others may have many. Because components typically have pre configured defaults that are commonly used, then you may often only need to configure a few options on a component; or none at all.

Configuring components can be done with the [Component DSL](#), in a configuration file (application.properties|yaml), or directly with Java code.

54.2.2. Configuring Endpoint Options

Where you find yourself configuring the most is on endpoints, as endpoints often have many options, which allows you to configure what you need the endpoint to do. The options are also categorized into whether the endpoint is used as consumer (from) or as a producer (to), or used for both.

Configuring endpoints is most often done directly in the endpoint URI as path and query parameters. You can also use the [Endpoint DSL](#) as a type safe way of configuring endpoints.

A good practice when configuring options is to use [Property Placeholders](#), which allows to not hardcode urls, port numbers, sensitive information, and other settings. In other words placeholders allows to externalize the configuration from your code, and gives more flexibility and reuse.

The following two sections lists all the options, firstly for the component followed by the endpoint.

54.3. COMPONENT OPTIONS

The XSLT component supports 7 options, which are listed below.

Name	Description	Default	Type
contentCache (producer)	Cache for the resource content (the stylesheet file) when it is loaded. If set to false Camel will reload the stylesheet file on each message processing. This is good for development. A cached stylesheet can be forced to reload at runtime via JMX using the <code>clearCachedStylesheet</code> operation.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
autowiredEnabled (advanced)	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	boolean
transformerFactoryClass (advanced)	To use a custom XSLT transformer factory, specified as a FQN class name.		String

Name	Description	Default	Type
transformerFactoryConfigurationStrategy (advanced)	A configuration strategy to apply on freshly created instances of TransformerFactory.		TransformerFactoryConfigurationStrategy
uriResolver (advanced)	To use a custom UriResolver. Should not be used together with the option 'uriResolverFactory'.		UriResolver
uriResolverFactory (advanced)	To use a custom UriResolver which depends on a dynamic endpoint resource URI. Should not be used together with the option 'uriResolver'.		XsltUriResolverFactory

54.4. ENDPOINT OPTIONS

The XSLT endpoint is configured using URI syntax:

```
xslt:resourceUri
```

with the following path and query parameters:

54.4.1. Path Parameters (1 parameters)

Name	Description	Default	Type
resourceUri (producer)	Required Path to the template. The following is supported by the default UriResolver. You can prefix with: classpath, file, http, ref, or bean. classpath, file and http loads the resource using these protocols (classpath is default). ref will lookup the resource in the registry. bean will call a method on a bean to be used as the resource. For bean you can specify the method name after dot, eg bean:myBean.myMethod.		String

54.4.2. Query Parameters (13 parameters)

Name	Description	Default	Type
contentCache (producer)	Cache for the resource content (the stylesheet file) when it is loaded. If set to false Camel will reload the stylesheet file on each message processing. This is good for development. A cached stylesheet can be forced to reload at runtime via JMX using the clearCachedStylesheet operation.	true	boolean

Name	Description	Default	Type
deleteOutputFile (producer)	If you have output=file then this option dictates whether or not the output file should be deleted when the Exchange is done processing. For example suppose the output file is a temporary file, then it can be a good idea to delete it after use.	false	boolean
failOnNullBody (producer)	Whether or not to throw an exception if the input body is null.	true	boolean
lazyStartProducer (producer)	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	boolean
output (producer)	Option to specify which output type to use. Possible values are: string, bytes, DOM, file. The first three options are all in memory based, where as file is streamed directly to a java.io.File. For file you must specify the filename in the IN header with the key Exchange.XSLT_FILE_NAME which is also CamelXsltFileName. Also any paths leading to the filename must be created beforehand, otherwise an exception is thrown at runtime. Enum values: <ul style="list-style-type: none"> ● string ● bytes ● DOM ● file 	string	XsltOutput
transformerCacheSize (producer)	The number of javax.xml.transform.Transformer object that are cached for reuse to avoid calls to Template.newTransformer().	0	int
entityResolver (advanced)	To use a custom org.xml.sax.EntityResolver with javax.xml.transform.sax.SAXSource.		EntityResolver

Name	Description	Default	Type
errorListener (advanced)	Allows to configure to use a custom <code>javax.xml.transform.ErrorListener</code> . Beware when doing this then the default error listener which captures any errors or fatal errors and store information on the Exchange as properties is not in use. So only use this option for special use-cases.		ErrorListener
resultHandlerFactory (advanced)	Allows you to use a custom <code>org.apache.camel.builder.xml.ResultHandlerFactory</code> which is capable of using custom <code>org.apache.camel.builder.xml.ResultHandler</code> types.		ResultHandlerFactory
transformerFactory (advanced)	To use a custom XSLT transformer factory.		TransformerFactory
transformerFactoryClass (advanced)	To use a custom XSLT transformer factory, specified as a FQN class name.		String
transformerFactoryConfigurationStrategy (advanced)	A configuration strategy to apply on freshly created instances of TransformerFactory.		TransformerFactoryConfigurationStrategy
uriResolver (advanced)	To use a custom <code>javax.xml.transform.URIResolver</code> .		URIResolver

54.5. USING XSLT ENDPOINTS

The following format is an example of using an XSLT template to formulate a response for a message for InOut message exchanges (where there is a **JMSReplyTo** header)

```
from("activemq:My.Queue").
to("xslt:com/acme/mytransform.xml");
```

If you want to use InOnly and consume the message and send it to another destination you could use the following route:

```
from("activemq:My.Queue").
to("xslt:com/acme/mytransform.xml").
to("activemq:Another.Queue");
```

54.6. GETTING USEABLE PARAMETERS INTO THE XSLT

By default, all headers are added as parameters which are then available in the XSLT. To make the parameters useable, you will need to declare them.

```
<setHeader name="myParam"><constant>42</constant></setHeader>
<to uri="xslt:MyTransform.xsl"/>
```

The parameter also needs to be declared in the top level of the XSLT for it to be available:

```
<xsl: ..... >

  <xsl:param name="myParam"/>

  <xsl:template ...>
```

54.7. SPRING XML VERSIONS

To use the above examples in Spring XML you would use something like the following code:

```
<camelContext xmlns="http://activemq.apache.org/camel/schema/spring">
  <route>
    <from uri="activemq:My.Queue"/>
    <to uri="xslt:org/apache/camel/spring/processor/example.xsl"/>
    <to uri="activemq:Another.Queue"/>
  </route>
</camelContext>
```

54.8. USING XSL:INCLUDE

Camel provides its own implementation of **URIResolver**. This allows Camel to load included files from the classpath.

For example the include file in the following code will be located relative to the starting endpoint.

```
<xsl:include href="staff_template.xsl"/>
```

This means that Camel will locate the file in the **classpath** as **org/apache/camel/component/xslt/staff_template.xsl**

You can use **classpath:** or **file:** to instruct Camel to look either in the classpath or file system. If you omit the prefix then Camel uses the prefix from the endpoint configuration. If no prefix is specified in the endpoint configuration, the default is **classpath:**.

You can also refer backwards in the include paths. In the following example, the xsl file will be resolved under **org/apache/camel/component**.

```
<xsl:include href="../../staff_other_template.xsl"/>
```

54.9. USING XSL:INCLUDE AND DEFAULT PREFIX

Camel will use the prefix from the endpoint configuration as the default prefix.

You can explicitly specify **file:** or **classpath:** loading. The two loading types can be mixed in a XSLT script, if necessary.

54.10. DYNAMIC STYLESHEETS

To provide a dynamic stylesheet at runtime you can define a dynamic URI. See [How to use a dynamic URI in to\(\)](#) for more information.

54.11. ACCESSING WARNINGS, ERRORS AND FATALERRORS FROM XSLT ERRORLISTENER

Any warning/error or fatalError is stored on the current Exchange as a property with the keys **Exchange.XSLT_ERROR**, **Exchange.XSLT_FATAL_ERROR**, or **Exchange.XSLT_WARNING** which allows end users to get hold of any errors happening during transformation.

For example in the stylesheet below, we want to terminate if a staff has an empty dob field. And to include a custom error message using `xsl:message`.

```
<xsl:template match="/">
  <html>
    <body>
      <xsl:for-each select="staff/programmer">
        <p>Name: <xsl:value-of select="name"/><br />
          <xsl:if test="dob="">
            <xsl:message terminate="yes">Error: DOB is an empty string!</xsl:message>
          </xsl:if>
        </p>
      </xsl:for-each>
    </body>
  </html>
</xsl:template>
```

The exception is stored on the Exchange as a warning with the key **Exchange.XSLT_WARNING**.

54.12. SPRING BOOT AUTO-CONFIGURATION

When using xslt with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-xslt-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

The component supports 8 options, which are listed below.

Name	Description	Default	Type
		t	

Name	Description	Default	Type
<code>camel.component.xslt.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.xslt.content-cache</code>	Cache for the resource content (the stylesheet file) when it is loaded. If set to false Camel will reload the stylesheet file on each message processing. This is good for development. A cached stylesheet can be forced to reload at runtime via JMX using the <code>clearCachedStylesheet</code> operation.	true	Boolean
<code>camel.component.xslt.enabled</code>	Whether to enable auto configuration of the xslt component. This is enabled by default.		Boolean
<code>camel.component.xslt.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.xslt.transformer-factory-class</code>	To use a custom XSLT transformer factory, specified as a FQN class name.		String
<code>camel.component.xslt.transformer-factory-configuration-strategy</code>	A configuration strategy to apply on freshly created instances of TransformerFactory. The option is a <code>org.apache.camel.component.xslt.TransformerFactoryConfigurationStrategy</code> type.		TransformerFactoryConfigurationStrategy
<code>camel.component.xslt.uri-resolver</code>	To use a custom UriResolver. Should not be used together with the option 'uriResolverFactory'. The option is a <code>javax.xml.transform.URIResolver</code> type.		URIResolver

Name	Description	Default	Type
camel.component.xslt.uri-resolver-factory	To use a custom UriResolver which depends on a dynamic endpoint resource URI. Should not be used together with the option 'uriResolver'. The option is a org.apache.camel.component.xslt.XsltUriResolverFactory type.		XsltUriResolverFactory

CHAPTER 55. AVRO

This component provides a dataformat for avro, which allows serialization and deserialization of messages using Apache Avro's binary dataformat. Since Camel 3.2 rpc functionality was moved into separate **camel-avro-rpc** component.

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-avro</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

You can easily generate classes from a schema, using maven, ant etc. More details can be found at the [Apache Avro documentation](#).

55.1. AVRO DATAFORMAT OPTIONS

The Avro dataformat supports 1 options, which are listed below.

Name	Default	Java Type	Description
instanceClassName		String	Class name to use for marshal and unmarshalling.

55.2. AVRO DATA FORMAT USAGE

Using the avro data format is as easy as specifying that the class that you want to marshal or unmarshal in your route.

```
AvroDataFormat format = new AvroDataFormat(Value.SCHEMA$);

from("direct:in").marshal(format).to("direct:marshal");
from("direct:back").unmarshal(format).to("direct:unmarshal");
```

Where Value is an Avro Maven Plugin Generated class.

or in XML

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:in"/>
    <marshal>
      <avro instanceClass="org.apache.camel.dataformat.avro.Message"/>
    </marshal>
    <to uri="log:out"/>
  </route>
</camelContext>
```


An alternative can be to specify the dataformat inside the context and reference it from your route.

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <dataFormats>
    <avro id="avro" instanceClass="org.apache.camel.dataformat.avro.Message"/>
  </dataFormats>
  <route>
    <from uri="direct:in"/>
    <marshal><custom ref="avro"/></marshal>
    <to uri="log:out"/>
  </route>
</camelContext>
```

In the same manner you can unmarshal using the avro data format.

55.3. SPRING BOOT AUTO-CONFIGURATION

When using avro with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-avro-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 2 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.avro.enabled</code>	Whether to enable auto configuration of the avro data format. This is enabled by default.		Boolean
<code>camel.dataformat.avro.instance-class-name</code>	Class name to use for marshal and unmarshalling.		String

CHAPTER 56. AVRO JACKSON

Jackson Avro is a Data Format which uses the [Jackson library](#) with the [Avro extension](#) to unmarshal an Avro payload into Java objects or to marshal Java objects into an Avro payload.



NOTE

If you are familiar with Jackson, this Avro data format behaves in the same way as its JSON counterpart, and thus can be used with classes annotated for JSON serialization/deserialization.

```
from("kafka:topic").
  unmarshal().avro(AvroLibrary.Jackson, JsonNode.class).
  to("log:info");
```

56.1. CONFIGURING THE SCHEMARESOLVER

Since Avro serialization is schema-based, this data format requires that you provide a SchemaResolver object that is able to lookup the schema for each exchange that is going to be marshalled/unmarshalled.

You can add a single SchemaResolver to the registry and it will be looked up automatically. Or you can explicitly specify the reference to a custom SchemaResolver.

56.2. AVRO JACKSON OPTIONS

The Avro Jackson dataformat supports 18 options, which are listed below.

Name	Default	Java Type	Description
<code>objectMapper</code>		String	Lookup and use the existing ObjectMapper with the given id when using Jackson.
<code>useDefaultObjectMapper</code>		Boolean	Whether to lookup and use default Jackson ObjectMapper from the registry.
<code>unmarshalType</code>		String	Class name of the java type to use when unmarshalling.
<code>jsonView</code>		String	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has JsonView annotations.
<code>include</code>		String	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to NON_NULL.
<code>allowJmsType</code>		Boolean	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.

Name	Default	Java Type	Description
collectionType		String	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.
useList		Boolean	To unmarshal to a List of Map or a List of Pojo.
moduleClassNames		String	To use custom Jackson modules com.fasterxml.jackson.databind.Module specified as a String with FQN class names. Multiple classes can be separated by comma.
moduleRefs		String	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.
enableFeatures		String	Set of features to enable on the Jackson com.fasterxml.jackson.databind.ObjectMapper. The features should be a name that matches a enum from com.fasterxml.jackson.databind.SerializationFeature, com.fasterxml.jackson.databind.DeserializationFeature, or com.fasterxml.jackson.databind.MapperFeature Multiple features can be separated by comma.
disableFeatures		String	Set of features to disable on the Jackson com.fasterxml.jackson.databind.ObjectMapper. The features should be a name that matches a enum from com.fasterxml.jackson.databind.SerializationFeature, com.fasterxml.jackson.databind.DeserializationFeature, or com.fasterxml.jackson.databind.MapperFeature Multiple features can be separated by comma.
allowUnmarshalType		Boolean	If enabled then Jackson is allowed to attempt to use the CamelJacksonUnmarshalType header during the unmarshalling. This should only be enabled when desired to be used.
timezone		String	If set then Jackson will use the Timezone when marshalling/unmarshalling.
autoDiscoverObjectMapper		Boolean	If set to true then Jackson will lookup for an objectMapper into the registry.
contentTypeHeader		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.

Name	Default	Java Type	Description
schemaResolver		String	Optional schema resolver used to lookup schemas for the data in transit.
autoDiscoverSchemaResolver		Boolean	When not disabled, the SchemaResolver will be looked up into the registry.

56.3. USING CUSTOM AVROMAPPER

You can configure **JacksonAvroDataFormat** to use a custom **AvroMapper** in case you need more control of the mapping configuration.

If you setup a single **AvroMapper** in the registry, then Camel will automatic lookup and use this **AvroMapper**.

56.4. DEPENDENCIES

To use Avro Jackson in your camel routes you need to add the dependency on **camel-jackson-avro** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jackson-avro</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

56.5. SPRING BOOT AUTO-CONFIGURATION

When using avro-jackson with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jackson-avro-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 19 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.avro-jackson.allow-jms-type</code>	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.	false	Boolean
<code>camel.dataformat.avro-jackson.allow-unmarshall-type</code>	If enabled then Jackson is allowed to attempt to use the CamelJacksonUnmarshalType header during the unmarshalling. This should only be enabled when desired to be used.	false	Boolean
<code>camel.dataformat.avro-jackson.auto-discover-object-mapper</code>	If set to true then Jackson will lookup for an objectMapper into the registry.	false	Boolean
<code>camel.dataformat.avro-jackson.auto-discover-schema-resolver</code>	When not disabled, the SchemaResolver will be looked up into the registry.	true	Boolean
<code>camel.dataformat.avro-jackson.collection-type</code>	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.		String
<code>camel.dataformat.avro-jackson.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.avro-jackson.disable-features</code>	Set of features to disable on the Jackson com.fasterxml.jackson.databind.ObjectMapper. The features should be a name that matches a enum from com.fasterxml.jackson.databind.SerializationFeature, com.fasterxml.jackson.databind.DeserializationFeature, or com.fasterxml.jackson.databind.MapperFeature. Multiple features can be separated by comma.		String

Name	Description	Default	Type
<code>camel.dataformat.avro-jackson.enable-features</code>	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.avro-jackson.enabled</code>	Whether to enable auto configuration of the avro-jackson data format. This is enabled by default.		Boolean
<code>camel.dataformat.avro-jackson.include</code>	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .		String
<code>camel.dataformat.avro-jackson.json-view</code>	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.		String
<code>camel.dataformat.avro-jackson.module-class-names</code>	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.		String
<code>camel.dataformat.avro-jackson.module-refs</code>	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.		String
<code>camel.dataformat.avro-jackson.object-mapper</code>	Lookup and use the existing <code>ObjectMapper</code> with the given id when using Jackson.		String
<code>camel.dataformat.avro-jackson.schema-resolver</code>	Optional schema resolver used to lookup schemas for the data in transit.		String
<code>camel.dataformat.avro-jackson.timezone</code>	If set then Jackson will use the <code>Timezone</code> when marshalling/unmarshalling.		String

Name	Description	Default	Type
<code>camel.dataformat.avro-jackson.unmarshal-type</code>	Class name of the java type to use when unmarshalling.		String
<code>camel.dataformat.avro-jackson.use-default-object-mapper</code>	Whether to lookup and use default Jackson ObjectMapper from the registry.	true	Boolean
<code>camel.dataformat.avro-jackson.use-list</code>	To unmarshal to a List of Map or a List of Pojo.	false	Boolean

CHAPTER 57. BINDY

The goal of this component is to allow the parsing/binding of non-structured data (or to be more precise non-XML data) to/from Java Beans that have binding mappings defined with annotations. Using Bindy, you can bind data from sources such as :

- CSV records,
- Fixed-length records,
- FIX messages,
- or almost any other non-structured data

to one or many Plain Old Java Object (POJO). Bindy converts the data according to the type of the java property. POJOs can be linked together with one-to-many relationships available in some cases. Moreover, for data type like Date, Double, Float, Integer, Short, Long and BigDecimal, you can provide the pattern to apply during the formatting of the property.

For the BigDecimal numbers, you can also define the precision and the decimal or grouping separators.

Type	Format Type	Pattern example	Link
Date	DateFormat	dd-MM-yyyy	https://docs.oracle.com/en/java/javase/11/docs/api/java.base/java/text/SimpleDateFormat.html
Decimal*	DecimalFormat	..##	https://docs.oracle.com/en/java/javase/11/docs/api/java.base/java/text/DecimalFormat.html

Where Decimal = Double, Integer, Float, Short, Long

Format supported

This first release only support comma separated values fields and key value pair fields (e.g. : FIX messages).

To work with camel-bindy, you must first define your model in a package (e.g. com.acme.model) and for each model class (e.g. Order, Client, Instrument, ...) add the required annotations (described hereafter) to the Class or field.

Multiple models

As you configure bindy using class names instead of package names you can put multiple models in the same package.

57.1. OPTIONS

The Bindy dataformat supports 5 options, which are listed below.

Name	Default	Java Type	Description
<code>type</code>		Enum	<p>Required Whether to use Csv, Fixed, or KeyValue.</p> <p>Enum values:</p> <ul style="list-style-type: none"> • Csv • Fixed • KeyValue
<code>classType</code>		String	Name of model class to use.
<code>locale</code>		String	To configure a default locale to use, such as us for united states. To use the JVM platform default locale then use the name default.
<code>unwrapSingleInstance</code>		Boolean	When unmarshalling should a single instance be unwrapped and returned instead of wrapped in a java.util.List.
<code>allowEmptyStream</code>		Boolean	Whether to allow empty streams in the unmarshal process. If true, no exception will be thrown when a body without records is provided.

57.2. ANNOTATIONS

The annotations created allow to map different concept of your model to the POJO like:

- Type of record (CSV, key value pair (e.g. FIX message), fixed length ...),
- Link (to link object in another object),
- DataField and their properties (int, type, ...),
- KeyValuePairField (for key = value format like we have in FIX financial messages),
- Section (to identify header, body and footer section),
- OneToMany,
- BindyConverter,
- FormatFactories

This section will describe them.

57.2.1.1. CsvRecord

The CsvRecord annotation is used to identify the root class of the model. It represents a record = "a line of a CSV file" and can be linked to several children model classes.

Annotation name	Record type	Level
CsvRecord	CSV	Class

Parameter name	Type	Required	Default value	Info
separator	String	✓		Separator used to split a record in tokens (mandatory) - can be ',' or ';' or 'anything'. The only whitespace character supported is tab (\t). No other whitespace characters (spaces) are not supported. This value is interpreted as a regular expression. If you want to use a sign which has a special meaning in regular expressions, e.g. the ' ' sign, then you have to mask it, like ' '.
allowEmptyStream	boolean		false	The allowEmptyStream parameter will allow to process the unavailable stream for CSV file.
autospanLine	boolean		false	Last record spans rest of line (optional) - if enabled then the last column is auto spanned to end of line, for example if it's a comment, etc this allows the line to contain all characters, also the delimiter char.
crlf	String		WINDOWS	Character to be used to add a carriage return after each record (optional) - allow to define the carriage return character to use. If you specify a value other than the three listed before, the value you enter (custom) will be used as the CRLF character(s). Three values can be used : WINDOWS, UNIX, MAC, or custom.
endWithLineBreak	boolean		true	The endWithLineBreak parameter flags if the CSV file should end with a line break or not (optional)
generateHeaderColumns	boolean		false	The generateHeaderColumns parameter allow to add in the CSV generated the header containing names of the columns
isOrdered	boolean		false	Indicates if the message must be ordered in output
name	String			Name describing the record (optional)

Parameter name	Type	Required	Default value	Info
quote	String		"	Whether to marshal columns with the given quote character (optional) - allow to specify a quote character of the fields when CSV is generated. This annotation is associated to the root class of the model and must be declared one time.
quoting	boolean		false	Indicate if the values (and headers) must be quoted when marshaling (optional)
quotingEscaped	boolean		false	Indicate if the values must be escaped when quoting (optional)
removeQuotes	boolean		true	The remove quotes parameter flags if unmarshalling should try to remove quotes for each field
skipField	boolean		false	The skipField parameter will allow to skip fields of a CSV file. If some fields are not necessary, they can be skipped.
skipFirstLine	boolean		false	The skipFirstLine parameter will allow to skip or not the first line of a CSV file. This line often contains columns definition

case 1 : separator = ','

The separator used to segregate the fields in the CSV record is , :

```
10, J, Pauline, M, XD12345678, Fortis Dynamic 15/15, 2500, USD, 08-01-2009
```

```
@CsvRecord( separator = "," )
public Class Order {
}

```

case 2 : separator = ';' ;

Compare to the previous case, the separator here is ; instead of , :

```
10; J; Pauline; M; XD12345678; Fortis Dynamic 15/15; 2500; USD; 08-01-2009
```

```
@CsvRecord( separator = ";" )
public Class Order {
}

```

case 3 : separator = '|'

Compare to the previous case, the separator here is | instead of ; :

```
10| J| Pauline| M| XD12345678| Fortis Dynamic 15/15| 2500| USD| 08-01-2009
```

```
@CsvRecord( separator = "\\|" )
public Class Order {
}

```

case 4 : separator = "\",\""

Applies for Camel 2.8.2 or older

When the field to be parsed of the CSV record contains , or ; which is also used as separator, we should find another strategy to tell camel bindy how to handle this case. To define the field containing the data with a comma, you will use single or double quotes as delimiter (e.g : '10', 'Street 10, NY', 'USA' or "10", "Street 10, NY", "USA").

—	In this case, the first and last character of the line which are a single or double quotes will be removed by bindy.
---	--

```
"10","J","Pauline"," M","XD12345678","Fortis Dynamic 15,15","2500","USD","08-01-2009"
```

```
@CsvRecord( separator = "\",\"" )
public Class Order {
}

```

Bindy automatically detects if the record is enclosed with either single or double quotes and automatic remove those quotes when unmarshalling from CSV to Object. Therefore do **not** include the quotes in the separator, but simply do as below:

```
"10","J","Pauline"," M","XD12345678","Fortis Dynamic 15,15","2500","USD","08-01-2009"
```

```
@CsvRecord( separator = "," )
public Class Order {
}

```

Notice that if you want to marshal from Object to CSV and use quotes, then you need to specify which quote character to use, using the **quote** attribute on the **@CsvRecord** as shown below:

```
@CsvRecord( separator = ",", quote = "\"" )
public Class Order {
}

```

case 5 : separator & skipFirstLine

The feature is interesting when the client wants to have in the first line of the file, the name of the data fields :

```
order id, client id, first name, last name, isin code, instrument name, quantity, currency, date
```

To inform bindy that this first line must be skipped during the parsing process, then we use the attribute :

```
@CsvRecord(separator = ",", skipFirstLine = true)
public Class Order {
}
}
```

case 6 : generateHeaderColumns

To add at the first line of the CSV generated, the attribute generateHeaderColumns must be set to true in the annotation like this :

```
@CsvRecord( generateHeaderColumns = true )
public Class Order {
}
}
```

As a result, Bindy during the unmarshaling process will generate CSV like this :

```
order id, client id, first name, last name, isin code, instrument name, quantity, currency, date
10, J, Pauline, M, XD12345678, Fortis Dynamic 15/15, 2500, USD, 08-01-2009
```

case 7 : carriage return

If the platform where camel-bindy will run is not Windows but Macintosh or Unix, then you can change the crlf property like this. Three values are available : WINDOWS, UNIX or MAC

```
@CsvRecord(separator = ",", crlf="MAC")
public Class Order {
}
}
```

Additionally, if for some reason you need to add a different line ending character, you can opt to specify it using the crlf parameter. In the following example, we can end the line with a comma followed by the newline character:

```
@CsvRecord(separator = ",", crlf=",\n")
public Class Order {
}
}
```

case 8 : isOrdered

Sometimes, the order to follow during the creation of the CSV record from the model is different from the order used during the parsing. Then, in this case, we can use the attribute **isOrdered = true** to indicate this in combination with attribute **position** of the DataField annotation.

```
@CsvRecord(isOrdered = true)
```

```
public Class Order {

    @DataField(pos = 1, position = 11)
    private int orderNr;

    @DataField(pos = 2, position = 10)
    private String clientNr;

}
```

pos is used to parse the file stream, while **position** is used to generate the CSV.

57.2.2. 2. Link

The link annotation will allow to link objects together.

Annotation name	Record type	Level
Link	all	Class & Property

Parameter name	Type	Required	Default value	Info
linkType	LinkType		OneToOne	Type of link identifying the relation between the classes

Only one-to-one relation is allowed as of the current version.

E.g : If the model class Client is linked to the Order class, then use annotation Link in the Order class like this :

Property Link

```
@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @Link
    private Client client;

}
```

And for the class Client :

Class Link

```
@Link
public class Client {
}
```

57.2.3. 3. DataField

The DataField annotation defines the property of the field. Each datafield is identified by its position in the record, a type (string, int, date, ...) and optionally of a pattern.

Annotation name	Record type	Level
DataField	all	Property

Parameter name	Type	Required	Default value	Info
pos	int	✓		Position of the data in the input record, must start from 1 (mandatory). See the position parameter.
align	String		R	Align the text to the right or left. Use values <code><tt>R</tt></code> or <code><tt>L</tt></code> .
clip	boolean		false	Indicates to clip data in the field if it exceeds the allowed length when using fixed length.
columnName	String			Name of the header column (optional). Uses the name of the property as default. Only applicable when CsvRecord has generateHeaderColumns = true
decimalSeparator	String			Decimal Separator to be used with BigDecimal number
defaultValue	String			Field's default value in case no value is set
delimiter	String			Optional delimiter to be used if the field has a variable length
groupingSeparator	String			Grouping Separator to be used with BigDecimal number when we would like to format/parse to number with grouping e.g. 123,456.789

Parameter name	Type	Required	Default value	Info
impliedDecimalSeparator	boolean		false	Indicates if there is a decimal point implied at a specified location
length	int		0	Length of the data block (number of characters) if the record is set to a fixed length
lengthPos	int		0	Identifies a data field in the record that defines the expected fixed length for this field
method	String			Method name to call to apply such customization on DataField. This must be the method on the datafield itself or you must provide static fully qualified name of the class's method e.g: see unit test org.apache.camel.dataformat.bindy.csv.BindySimpleCsvFunctionWithExternalMethodTest.replaceToBar
name	String			Name of the field (optional)
paddingChar	char			The char to pad with if the record is set to a fixed length
pattern	String			Pattern that the Java formatter (SimpleDateFormat by example) will use to transform the data (optional). If using pattern, then setting locale on bindy data format is recommended. Either set to a known locale such as "us" or use "default" to use platform default locale.
position	int		0	Position of the field in the output message generated (should start from 1). Must be used when the position of the field in the CSV generated (output message) must be different compare to input position (pos). See the pos parameter.
precision	int		0	precision of the {@link java.math.BigDecimal} number to be created
required	boolean		false	Indicates if the field is mandatory
rounding	String		CEILING	Round mode to be used to round/scale a BigDecimal Values : UP, DOWN, CEILING, FLOOR, HALF_UP, HALF_DOWN, HALF_EVEN, UNNECESSARY e.g : Number = 123456.789, Precision = 2, Rounding = CEILING Result : 123456.79

Parameter name	Type	Required	Default value	Info
timezone	String			Timezone to be used.
trim	boolean		false	Indicates if the value should be trimmed

case 1: pos

This parameter/attribute represents the position of the field in the CSV record.

Position

```
@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @DataField(pos = 5)
    private String isinCode;

}
```

As you can see in this example the position starts at **1** but continues at **5** in the class Order. The numbers from **2** to **4** are defined in the class Client (see here after).

Position continues in another model class

```
public class Client {

    @DataField(pos = 2)
    private String clientNr;

    @DataField(pos = 3)
    private String firstName;

    @DataField(pos = 4)
    private String lastName;

}
```

case 2: pattern

The pattern allows to enrich or validates the format of your data

Pattern

```
@CsvRecord(separator = ",")
public class Order {
```

```

@DataField(pos = 1)
private int orderNr;

@DataField(pos = 5)
private String isinCode;

@DataField(name = "Name", pos = 6)
private String instrumentName;

@DataField(pos = 7, precision = 2)
private BigDecimal amount;

@DataField(pos = 8)
private String currency;

// pattern used during parsing or when the date is created
@DataField(pos = 9, pattern = "dd-MM-yyyy")
private Date orderDate;
}

```

case 3 : precision

The precision is helpful when you want to define the decimal part of your number.

Precision

```

@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @Link
    private Client client;

    @DataField(pos = 5)
    private String isinCode;

    @DataField(name = "Name", pos = 6)
    private String instrumentName;

    @DataField(pos = 7, precision = 2)
    private BigDecimal amount;

    @DataField(pos = 8)
    private String currency;

    @DataField(pos = 9, pattern = "dd-MM-yyyy")
    private Date orderDate;
}

```

case 4 : Position is different in output

The position attribute will inform bindy how to place the field in the CSV record generated. By default, the position used corresponds to the position defined with the attribute **pos**. If the position is different (that means that we have an asymmetric process comparing marshaling from unmarshaling) then we

can use **position** to indicate this.

Here is an example:

Position is different in output

```
@CsvRecord(separator = ",", isOrdered = true)
public class Order {

    // Positions of the fields start from 1 and not from 0

    @DataField(pos = 1, position = 11)
    private int orderNr;

    @DataField(pos = 2, position = 10)
    private String clientNr;

    @DataField(pos = 3, position = 9)
    private String firstName;

    @DataField(pos = 4, position = 8)
    private String lastName;

    @DataField(pos = 5, position = 7)
    private String instrumentCode;

    @DataField(pos = 6, position = 6)
    private String instrumentNumber;
}
```

This attribute of the annotation **@DataField** must be used in combination with attribute **isOrdered = true** of the annotation **@CsvRecord**.

case 5 : required

If a field is mandatory, simply use the attribute **required** set to true.

Required

```
@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @DataField(pos = 2, required = true)
    private String clientNr;

    @DataField(pos = 3, required = true)
    private String firstName;

    @DataField(pos = 4, required = true)
    private String lastName;
}
```

If this field is not present in the record, then an error will be raised by the parser with the following information :

Some fields are missing (optional or mandatory), line :

case 6 : trim

If a field has leading and/or trailing spaces which should be removed before they are processed, simply use the attribute **trim** set to true.

Trim

```
@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1, trim = true)
    private int orderNr;

    @DataField(pos = 2, trim = true)
    private Integer clientNr;

    @DataField(pos = 3, required = true)
    private String firstName;

    @DataField(pos = 4)
    private String lastName;
}
```

case 7 : defaultValue

If a field is not defined then uses the value indicated by the **defaultValue** attribute.

Default value

```
@CsvRecord(separator = ",")
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @DataField(pos = 2)
    private Integer clientNr;

    @DataField(pos = 3, required = true)
    private String firstName;

    @DataField(pos = 4, defaultValue = "Barin")
    private String lastName;
}
```

case 8 : columnName

Specifies the column name for the property only if **@CsvRecord** has annotation **generateHeaderColumns = true**.

Column Name

```
@CsvRecord(separator = ";", generateHeaderColumns = true)
public class Order {

    @DataField(pos = 1)
    private int orderNr;

    @DataField(pos = 5, columnName = "ISIN")
    private String isinCode;

    @DataField(name = "Name", pos = 6)
    private String instrumentName;
}
```

This attribute is only applicable to optional fields.

57.2.4. 4. FixedLengthRecord

The `FixedLengthRecord` annotation is used to identify the root class of the model. It represents a record = "a line of a file/message containing data fixed length (number of characters) formatted" and can be linked to several children model classes. This format is a bit particular because data of a field can be aligned to the right or to the left.

When the size of the data does not fill completely the length of the field, we can then add 'pad' characters.

Annotation name	Record type	Level
<code>FixedLengthRecord</code>	fixed	Class

Parameter name	Type	Required	Default value	Info
<code>countGrapheme</code>	boolean		false	Indicates how chars are counted
<code>crlf</code>	String		WINDOWS	Character to be used to add a carriage return after each record (optional). Possible values: WINDOWS, UNIX, MAC, or custom. This option is used only during marshalling, whereas unmarshalling uses system default JDK provided line delimiter unless eol is customized.

Parameter name	Type	Required	Default value	Info
eol	String			Character to be used to process considering end of line after each record while unmarshalling (optional - default: "", which help default JDK provided line delimiter to be used unless any other line delimiter provided) This option is used only during unmarshalling, where marshalling uses system default provided line delimiter as "WINDOWS" unless any other value is provided.
footer	Class		void	Indicates that the record(s) of this type may be followed by a single footer record at the end of the file
header	Class		void	Indicates that the record(s) of this type may be preceded by a single header record at the beginning of in the file
ignoreMissingChars	boolean		false	Indicates whether too short lines will be ignored
ignoreTrailingChars	boolean		false	Indicates that characters beyond the last mapped field can be ignored when unmarshalling / parsing. This annotation is associated to the root class of the model and must be declared one time.
length	int		0	The fixed length of the record (number of characters). It means that the record will always be that long padded with <code>\{#paddingChar()\}'s</code>
name	String			Name describing the record (optional)
paddingChar	char			The char to pad with.
skipFooter	boolean		false	Configures the data format to skip marshalling / unmarshalling of the footer record. Configure this parameter on the primary record (e.g., not the header or footer).
skipHeader	boolean		false	Configures the data format to skip marshalling / unmarshalling of the header record. Configure this parameter on the primary record (e.g., not the header or footer).

A record may not be both a header/footer and a primary fixed-length record.

case 1: Simple fixed length record

This simple example shows how to design the model to parse/format a fixed message

```
10A9PaulineMISINXD12345678BUYShare2500.45USD01-08-2009
```

Fixed-simple

```
@FixedLengthRecord(length=54, paddingChar=' ')
public static class Order {

    @DataField(pos = 1, length=2)
    private int orderNr;

    @DataField(pos = 3, length=2)
    private String clientNr;

    @DataField(pos = 5, length=7)
    private String firstName;

    @DataField(pos = 12, length=1, align="L")
    private String lastName;

    @DataField(pos = 13, length=4)
    private String instrumentCode;

    @DataField(pos = 17, length=10)
    private String instrumentNumber;

    @DataField(pos = 27, length=3)
    private String orderType;

    @DataField(pos = 30, length=5)
    private String instrumentType;

    @DataField(pos = 35, precision = 2, length=7)
    private BigDecimal amount;

    @DataField(pos = 42, length=3)
    private String currency;

    @DataField(pos = 45, length=10, pattern = "dd-MM-yyyy")
    private Date orderDate;
}
```

case 2 : Fixed length record with alignment and padding

This more elaborated example show how to define the alignment for a field and how to assign a padding character which is ' ' here:

```
10A9 PaulineM ISINXD12345678BUYShare2500.45USD01-08-2009
```

Fixed-padding-align

```
@FixedLengthRecord(length=60, paddingChar=' ')
public static class Order {

    @DataField(pos = 1, length=2)
    private int orderNr;
```

```

@DataField(pos = 3, length=2)
private String clientNr;

@DataField(pos = 5, length=9)
private String firstName;

@DataField(pos = 14, length=5, align="L") // align text to the LEFT zone of the block
private String lastName;

@DataField(pos = 19, length=4)
private String instrumentCode;

@DataField(pos = 23, length=10)
private String instrumentNumber;

@DataField(pos = 33, length=3)
private String orderType;

@DataField(pos = 36, length=5)
private String instrumentType;

@DataField(pos = 41, precision = 2, length=7)
private BigDecimal amount;

@DataField(pos = 48, length=3)
private String currency;

@DataField(pos = 51, length=10, pattern = "dd-MM-yyyy")
private Date orderDate;
}

```

case 3 : Field padding

Sometimes, the default padding defined for record cannot be applied to the field as we have a number format where we would like to pad with '0' instead of ' '. In this case, you can use in the model the attribute **paddingChar** on **@DataField** to set this value.

```
10A9 PaulineM ISINXD12345678BUYShare000002500.45USD01-08-2009
```

Fixed-padding-field

```

@FixedLengthRecord(length = 65, paddingChar = ' ')
public static class Order {

    @DataField(pos = 1, length = 2)
    private int orderNr;

    @DataField(pos = 3, length = 2)
    private String clientNr;

    @DataField(pos = 5, length = 9)
    private String firstName;

    @DataField(pos = 14, length = 5, align = "L")

```



```

private String lastName;

@DataField(pos = 19, length = 4)
private String instrumentCode;

@DataField(pos = 23, length = 10)
private String instrumentNumber;

@DataField(pos = 33, length = 3)
private String orderType;

@DataField(pos = 36, length = 5)
private String instrumentType;

@DataField(pos = 41, precision = 2, length = 12, paddingChar = '0')
private BigDecimal amount;

@DataField(pos = 53, length = 3)
private String currency;

@DataField(pos = 56, length = 10, pattern = "dd-MM-yyyy")
private Date orderDate;
}

```

case 4: Fixed length record with delimiter

Fixed-length records sometimes have delimited content within the record. The firstName and lastName fields are delimited with the ^ character in the following example:

```
10A9Pauline^M^ISINXD12345678BUYShare000002500.45USD01-08-2009
```

Fixed-delimited

```

@FixedLengthRecord
public static class Order {

    @DataField(pos = 1, length = 2)
    private int orderNr;

    @DataField(pos = 2, length = 2)
    private String clientNr;

    @DataField(pos = 3, delimiter = "^")
    private String firstName;

    @DataField(pos = 4, delimiter = "^")
    private String lastName;

    @DataField(pos = 5, length = 4)
    private String instrumentCode;

    @DataField(pos = 6, length = 10)
    private String instrumentNumber;

    @DataField(pos = 7, length = 3)

```

```

private String orderType;

@DataField(pos = 8, length = 5)
private String instrumentType;

@DataField(pos = 9, precision = 2, length = 12, paddingChar = '0')
private BigDecimal amount;

@DataField(pos = 10, length = 3)
private String currency;

@DataField(pos = 11, length = 10, pattern = "dd-MM-yyyy")
private Date orderDate;
}

```

The **pos** value(s) in a fixed-length record may optionally be defined using ordinal, sequential values instead of precise column numbers.

case 5 : Fixed length record with record-defined field length

Occasionally a fixed-length record may contain a field that define the expected length of another field within the same record. In the following example the length of the **instrumentNumber** field value is defined by the value of **instrumentNumberLen** field in the record.

```
10A9Pauline^M^ISIN10XD12345678BUYShare000002500.45USD01-08-2009
```

Fixed-delimited

```

@FixedLengthRecord
public static class Order {

    @DataField(pos = 1, length = 2)
    private int orderNr;

    @DataField(pos = 2, length = 2)
    private String clientNr;

    @DataField(pos = 3, delimiter = "^")
    private String firstName;

    @DataField(pos = 4, delimiter = "^")
    private String lastName;

    @DataField(pos = 5, length = 4)
    private String instrumentCode;

    @DataField(pos = 6, length = 2, align = "R", paddingChar = '0')
    private int instrumentNumberLen;

    @DataField(pos = 7, lengthPos=6)
    private String instrumentNumber;

    @DataField(pos = 8, length = 3)
    private String orderType;
}

```

```

@DataField(pos = 9, length = 5)
private String instrumentType;

@DataField(pos = 10, precision = 2, length = 12, paddingChar = '0')
private BigDecimal amount;

@DataField(pos = 11, length = 3)
private String currency;

@DataField(pos = 12, length = 10, pattern = "dd-MM-yyyy")
private Date orderDate;
}

```

case 6 : Fixed length record with header and footer

Bindy will discover fixed-length header and footer records that are configured as part of the model – provided that the annotated classes exist either in the same package as the primary **@FixedLengthRecord** class, or within one of the configured scan packages. The following text illustrates two fixed-length records that are bracketed by a header record and footer record.

```

101-08-2009
10A9 PaulineM ISINXD12345678BUYShare000002500.45USD01-08-2009
10A9 RichN ISINXD12345678BUYShare000002700.45USD01-08-2009
9000000002

```

Fixed-header-and-footer-main-class

```

@FixedLengthRecord(header = OrderHeader.class, footer = OrderFooter.class)
public class Order {

    @DataField(pos = 1, length = 2)
    private int orderNr;

    @DataField(pos = 2, length = 2)
    private String clientNr;

    @DataField(pos = 3, length = 9)
    private String firstName;

    @DataField(pos = 4, length = 5, align = "L")
    private String lastName;

    @DataField(pos = 5, length = 4)
    private String instrumentCode;

    @DataField(pos = 6, length = 10)
    private String instrumentNumber;

    @DataField(pos = 7, length = 3)
    private String orderType;

    @DataField(pos = 8, length = 5)
    private String instrumentType;

    @DataField(pos = 9, precision = 2, length = 12, paddingChar = '0')

```

```

private BigDecimal amount;

@DataField(pos = 10, length = 3)
private String currency;

@DataField(pos = 11, length = 10, pattern = "dd-MM-yyyy")
private Date orderDate;
}

@FixedLengthRecord
public class OrderHeader {
    @DataField(pos = 1, length = 1)
    private int recordType = 1;

    @DataField(pos = 2, length = 10, pattern = "dd-MM-yyyy")
    private Date recordDate;
}

@FixedLengthRecord
public class OrderFooter {

    @DataField(pos = 1, length = 1)
    private int recordType = 9;

    @DataField(pos = 2, length = 9, align = "R", paddingChar = '0')
    private int numberOfRecordsInTheFile;
}

```

case 7 : Skipping content when parsing a fixed length record

It is common to integrate with systems that provide fixed-length records containing more information than needed for the target use case. It is useful in this situation to skip the declaration and parsing of those fields that we do not need. To accommodate this, Bindy will skip forward to the next mapped field within a record if the **pos** value of the next declared field is beyond the cursor position of the last parsed field. Using absolute **pos** locations for the fields of interest (instead of ordinal values) causes Bindy to skip content between two fields.

Similarly, it is possible that none of the content beyond some field is of interest. In this case, you can tell Bindy to skip parsing of everything beyond the last mapped field by setting the **ignoreTrailingChars** property on the **@FixedLengthRecord** declaration.

```

@FixedLengthRecord(ignoreTrailingChars = true)
public static class Order {

    @DataField(pos = 1, length = 2)
    private int orderNr;

    @DataField(pos = 3, length = 2)
    private String clientNr;

    // any characters that appear beyond the last mapped field will be ignored
}

```

57.2.5. 5. Message

The Message annotation is used to identify the class of your model who will contain key value pairs fields. This kind of format is used mainly in Financial Exchange Protocol Messages (FIX). Nevertheless, this annotation can be used for any other format where data are identified by keys. The key pair values are separated each other by a separator which can be a special character like a tab delimiter (unicode representation : `\u0009`) or a start of heading (unicode representation : `\u0001`)



NOTE

To work with FIX messages, the model must contain a Header and Trailer classes linked to the root message class which could be a Order class. This is not mandatory but will be very helpful when you will use camel-bindy in combination with camel-fix which is a Fix gateway based on quickFix project .

Annotation name	Record type	Level
Message	key value pair	Class

Parameter name	Type	Required	Default value	Info
keyValuePairSeparator	String	✓		Key value pair separator is used to split the values from their keys (mandatory). Can be <code>'\u0001'</code> , <code>'\u0009'</code> , <code>'#'</code> , or <code>'anything'</code> .
pairSeparator	String	✓		Pair separator used to split the key value pairs in tokens (mandatory). Can be <code>'='</code> , <code>','</code> , or <code>'anything'</code> .
crlf	String		WINDOWS	Character to be used to add a carriage return after each record (optional). Possible values = WINDOWS, UNIX, MAC, or custom. If you specify a value other than the three listed before, the value you enter (custom) will be used as the CRLF character(s).
isOrdered	boolean		false	Indicates if the message must be ordered in output. This annotation is associated to the message class of the model and must be declared one time.
name	String			Name describing the message (optional)
type	String		FIX	type is used to define the type of the message (e.g. FIX, EMX, ...) (optional)
version	String		4.1	version defines the version of the message (e.g. 4.1, ...) (optional)

case 1: separator = `'\u0001'`

The separator used to segregate the key value pair fields in a FIX message is the ASCII **01** character or in unicode format **\u0001**. This character must be escaped a second time to avoid a java runtime error. Here is an example :

```
8=FIX.4.1 9=20 34=1 35=0 49=INVMGR 56=BRKR 1=BE.CHM.001 11=CHM0001-01 22=4 ...
```

and how to use the annotation:

FIX - message

```
@Message(keyValuePairSeparator = "=", pairSeparator = "\u0001", type="FIX", version="4.1")
public class Order {
}

```

Look at test cases

The ASCII character like tab, ... cannot be displayed in WIKI page. So, have a look to the test case of camel-bindy to see exactly how the FIX message looks like (<https://github.com/apache/camel/blob/main/components/camel-bindy/src/test/data/fix/fix.txt>) and the Order, Trailer, Header classes (<https://github.com/apache/camel/blob/main/components/camel-bindy/src/test/java/org/apache/camel/dataformat/bindy/model/fix/simple/Order.java>).

57.2.6. 6. KeyValuePairField

The KeyValuePairField annotation defines the property of a key value pair field. Each KeyValuePairField is identified by a tag (= key) and its value associated, a type (string, int, date, ...), optionally a pattern and if the field is required.

Annotation name	Record type	Level
KeyValuePairField	Key Value Pair - FIX	Property

Parameter name	Type	Required	Default value	Info
tag	int	✓		tag identifying the field in the message (mandatory) - must be unique
impliedDecimalSeparator	boolean		false	Camel 2.11: Indicates if there is a decimal point implied at a specified location
name	String			name of the field (optional)
pattern	String			pattern that the formater will use to transform the data (optional)

Parameter name	Type	Required	Default value	Info
position	int		0	Position of the field in the message generated - must be used when the position of the key/tag in the FIX message must be different
precision	int		0	precision of the BigDecimal number to be created
required	boolean		false	Indicates if the field is mandatory
timezone	String			Timezone to be used.

case 1: tag

This parameter represents the key of the field in the message:

FIX message - Tag

```
@Message(keyValuePairSeparator = "=", pairSeparator = "\u0001", type="FIX", version="4.1")
public class Order {

    @Link Header header;

    @Link Trailer trailer;

    @KeyValuePairField(tag = 1) // Client reference
    private String Account;

    @KeyValuePairField(tag = 11) // Order reference
    private String ClOrdId;

    @KeyValuePairField(tag = 22) // Fund ID type (Sedol, ISIN, ...)
    private String IDSource;

    @KeyValuePairField(tag = 48) // Fund code
    private String SecurityId;

    @KeyValuePairField(tag = 54) // Movement type ( 1 = Buy, 2 = sell)
    private String Side;

    @KeyValuePairField(tag = 58) // Free text
    private String Text;
}
```

case 2 : Different position in output

If the tags/keys that we will put in the FIX message must be sorted according to a predefined order, then use the attribute **position** of the annotation **@KeyValuePairField**.

FIX message - Tag - sort

```

@Message(keyValuePairSeparator = "=", pairSeparator = "\\u0001", type = "FIX", version = "4.1",
isOrdered = true)
public class Order {

    @Link Header header;

    @Link Trailer trailer;

    @KeyValuePairField(tag = 1, position = 1) // Client reference
    private String account;

    @KeyValuePairField(tag = 11, position = 3) // Order reference
    private String clOrdId;
}

```

57.2.7. 7. Section

In FIX message of fixed length records, it is common to have different sections in the representation of the information : header, body and section. The purpose of the annotation **@Section** is to inform bindy about which class of the model represents the header (= section 1), body (= section 2) and footer (= section 3)

Only one attribute/parameter exists for this annotation.

Annotation name	Record type	Level
Section	FIX	Class

Parameter name	Type	Required	Default value	Info
number	int	✓		Number of the section

case 1: Section

Definition of the header section:

FIX message - Section - Header

```

@Section(number = 1)
public class Header {

    @KeyValuePairField(tag = 8, position = 1) // Message Header
    private String beginString;
}

```



```

    @KeyValuePairField(tag = 9, position = 2) // Checksum
    private int bodyLength;
}

```

Definition of the body section:

FIX message - Section - Body

```

@Section(number = 2)
@Message(keyValuePairSeparator = "=", pairSeparator = "\\u0001", type = "FIX", version = "4.1",
isOrdered = true)
public class Order {

    @Link Header header;

    @Link Trailer trailer;

    @KeyValuePairField(tag = 1, position = 1) // Client reference
    private String account;

    @KeyValuePairField(tag = 11, position = 3) // Order reference
    private String clOrdId;
}

```

Definition of the footer section:

FIX message - Section - Footer

```

@Section(number = 3)
public class Trailer {

    @KeyValuePairField(tag = 10, position = 1)
    // CheckSum
    private int checkSum;

    public int getCheckSum() {
        return checkSum;
    }
}

```

57.2.8. 8. OneToMany

The purpose of the annotation **@OneToMany** is to allow to work with a **List<?>** field defined a POJO class or from a record containing repetitive groups.



NOTE

Restrictions for OneToMany

Be careful, the one to many of bindy does not allow to handle repetitions defined on several levels of the hierarchy.

The relation OneToMany ONLY WORKS in the following cases :

- Reading a FIX message containing repetitive groups (= group of tags/keys)
- Generating a CSV with repetitive data

Annotation name	Record type	Level
OneToMany	all	Property

Parameter name	Type	Required	Default value	Info
mappedTo	String			Class name associated to the type of the List<Type of the Class>

case 1: Generating CSV with repetitive data

Here is the CSV output that we want :

```
Claus,Ibsen,Camel in Action 1,2010,35
Claus,Ibsen,Camel in Action 2,2012,35
Claus,Ibsen,Camel in Action 3,2013,35
Claus,Ibsen,Camel in Action 4,2014,35
```



NOTE

The repetitive data concern the title of the book and its publication date while first, last name and age are common and the classes used to modeling this. The Author class contains a List of Book.

Generate CSV with repetitive data

```
@CsvRecord(separator=",")
public class Author {

    @DataField(pos = 1)
    private String firstName;

    @DataField(pos = 2)
    private String lastName;

    @OneToMany
    private List<Book> books;

    @DataField(pos = 5)
    private String Age;
}

public class Book {

    @DataField(pos = 3)
    private String title;
```

```

    @DataField(pos = 4)
    private String year;
}

```

case 2 : Reading FIX message containing group of tags/keys

Here is the message that we would like to process in our model :

```

8=FIX 4.19=2034=135=049=INVMGR56=BRKR
1=BE.CHM.00111=CHM0001-0158=this is a camel - bindy test
22=448=BE000124567854=1
22=548=BE000987654354=2
22=648=BE000999999954=3
10=220

```

Tags 22, 48 and 54 are repeated.

And the code:

Reading FIX message containing group of tags/keys

```

public class Order {

    @Link Header header;

    @Link Trailer trailer;

    @KeyValuePairField(tag = 1) // Client reference
    private String account;

    @KeyValuePairField(tag = 11) // Order reference
    private String clOrdId;

    @KeyValuePairField(tag = 58) // Free text
    private String text;

    @OneToMany(mappedTo =
"org.apache.camel.dataformat.bindy.model.fix.complex.onetomany.Security")
    List<Security> securities;
}

public class Security {

    @KeyValuePairField(tag = 22) // Fund ID type (Sedol, ISIN, ...)
    private String idSource;

    @KeyValuePairField(tag = 48) // Fund code
    private String securityCode;

    @KeyValuePairField(tag = 54) // Movement type ( 1 = Buy, 2 = sell)
    private String side;
}

```

57.2.9. 9. BindyConverter

The purpose of the annotation **@BindyConverter** is define a converter to be used on field level. The provided class must implement the `Format` interface.

```
@FixedLengthRecord(length = 10, paddingChar = ' ')
public static class DataModel {
    @DataField(pos = 1, length = 10, trim = true)
    @BindyConverter(CustomConverter.class)
    public String field1;
}

public static class CustomConverter implements Format<String> {
    @Override
    public String format(String object) throws Exception {
        return (new StringBuilder(object)).reverse().toString();
    }

    @Override
    public String parse(String string) throws Exception {
        return (new StringBuilder(string)).reverse().toString();
    }
}
```

57.2.10.10. FormatFactories

The purpose of the annotation **@FormatFactories** is to define a set of converters at record-level. The provided classes must implement the **FormatFactoryInterface** interface.

```
@CsvRecord(separator = ",")
@FormatFactories({OrderNumberFormatFactory.class})
public static class Order {

    @DataField(pos = 1)
    private OrderNumber orderNr;

    @DataField(pos = 2)
    private String firstName;
}

public static class OrderNumber {
    private int orderNr;

    public static OrderNumber ofString(String orderNumber) {
        OrderNumber result = new OrderNumber();
        result.orderNr = Integer.valueOf(orderNumber);
        return result;
    }
}

public static class OrderNumberFormatFactory extends AbstractFormatFactory {

    {
        supportedClasses.add(OrderNumber.class);
    }

    @Override
```

```
public Format<?> build(FormattingOptions formattingOptions) {
    return new Format<OrderNumber>() {
        @Override
        public String format(OrderNumber object) throws Exception {
            return String.valueOf(object.orderNr);
        }

        @Override
        public OrderNumber parse(String string) throws Exception {
            return OrderNumber.ofString(string);
        }
    };
}
```

57.3. SUPPORTED DATATYPES

The DefaultFormatFactory makes formatting of the following datatype available by returning an instance of the interface FormatFactoryInterface based on the provided FormattingOptions:

- BigDecimal
- BigInteger
- Boolean
- Byte
- Character
- Date
- Double
- Enums
- Float
- Integer
- LocalDate
- LocalDateTime
- LocalTime
- Long
- Short
- String

The DefaultFormatFactory can be overridden by providing an instance of FactoryRegistry in the registry in use (e.g. spring or JNDI).

57.4. USING THE JAVA DSL

The next step instantiates the DataFormat *bindy* class associated with this record type and providing a class as a parameter.

For example the following uses the class **BindyCsvDataFormat** (which corresponds to the class associated with the CSV record type) which is configured with *com.acme.model.MyModel.class* to initialize the model objects configured in this package.

```
DataFormat bindy = new BindyCsvDataFormat(com.acme.model.MyModel.class);
```

57.4.1. Setting locale

Bindy supports configuring the locale on the dataformat, such as

```
BindyCsvDataFormat bindy = new BindyCsvDataFormat(com.acme.model.MyModel.class);
bindy.setLocale("us");
```

Or to use the platform default locale then use "default" as the locale name.

```
BindyCsvDataFormat bindy = new BindyCsvDataFormat(com.acme.model.MyModel.class);
bindy.setLocale("default");
```

57.4.2. Unmarshaling

```
from("file://inbox")
  .unmarshal(bindy)
  .to("direct:handleOrders");
```

Alternatively, you can use a named reference to a data format which can then be defined in your Registry e.g. your Spring XML file:

```
from("file://inbox")
  .unmarshal("myBindyDataFormat")
  .to("direct:handleOrders");
```

The Camel route will pick-up files in the inbox directory, unmarshall CSV records into a collection of model objects and send the collection to the route referenced by **handleOrders**.

The collection returned is a **List of Map** objects. Each Map within the list contains the model objects that were marshalled out of each line of the CSV. The reason behind this is that *each line can correspond to more than one object*. This can be confusing when you simply expect one object to be returned per line.

Each object can be retrieve using its class name.

```
List<Map<String, Object>> unmarshaledModels = (List<Map<String, Object>>)
exchange.getIn().getBody();

int modelCount = 0;
for (Map<String, Object> model : unmarshaledModels) {
  for (String className : model.keySet()) {
```

```

    Object obj = model.get(className);
    LOG.info("Count : " + modelCount + ", " + obj.toString());
  }
  modelCount++;
}

LOG.info("Total CSV records received by the csv bean : " + modelCount);

```

Assuming that you want to extract a single Order object from this map for processing in a route, you could use a combination of a Splitter and a Processor as per the following:

```

from("file://inbox")
  .unmarshal(bindy)
  .split(body())
  .process(new Processor() {
    public void process(Exchange exchange) throws Exception {
      Message in = exchange.getIn();
      Map<String, Object> modelMap = (Map<String, Object>) in.getBody();
      in.setBody(modelMap.get(Order.class.getCanonicalName()));
    }
  })
  .to("direct:handleSingleOrder")
  .end();

```

Take care of the fact that Bindy uses `CHARSET_NAME` property or the `CHARSET_NAME` header as define in the Exchange interface to do a charset conversion of the inputstream received for unmarshalling. In some producers (e.g. file-endpoint) you can define a charset. The charset conversion can already been done by this producer. Sometimes you need to remove this property or header from the exchange before sending it to the unmarshal. If you don't remove it the conversion might be done twice which might lead to unwanted results.

```

from("file://inbox?charset=Cp922")
  .removeProperty(Exchange.CHARSET_NAME)
  .unmarshal("myBindyDataFormat")
  .to("direct:handleOrders");

```

57.4.3. Marshaling

To generate CSV records from a collection of model objects, you create the following route :

```

from("direct:handleOrders")
  .marshal(bindy)
  .to("file://outbox")

```

57.5. USING SPRING XML

This is really easy to use Spring as your favorite DSL language to declare the routes to be used for camel-bindy. The following example shows two routes where the first will pick-up records from files, unmarshal the content and bind it to their model. The result is then send to a pojo (doing nothing special) and place them into a queue.

The second route will extract the pojos from the queue and marshal the content to generate a file containing the CSV record.

Spring DSL

```
<?xml version="1.0" encoding="UTF-8"?>

<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring
    http://camel.apache.org/schema/spring/camel-spring.xsd">

  <!-- Queuing engine - ActiveMq - work locally in mode virtual memory -->
  <bean id="activemq" class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="vm://localhost:61616"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <dataFormats>
      <bindy id="bindyDataformat" type="Csv" classType="org.apache.camel.bindy.model.Order"/>
    </dataFormats>

    <route>
      <from uri="file://src/data/csv/?noop=true" />
      <unmarshal ref="bindyDataformat" />
      <to uri="bean:csv" />
      <to uri="activemq:queue:in" />
    </route>

    <route>
      <from uri="activemq:queue:in" />
      <marshal ref="bindyDataformat" />
      <to uri="file://src/data/csv/out" />
    </route>
  </camelContext>
</beans>
```



NOTE

Please verify that your model classes implements serializable otherwise the queue manager will raise an error.

57.6. DEPENDENCIES

To use Bindy in your camel routes you need to add the a dependency on **camel-bindy** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-bindy</artifactId>
```



```
<version>3.14.5.redhat-00018</version>
</dependency>
```

57.7. SPRING BOOT AUTO-CONFIGURATION

When using `bindy-csv` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-bindy-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 18 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.bindy-csv.allow-empty-stream</code>	Whether to allow empty streams in the unmarshal process. If true, no exception will be thrown when a body without records is provided.	false	Boolean
<code>camel.dataformat.bindy-csv.class-type</code>	Name of model class to use.		String
<code>camel.dataformat.bindy-csv.enabled</code>	Whether to enable auto configuration of the <code>bindy-csv</code> data format. This is enabled by default.		Boolean
<code>camel.dataformat.bindy-csv.locale</code>	To configure a default locale to use, such as <code>us</code> for united states. To use the JVM platform default locale then use the name <code>default</code> .		String
<code>camel.dataformat.bindy-csv.type</code>	Whether to use <code>Csv</code> , <code>Fixed</code> , or <code>KeyValue</code> .		String
<code>camel.dataformat.bindy-csv.unwrap-single-instance</code>	When unmarshalling should a single instance be unwrapped and returned instead of wrapped in a <code>java.util.List</code> .	true	Boolean
<code>camel.dataformat.bindy-fixed.allow-empty-stream</code>	Whether to allow empty streams in the unmarshal process. If true, no exception will be thrown when a body without records is provided.	false	Boolean

Name	Description	Default	Type
<code>camel.dataformat.bindy-fixed.class-type</code>	Name of model class to use.		String
<code>camel.dataformat.bindy-fixed.enabled</code>	Whether to enable auto configuration of the bindy-fixed data format. This is enabled by default.		Boolean
<code>camel.dataformat.bindy-fixed.locale</code>	To configure a default locale to use, such as us for united states. To use the JVM platform default locale then use the name default.		String
<code>camel.dataformat.bindy-fixed.type</code>	Whether to use Csv, Fixed, or KeyValue.		String
<code>camel.dataformat.bindy-fixed.unwrap-single-instance</code>	When unmarshalling should a single instance be unwrapped and returned instead of wrapped in a <code>java.util.List</code> .	true	Boolean
<code>camel.dataformat.bindy-kvp.allow-empty-stream</code>	Whether to allow empty streams in the unmarshal process. If true, no exception will be thrown when a body without records is provided.	false	Boolean
<code>camel.dataformat.bindy-kvp.class-type</code>	Name of model class to use.		String
<code>camel.dataformat.bindy-kvp.enabled</code>	Whether to enable auto configuration of the bindy-kvp data format. This is enabled by default.		Boolean
<code>camel.dataformat.bindy-kvp.locale</code>	To configure a default locale to use, such as us for united states. To use the JVM platform default locale then use the name default.		String
<code>camel.dataformat.bindy-kvp.type</code>	Whether to use Csv, Fixed, or KeyValue.		String
<code>camel.dataformat.bindy-kvp.unwrap-single-instance</code>	When unmarshalling should a single instance be unwrapped and returned instead of wrapped in a <code>java.util.List</code> .	true	Boolean

CHAPTER 58. HL7

The HL7 component is used for working with the HL7 MLLP protocol and [HL7 v2 messages](#) using the [HAPI library](#).

This component supports the following:

- HL7 MLLP codec for [Mina](#)
- HL7 MLLP codec for [Netty](#)
- Type Converter from/to HAPI and String
- HL7 DataFormat using the HAPI library

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-hl7</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

58.1. HL7 MLLP PROTOCOL

HL7 is often used with the HL7 MLLP protocol, which is a text based TCP socket based protocol. This component ships with a Mina and Netty Codec that conforms to the MLLP protocol so you can easily expose an HL7 listener accepting HL7 requests over the TCP transport layer. To expose a HL7 listener service, the [camel-mina](#) or link:[camel-netty](#) component is used with the HL7MLLPCodec (mina) or HL7MLLPNettyDecoder/HL7MLLPNettyEncoder (Netty).

HL7 MLLP codec can be configured as follows:

Name	Default Value	Description
startByte	0x0b	The start byte spanning the HL7 payload.
endByte1	0x1c	The first end byte spanning the HL7 payload.
endByte2	0x0d	The 2nd end byte spanning the HL7 payload.
charset	JVM Default	The encoding (a charset name) to use for the codec. If not provided, Camel will use the JVM default Charset .
produceString	true	If true, the codec creates a string using the defined charset. If false, the codec sends a plain byte array into the route, so that the HL7 Data Format can determine the actual charset from the HL7 message content.
convertLFtoCR	false	Will convert <code>\n</code> to <code>\r</code> (0x0d , 13 decimal) as HL7 stipulates <code>\r</code> as segment terminators. The HAPI library requires the use of <code>\r</code> .

58.1.1. Exposing an HL7 listener using Mina

In the Spring XML file, we configure a mina endpoint to listen for HL7 requests using TCP on port **8888**:

```
<endpoint id="hl7MinaListener" uri="mina:tcp://localhost:8888?sync=true&codec=#hl7codec"/>
```

sync=true indicates that this listener is synchronous and therefore will return a HL7 response to the caller. The HL7 codec is setup with **codec=#hl7codec**. Note that **hl7codec** is just a Spring bean ID, so it could be named **mygreatcodecforhl7** or whatever. The codec is also set up in the Spring XML file:

```
<bean id="hl7codec" class="org.apache.camel.component.hl7.HL7MLLPCodec">
  <property name="charset" value="iso-8859-1"/>
</bean>
```

The endpoint **hl7MinaListener** can then be used in a route as a consumer, as this Java DSL example illustrates:

```
from("hl7MinaListener")
  .bean("patientLookupService");
```

This is a very simple route that will listen for HL7 and route it to a service named **patientLookupService**. This is also Spring bean ID, configured in the Spring XML as:

```
<bean id="patientLookupService"
class="com.mycompany.healthcare.service.PatientLookupService"/>
```

The business logic can be implemented in POJO classes that do not depend on Camel, as shown here:

```
import ca.uhn.hl7v2.HL7Exception;
import ca.uhn.hl7v2.model.Message;
import ca.uhn.hl7v2.model.v24.segment.QRD;

public class PatientLookupService {
  public Message lookupPatient(Message input) throws HL7Exception {
    QRD qrd = (QRD)input.get("QRD");
    String patientId = qrd.getWhoSubjectFilter(0).getIDNumber().getValue();

    // find patient data based on the patient id and create a HL7 model object with the response
    Message response = ... create and set response data
    return response
  }
}
```

58.1.2. Exposing an HL7 listener using Netty (available from Camel 2.15 onwards)

In the Spring XML file, we configure a netty endpoint to listen for HL7 requests using TCP on port **8888**:

```
<endpoint id="hl7NettyListener" uri="netty:tcp://localhost:8888?
sync=true&encoders=#hl7encoder&decoders=#hl7decoder"/>
```

sync=true indicates that this listener is synchronous and therefore will return a HL7 response to the caller. The HL7 codec is setup with **encoders=#hl7encoder*and*decoders=#hl7decoder**. Note that **hl7encoder** and **hl7decoder** are just bean IDs, so they could be named differently. The beans can be set in the Spring XML file:

```
<bean id="hl7decoder" class="org.apache.camel.component.hl7.HL7MLLPNettyDecoderFactory"/>
<bean id="hl7encoder" class="org.apache.camel.component.hl7.HL7MLLPNettyEncoderFactory"/>
```

The endpoint **hl7NettyListener** can then be used in a route as a consumer, as this Java DSL example illustrates:

```
from("hl7NettyListener")
    .bean("patientLookupService");
```

58.2. HL7 MODEL USING JAVA.LANG.STRING OR BYTE[]

The HL7 MLLP codec uses plain `String` as its data format. Camel uses its `Type Converter` to convert to/from strings to the HAPI HL7 model objects, but you can use the plain `String` objects if you prefer, for instance if you wish to parse the data yourself.

You can also let both the Mina and Netty codecs use a plain **byte[]** as its data format by setting the **produceString** property to `false`. The `Type Converter` is also capable of converting the **byte[]** to/from HAPI HL7 model objects.

58.3. HL7V2 MODEL USING HAPI

The HL7v2 model uses Java objects from the HAPI library. Using this library, you can encode and decode from the EDI format (ER7) that is mostly used with HL7v2.

The sample below is a request to lookup a patient with the patient ID **0101701234**.

```
MSH|^~\&|MYSENDER|MYRECEIVER|MYAPPLICATION||200612211200||QRY^A19|1234|P|2.4
QRD|200612211200|R||GetPatient|||1^RD|0101701234|DEM||
```

Using the HL7 model you can work with a **ca.uhn.hl7v2.model.Message** object, e.g. to retrieve a patient ID:

```
Message msg = exchange.getIn().getBody(Message.class);
QRD qrd = (QRD)msg.get("QRD");
String patientId = qrd.getWhoSubjectFilter(0).getIDNumber().getValue(); // 0101701234
```

This is powerful when combined with the HL7 listener, because you don't have to work with **byte[]**, **String** or any other simple object formats. You can just use the HAPI HL7v2 model objects. If you know the message type in advance, you can be more type-safe:

```
QRY_A19 msg = exchange.getIn().getBody(QRY_A19.class);
String patientId = msg.getQRD().getWhoSubjectFilter(0).getIDNumber().getValue();
```

58.4. HL7 DATAFORMAT

The **camel-hl7** JAR ships with a HL7 data format that can be used to marshal or unmarshal HL7 model objects.

The HL7 dataformat supports 1 options, which are listed below.

Name	Default	Java Type	Description
validate		Boolean	Whether to validate the HL7 message is by default true.

- **marshal** = from Message to byte stream (can be used when responding using the HL7 MLLP codec)
- **unmarshal** = from byte stream to Message (can be used when receiving streamed data from the HL7 MLLP)

To use the data format, simply instantiate an instance and invoke the marshal or unmarshal operation in the route builder:

```
DataFormat hl7 = new HL7DataFormat();

from("direct:hl7in")
    .marshal(hl7)
    .to("jms:queue:hl7out");
```

In the sample above, the HL7 is marshalled from a HAPI Message object to a byte stream and put on a JMS queue.

The next example is the opposite:

```
DataFormat hl7 = new HL7DataFormat();

from("jms:queue:hl7out")
    .unmarshal(hl7)
    .to("patientLookupService");
```

Here we unmarshal the byte stream into a HAPI Message object that is passed to our patient lookup service.

58.4.1. Segment separators

Unmarshalling does not automatically fix segment separators anymore by converting `\n` to `\r`. If you need this conversion, `org.apache.camel.component.hl7.HL7#convertLFToCR` provides a handy **Expression** for this purpose.

58.4.2. Charset

Both **marshal** and **unmarshal** evaluate the charset provided in the field **MSH-18**. If this field is empty, by default the charset contained in the corresponding Camel charset property/header is assumed. You can even change this default behavior by overriding the **guessCharsetName** method when inheriting from the **HL7DataFormat** class.

There is a shorthand syntax in Camel for well-known data formats that are commonly used. Then you don't need to create an instance of the **HL7DataFormat** object:

```
from("direct:hl7in")
    .marshal().hl7()
    .to("jms:queue:hl7out");
```

```

from("jms:queue:hl7out")
  .unmarshal().hl7()
  .to("patientLookupService");

```

58.5. MESSAGE HEADERS

The `unmarshal` operation adds these fields from the MSH segment as headers on the Camel message:

Key	MSH field	Example
CamelHL7SendingApplication	MSH-3	MYSERVER
CamelHL7SendingFacility	MSH-4	MYSERVERAPP
CamelHL7ReceivingApplication	MSH-5	MYCLIENT
CamelHL7ReceivingFacility	MSH-6	MYCLIENTAPP
CamelHL7Timestamp	MSH-7	20071231235900
CamelHL7Security	MSH-8	null
CamelHL7MessageType	MSH-9-1	ADT
CamelHL7TriggerEvent	MSH-9-2	A01
CamelHL7MessageControl	MSH-10	1234
CamelHL7ProcessingId	MSH-11	P
CamelHL7VersionId	MSH-12	2.4
CamelHL7Context	..	contains the that was used to parse the message

Key	MSH field	Example
CamelHL7Charset	MSH-18	UNICODE UTF-8

All headers except **CamelHL7Context** are **String** types. If a header value is missing, its value is **null**.

58.6. DEPENDENCIES

To use HL7 in your Camel routes you'll need to add a dependency on **camel-hl7** listed above, which implements this data format.

The HAPI library is split into a [base library](#) and several structure libraries, one for each HL7v2 message version:

- [v2.1 structures library](#)
- [v2.2 structures library](#)
- [v2.3 structures library](#)
- [v2.3.1 structures library](#)
- [v2.4 structures library](#)
- [v2.5 structures library](#)
- [v2.5.1 structures library](#)
- [v2.6 structures library](#)

By default **camel-hl7** only references the HAPI [base library](#). Applications are responsible for including structure libraries themselves. For example, if an application works with HL7v2 message versions 2.4 and 2.5 then the following dependencies must be added:

```
<dependency>
  <groupId>ca.uhn.hapi</groupId>
  <artifactId>hapi-structures-v24</artifactId>
  <version>2.2</version>
  <!-- use the same version as your hapi-base version -->
</dependency>
<dependency>
  <groupId>ca.uhn.hapi</groupId>
  <artifactId>hapi-structures-v25</artifactId>
  <version>2.2</version>
  <!-- use the same version as your hapi-base version -->
</dependency>
```

Alternatively, an OSGi bundle containing the base library, all structures libraries and required dependencies (on the bundle classpath) can be downloaded from the [central Maven repository](#).

```
<dependency>
  <groupId>ca.uhn.hapi</groupId>
```



```

<artifactId>hapi-osgi-base</artifactId>
<version>2.2</version>
</dependency>

```

58.7. SPRING BOOT AUTO-CONFIGURATION

When using hl7 with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-hl7-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.hl7.enabled</code>	Whether to enable auto configuration of the hl7 data format. This is enabled by default.		Boolean
<code>camel.dataformat.hl7.validate</code>	Whether to validate the HL7 message ls by default true.	true	Boolean
<code>camel.language.hl7terser.enabled</code>	Whether to enable auto configuration of the hl7terser language. This is enabled by default.		Boolean
<code>camel.language.hl7terser.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean

CHAPTER 59. JACKSONXML

Jackson XML is a Data Format which uses the [Jackson library](#) with the [XMLMapper extension](#) to unmarshal an XML payload into Java objects or to marshal Java objects into an XML payload. NOTE: If you are familiar with Jackson, this XML data format behaves in the same way as its JSON counterpart, and thus can be used with classes annotated for JSON serialization/deserialization.

This extension also mimics [JAXB's "Code first" approach](#).

This data format relies on [Woodstox](#) (especially for features like pretty printing), a fast and efficient XML processor.

```
from("activemq:My.Queue").
  unmarshal().jacksonxml().
  to("mqseries:Another.Queue");
```

59.1. JACKSONXML OPTIONS

The JacksonXML dataformat supports 15 options, which are listed below.

Name	Default	Java Type	Description
<code>xmlMapper</code>		String	Lookup and use the existing XmlMapper with the given id.
<code>prettyPrint</code>	false	Boolean	To enable pretty printing output nicely formatted. Is by default false.
<code>unmarshalType</code>		String	Class name of the java type to use when unmarshalling.
<code>jsonView</code>		String	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has JsonView annotations.
<code>include</code>		String	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to NON_NULL.
<code>allowJmsType</code>		Boolean	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.
<code>collectionType</code>		String	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.
<code>useList</code>		Boolean	To unmarshal to a List of Map or a List of Pojo.

Name	Default	Java Type	Description
<code>enableJaxbAnnotationModule</code>		Boolean	Whether to enable the JAXB annotations module when using Jackson. When enabled then JAXB annotations can be used by Jackson.
<code>moduleClassNames</code>		String	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.
<code>moduleRefs</code>		String	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.
<code>enableFeatures</code>		String	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> Multiple features can be separated by comma.
<code>disableFeatures</code>		String	Set of features to disable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> Multiple features can be separated by comma.
<code>allowUnmarshalType</code>		Boolean	If enabled then Jackson is allowed to attempt to use the <code>CamelJacksonUnmarshalType</code> header during the unmarshalling. This should only be enabled when desired to be used.
<code>contentTypeHeader</code>		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.

59.1.1. Using Jackson XML in Spring DSL

When using Data Format in Spring DSL you need to declare the data formats first. This is done in the `DataFormats` XML tag.

```

<dataFormats>
  <!-- here we define a Xml data format with the id jack and that it should use the TestPojo as
the class type when
doing unmarshal. The unmarshalType is optional, if not provided Camel will use a Map as
the type -->
  <jacksonxml id="jack" unmarshalType="org.apache.camel.component.jacksonxml.TestPojo"/>
</dataFormats>

```

And then you can refer to this id in the route:

```
<route>
  <from uri="direct:back"/>
  <unmarshal><custom ref="jack"/></unmarshal>
  <to uri="mock:reverse"/>
</route>
```

59.1.2. Excluding POJO fields from marshalling

When marshalling a POJO to XML you might want to exclude certain fields from the XML output. With Jackson you can use [JSON views](#) to accomplish this. First create one or more marker classes.

Use the marker classes with the `@JsonView` annotation to include/exclude certain fields. The annotation also works on getters.

Finally use the Camel `JacksonXMLDataFormat` to marshal the above POJO to XML.

Note that the weight field is missing in the resulting XML:

```
<pojo age="30" weight="70"/>
```

59.2. INCLUDE/EXCLUDE FIELDS USING THE JSONVIEW ATTRIBUTE WITH `JACKSONXML` DATAFORMAT

As an example of using this attribute you can instead of:

```
JacksonXMLDataFormat ageViewFormat = new JacksonXMLDataFormat(TestPojoView.class,
Views.Age.class);
from("direct:inPojoAgeView").
  marshal(ageViewFormat);
```

Directly specify your [JSON view](#) inside the Java DSL as:

```
from("direct:inPojoAgeView").
  marshal().jacksonxml(TestPojoView.class, Views.Age.class);
```

And the same in XML DSL:

```
<from uri="direct:inPojoAgeView"/>
<marshal>
  <jacksonxml unmarshalType="org.apache.camel.component.jacksonxml.TestPojoView"
jsonView="org.apache.camel.component.jacksonxml.Views$Age"/>
</marshal>
```

59.3. SETTING SERIALIZATION INCLUDE OPTION

If you want to marshal a pojo to XML, and the pojo has some fields with null values. And you want to skip these null values, then you need to set either an annotation on the pojo,

```
@JsonInclude(Include.NON_NULL)
public class MyPojo {
    ...
}
```

But this requires you to include that annotation in your pojo source code. You can also configure the Camel JacksonXMLDataFormat to set the include option, as shown below:

```
JacksonXMLDataFormat format = new JacksonXMLDataFormat();
format.setInclude("NON_NULL");
```

Or from XML DSL you configure this as

```
<dataFormats>
  <jacksonxml id="jacksonxml" include="NON_NULL"/>
</dataFormats>
```

59.4. UNMARSHALLING FROM XML TO POJO WITH DYNAMIC CLASS NAME

If you use jackson to unmarshal XML to POJO, then you can now specify a header in the message that indicate which class name to unmarshal to.

The header has key **CamelJacksonUnmarshalType** if that header is present in the message, then Jackson will use that as FQN for the POJO class to unmarshal the XML payload as.

For JMS end users there is the JMSType header from the JMS spec that indicates that also. To enable support for JMSType you would need to turn that on, on the jackson data format as shown:

```
JacksonDataFormat format = new JacksonDataFormat();
format.setAllowJmsType(true);
```

Or from XML DSL you configure this as

```
<dataFormats>
  <jacksonxml id="jacksonxml" allowJmsType="true"/>
</dataFormats>
```

59.5. UNMARSHALLING FROM XML TO LIST<MAP> OR LIST<POJO>

If you are using Jackson to unmarshal XML to a list of map/pojo, you can now specify this by setting **useList="true"** or use the **org.apache.camel.component.jacksonxml.ListJacksonXMLDataFormat**. For example with Java you can do as shown below:

```
JacksonXMLDataFormat format = new ListJacksonXMLDataFormat();
// or
JacksonXMLDataFormat format = new JacksonXMLDataFormat();
format.useList();
// and you can specify the pojo class type also
format.setUnmarshalType(MyPojo.class);
```

And if you use XML DSL then you configure to use list using **useList** attribute as shown below:

```
<dataFormats>
  <jacksonxml id="jack" useList="true"/>
</dataFormats>
```

And you can specify the pojo type also

```
<dataFormats>
  <jacksonxml id="jack" useList="true" unmarshalType="com.foo.MyPojo"/>
</dataFormats>
```

59.6. USING CUSTOM JACKSON MODULES

You can use custom Jackson modules by specifying the class names of those using the `moduleClassNames` option as shown below.

```
<dataFormats>
  <jacksonxml id="jack" useList="true" unmarshalType="com.foo.MyPojo"
  moduleClassNames="com.foo.MyModule,com.foo.MyOtherModule"/>
</dataFormats>
```

When using `moduleClassNames` then the custom jackson modules are not configured, by created using default constructor and used as-is. If a custom module needs any custom configuration, then an instance of the module can be created and configured, and then use `modulesRefs` to refer to the module as shown below:

```
<bean id="myJacksonModule" class="com.foo.MyModule">
  ... // configure the module as you want
</bean>

<dataFormats>
  <jacksonxml id="jacksonxml" useList="true" unmarshalType="com.foo.MyPojo"
  moduleRefs="myJacksonModule"/>
</dataFormats>
```

Multiple modules can be specified separated by comma, such as `moduleRefs="myJacksonModule,myOtherModule"`

59.7. ENABLING OR DISABLE FEATURES USING JACKSON

Jackson has a number of features you can enable or disable, which its `ObjectMapper` uses. For example to disable failing on unknown properties when marshalling, you can configure this using the `disableFeatures`:

```
<dataFormats>
  <jacksonxml id="jacksonxml" unmarshalType="com.foo.MyPojo"
  disableFeatures="FAIL_ON_UNKNOWN_PROPERTIES"/>
</dataFormats>
```

You can disable multiple features by separating the values using comma. The values for the features must be the name of the enums from Jackson from the following enum classes

- `com.fasterxml.jackson.databind.SerializationFeature`
- `com.fasterxml.jackson.databind.DeserializationFeature`
- `com.fasterxml.jackson.databind.MapperFeature`

To enable a feature use the `enableFeatures` options instead.

From Java code you can use the type safe methods from `camel-jackson` module:

```
JacksonDataFormat df = new JacksonDataFormat(MyPojo.class);
df.disableFeature(DeserializationFeature.FAIL_ON_UNKNOWN_PROPERTIES);
df.disableFeature(DeserializationFeature.FAIL_ON_NULL_FOR_PRIMITIVES);
```

59.8. CONVERTING MAPS TO POJO USING JACKSON

Jackson **ObjectMapper** can be used to convert maps to POJO objects. Jackson component comes with the data converter that can be used to convert `java.util.Map` instance to non-String, non-primitive and non-Number objects.

```
Map<String, Object> invoiceData = new HashMap<String, Object>();
invoiceData.put("netValue", 500);
producerTemplate.sendBody("direct:mapToInvoice", invoiceData);
...
// Later in the processor
Invoice invoice = exchange.getIn().getBody(Invoice.class);
```

If there is a single **ObjectMapper** instance available in the Camel registry, it will be used by the converter to perform the conversion. Otherwise the default mapper will be used.

59.9. FORMATTED XML MARSHALLING (PRETTY-PRINTING)

Using the **prettyPrint** option one can output a well formatted XML while marshalling:

```
<dataFormats>
  <jacksonxml id="jack" prettyPrint="true"/>
</dataFormats>
```

And in Java DSL:

```
from("direct:inPretty").marshal().jacksonxml(true);
```

Please note that there are 5 different overloaded **jacksonxml()** DSL methods which support the **prettyPrint** option in combination with other settings for **unmarshalType**, **jsonView** etc.

59.10. DEPENDENCIES

To use Jackson XML in your camel routes you need to add the dependency on `camel-jacksonxml` which implements this data format. If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
```

```

<groupId>org.apache.camel</groupId>
<artifactId>camel-jacksonxml</artifactId>
<version>3.14.5.redhat-00018</version>
<!-- use the same version as your Camel core version -->
</dependency>

```

59.11. SPRING BOOT AUTO-CONFIGURATION

When using jacksonxml with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jacksonxml-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 16 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.jacksonxml.allow-jms-type</code>	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.	false	Boolean
<code>camel.dataformat.jacksonxml.allow-unmarshall-type</code>	If enabled then Jackson is allowed to attempt to use the CamelJacksonUnmarshalType header during the unmarshalling. This should only be enabled when desired to be used.	false	Boolean
<code>camel.dataformat.jacksonxml.collection-type</code>	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.		String
<code>camel.dataformat.jacksonxml.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.jacksonxml.disable-features</code>	Set of features to disable on the Jackson com.fasterxml.jackson.databind.ObjectMapper. The features should be a name that matches a enum from com.fasterxml.jackson.databind.SerializationFeature, com.fasterxml.jackson.databind.DeserializationFeature, or com.fasterxml.jackson.databind.MapperFeature. Multiple features can be separated by comma.		String

Name	Description	Default	Type
<code>camel.dataformat.jacksonxml.enable-features</code>	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.jacksonxml.enable-jaxb-annotation-module</code>	Whether to enable the JAXB annotations module when using jackson. When enabled then JAXB annotations can be used by Jackson.	false	Boolean
<code>camel.dataformat.jacksonxml.enabled</code>	Whether to enable auto configuration of the jacksonxml data format. This is enabled by default.		Boolean
<code>camel.dataformat.jacksonxml.include</code>	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .		String
<code>camel.dataformat.jacksonxml.json-view</code>	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.		String
<code>camel.dataformat.jacksonxml.module-class-names</code>	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.		String
<code>camel.dataformat.jacksonxml.module-refs</code>	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.		String
<code>camel.dataformat.jacksonxml.pretty-print</code>	To enable pretty printing output nicely formatted. Is by default false.	false	Boolean
<code>camel.dataformat.jacksonxml.unmarshal-type</code>	Class name of the java type to use when unmarshalling.		String

Name	Description	Default	Type
<code>camel.dataformat.jacksonxml.use-list</code>	To unmarshal to a List of Map or a List of Pojo.	false	Boolean
<code>camel.dataformat.jacksonxml.xml-mapper</code>	Lookup and use the existing XmlMapper with the given id.		String

CHAPTER 60. JAXB

JAXB is a Data Format which uses the JAXB2 XML marshalling standard which is included in Java 6 to unmarshal an XML payload into Java objects or to marshal Java objects into an XML payload.

60.1. OPTIONS

The JAXB dataformat supports 19 options, which are listed below.

Name	Default	Java Type	Description
<code>contextPath</code>		String	Required Package name where your JAXB classes are located.
<code>contextPathsClassName</code>		Boolean	This can be set to true to mark that the contextPath is referring to a classname and not a package name.
<code>schema</code>		String	To validate against an existing schema. You can use the prefix classpath:, file: or http: to specify how the resource should be resolved. You can separate multiple schema files by using the ',' character.
<code>schemaSeverityLevel</code>		Enum	Sets the schema severity level to use when validating against a schema. This level determines the minimum severity error that triggers JAXB to stop continue parsing. The default value of 0 (warning) means that any error (warning, error or fatal error) will trigger JAXB to stop. There are the following three levels: 0=warning, 1=error, 2=fatal error. Enum values: <ul style="list-style-type: none"> • 0 • 1 • 2
<code>prettyPrint</code>		Boolean	To enable pretty printing output nicely formatted. Is by default false.
<code>objectFactory</code>		Boolean	Whether to allow using ObjectFactory classes to create the POJO classes during marshalling. This only applies to POJO classes that has not been annotated with JAXB and providing jaxb.index descriptor files.
<code>ignoreJAXBElement</code>		Boolean	Whether to ignore JAXBElement elements - only needed to be set to false in very special use-cases.

Name	Default	Java Type	Description
<code>mustBeJAXBElement</code>		Boolean	Whether marshalling must be java objects with JAXB annotations. And if not then it fails. This option can be set to false to relax that, such as when the data is already in XML format.
<code>filterNonXmlChars</code>		Boolean	To ignore non xml characters and replace them with an empty space.
<code>encoding</code>		String	To overrule and use a specific encoding.
<code>fragment</code>		Boolean	To turn on marshalling XML fragment trees. By default JAXB looks for <code>XmlRootElement</code> annotation on given class to operate on whole XML tree. This is useful but not always - sometimes generated code does not have <code>XmlRootElement</code> annotation, sometimes you need unmarshal only part of tree. In that case you can use partial unmarshalling. To enable this behaviours you need set property <code>partClass</code> . Camel will pass this class to JAXB's unmarshaller.
<code>partClass</code>		String	Name of class used for fragment parsing. See more details at the <code>fragment</code> option.
<code>partNamespace</code>		String	XML namespace to use for fragment parsing. See more details at the <code>fragment</code> option.
<code>namespacePrefixRef</code>		String	When marshalling using JAXB or SOAP then the JAXB implementation will automatic assign namespace prefixes, such as <code>ns2</code> , <code>ns3</code> , <code>ns4</code> etc. To control this mapping, Camel allows you to refer to a map which contains the desired mapping.
<code>xmlStreamWriterWrapper</code>		String	To use a custom xml stream writer.
<code>schemaLocation</code>		String	To define the location of the schema.
<code>noNamespaceSchemaLocation</code>		String	To define the location of the namespaceless schema.
<code>jaxbProviderProperties</code>		String	Refers to a custom <code>java.util.Map</code> to lookup in the registry containing custom JAXB provider properties to be used with the JAXB marshaller.
<code>contentTypeHeader</code>		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.

60.2. USING THE JAVA DSL

For example the following uses a named DataFormat of **jaxb** which is configured with a number of Java package names to initialize the `JAXBContext`.

```
DataFormat jaxb = new JaxbDataFormat("com.acme.model");

from("activemq:My.Queue").
  unmarshal(jaxb).
  to("mqseries:Another.Queue");
```

You can if you prefer use a named reference to a data format which can then be defined in your Registry such as via your Spring XML file. e.g.

```
from("activemq:My.Queue").
  unmarshal("myJaxbDataType").
  to("mqseries:Another.Queue");
```

60.3. USING SPRING XML

The following example shows how to configure the **JaxbDataFormat** and use it in multiple routes.

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
  http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd">

  <bean id="myJaxb" class="org.apache.camel.converter.jaxb.JaxbDataFormat">
    <property name="contextPath" value="org.apache.camel.example"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="direct:start"/>
      <marshal><custom ref="myJaxb"/></marshal>
      <to uri="direct:marshalled"/>
    </route>
    <route>
      <from uri="direct:marshalled"/>
      <unmarshal><custom ref="myJaxb"/></unmarshal>
      <to uri="mock:result"/>
    </route>
  </camelContext>

</beans>
```

Multiple context paths

It is possible to use this data format with more than one context path. You can specify context path using `:` as separator, for example **com.mycompany:com.mycompany2**. Note that this is handled by JAXB implementation and might change if you use different vendor than RI.

60.4. PARTIAL MARSHALLING/UNMARSHALLING

JAXB 2 supports marshalling and unmarshalling XML tree fragments. By default JAXB looks for **@XmlRootElement** annotation on given class to operate on whole XML tree. This is useful but not always - sometimes generated code does not have **@XmlRootElement** annotation, sometimes you need unmarshall only part of tree.

In that case you can use partial unmarshalling. To enable this behaviours you need set property **partClass**. Camel will pass this class to JAXB's unmarshaller. If **JaxbConstants.JAXB_PART_CLASS** is set as one of headers, (even if partClass property is set on DataFormat), the property on DataFormat is surpassed and the one set in the headers is used.

For marshalling you have to add **partNamespace** attribute with QName of destination namespace. Example of Spring DSL you can find above.

If **JaxbConstants.JAXB_PART_NAMESPACE** is set as one of headers, (even if partNamespace property is set on DataFormat), the property on DataFormat is surpassed and the one set in the headers is used. While setting **partNamespace** through **JaxbConstants.JAXB_PART_NAMESPACE**, please note that you need to specify its value `\{[namespaceUri]\}[localPart]`

```
...
.setHeader(JaxbConstants.JAXB_PART_NAMESPACE, simple("
{http://www.camel.apache.org/jaxb/example/address/1}address"));
...
```

60.5. FRAGMENT

JaxbDataFormat has new property fragment which can set the the **Marshaller.JAXB_FRAGMENT** encoding property on the JAXB Marshaller. If you don't want the JAXB Marshaller to generate the XML declaration, you can set this option to be true. The default value of this property is false.

60.6. IGNORING THE NONXML CHARACTER

JaxbDataFormat supports to ignore the [NonXML Character](#), you just need to set the filterNonXmlChars property to be true, JaxbDataFormat will replace the NonXML character with " " when it is marshaling or unmarshaling the message. You can also do it by setting the Exchange property **Exchange.FILTER_NON_XML_CHARS**.

	JDK 1.5	JDK 1.6+
Filtering in use	StAX API and implementation	No
Filtering not in use	StAX API only	No

This feature has been tested with Woodstox 3.2.9 and Sun JDK 1.6 StAX implementation.

JaxbDataFormat now allows you to customize the XMLStreamWriter used to marshal the stream to XML. Using this configuration, you can add your own stream writer to completely remove, escape, or replace non-xml characters.

```
JaxbDataFormat customWriterFormat = new JaxbDataFormat("org.apache.camel.foo.bar");
customWriterFormat.setXmlStreamWriterWrapper(new TestXmlStreamWriter());
```

The following example shows using the Spring DSL and also enabling Camel's NonXML filtering:

```
<bean id="testXmlStreamWriterWrapper" class="org.apache.camel.jaxb.TestXmlStreamWriter"/>
<jaxb filterNonXmlChars="true" contextPath="org.apache.camel.foo.bar"
xmlStreamWriterWrapper="#testXmlStreamWriterWrapper" />
```

60.7. WORKING WITH THE OBJECTFACTORY

If you use XJC to create the java class from the schema, you will get an ObjectFactory for you JAXB context. Since the ObjectFactory uses [JAXBElement](#) to hold the reference of the schema and element instance value, jaxbDataformat will ignore the JAXBElement by default and you will get the element instance value instead of the JAXBElement object form the unmarshaled message body.

If you want to get the JAXBElement object form the unmarshaled message body, you need to set the JaxbDataFormat object's ignoreJAXBElement property to be false.

60.8. SETTING ENCODING

You can set the **encoding** option to use when marshalling. Its the **Marshaller.JAXB_ENCODING** encoding property on the JAXB Marshaller.

You can setup which encoding to use when you declare the JAXB data format. You can also provide the encoding in the Exchange property **Exchange.CHARSET_NAME**. This property will overrule the encoding set on the JAXB data format.

In this Spring DSL we have defined to use **iso-8859-1** as the encoding.

60.9. CONTROLLING NAMESPACE PREFIX MAPPING

When marshalling using [JAXB](#) or [SOAP](#) then the JAXB implementation will automatic assign namespace prefixes, such as ns2, ns3, ns4 etc. To control this mapping, Camel allows you to refer to a map which contains the desired mapping.

Notice this requires having JAXB-RI 2.1 or better (from SUN) on the classpath, as the mapping functionality is dependent on the implementation of JAXB, whether its supported.

For example in Spring XML we can define a Map with the mapping. In the mapping file below, we map SOAP to use soap as prefix. While our custom namespace "http://www.mycompany.com/foo/2" is not using any prefix.

```
<util:map id="myMap">
  <entry key="http://www.w3.org/2003/05/soap-envelope" value="soap"/>
  <!-- we dont want any prefix for our namespace -->
  <entry key="http://www.mycompany.com/foo/2" value=""/>
</util:map>
```

To use this in [JAXB](#) or [SOAP](#) you refer to this map, using the `namespacePrefixRef` attribute as shown below. Then Camel will lookup in the Registry a `java.util.Map` with the id "myMap", which was what we defined above.

```
<marshal>
  <soapjxb version="1.2" contextPath="com.mycompany.foo" namespacePrefixRef="myMap"/>
</marshal>
```

60.10. SCHEMA VALIDATION

The JAXB Data Format supports validation by marshalling and unmarshalling from/to XML. You can use the prefix `classpath:`, `file:` or `http:` to specify how the resource should be resolved. You can separate multiple schema files by using the `'`, `'` character.

Using the Java DSL, you can configure it in the following way:

```
JaxbDataFormat jaxbDataFormat = new JaxbDataFormat();
jaxbDataFormat.setContextPath(Person.class.getPackage().getName());
jaxbDataFormat.setSchema("classpath:person.xsd,classpath:address.xsd");
```

You can do the same using the XML DSL:

```
<marshal>
  <jaxb id="jxb" schema="classpath:person.xsd,classpath:address.xsd"/>
</marshal>
```

Camel will create and pool the underlying **SchemaFactory** instances on the fly, because the **SchemaFactory** shipped with the JDK is not thread safe.

However, if you have a **SchemaFactory** implementation which is thread safe, you can configure the JAXB data format to use this one:

```
JaxbDataFormat jaxbDataFormat = new JaxbDataFormat();
jaxbDataFormat.setSchemaFactory(threadSafeSchemaFactory);
```

60.11. SCHEMA LOCATION

The JAXB Data Format supports to specify the SchemaLocation when marshaling the XML.

Using the Java DSL, you can configure it in the following way:

```
JaxbDataFormat jaxbDataFormat = new JaxbDataFormat();
jaxbDataFormat.setContextPath(Person.class.getPackage().getName());
jaxbDataFormat.setSchemaLocation("schema/person.xsd");
```

You can do the same using the XML DSL:

```
<marshal>
  <jaxb id="jxb" schemaLocation="schema/person.xsd"/>
</marshal>
```

60.12. MARSHAL DATA THAT IS ALREADY XML

The JAXB marshaller requires that the message body is JAXB compatible, eg its a JAXBElement, eg a java instance that has JAXB annotations, or extend JAXBElement. There can be situations where the message body is already in XML, eg from a String type.

There is a new option **mustBeJAXBElement** you can set to false, to relax this check, so the JAXB marshaller only attempts to marshal JAXBElements (javax.xml.bind.JAXBIntrospector#isElement returns true). And in those situations the marshaller fallbacks to marshal the message body as-is.

60.13. DEPENDENCIES

To use JAXB in your camel routes you need to add the a dependency on **camel-jaxb** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jaxb</artifactId>
  <version>3.14.5.redhat-00018</version>
</dependency>
```

60.14. SPRING BOOT AUTO-CONFIGURATION

When using jaxb with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jaxb-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 20 options, which are listed below.

Name	Description	Default	Type
camel.dataformat.jaxb.content-type-header	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.	true	Boolean
camel.dataformat.jaxb.context-path	Package name where your JAXB classes are located.		String

Name	Description	Default	Type
<code>camel.dataformat.jaxb.context-path-is-class-name</code>	This can be set to true to mark that the contextPath is referring to a classname and not a package name.	false	Boolean
<code>camel.dataformat.jaxb.enabled</code>	Whether to enable auto configuration of the jaxb data format. This is enabled by default.		Boolean
<code>camel.dataformat.jaxb.encoding</code>	To overrule and use a specific encoding.		String
<code>camel.dataformat.jaxb.filter-non-xml-chars</code>	To ignore non xml characheters and replace them with an empty space.	false	Boolean
<code>camel.dataformat.jaxb.fragment</code>	To turn on marshalling XML fragment trees. By default JAXB looks for XmlRootElement annotation on given class to operate on whole XML tree. This is useful but not always - sometimes generated code does not have XmlRootElement annotation, sometimes you need unmarshall only part of tree. In that case you can use partial unmarshalling. To enable this behaviours you need set property partClass. Camel will pass this class to JAXB's unmarshaller.	false	Boolean
<code>camel.dataformat.jaxb.ignore-j-a-x-b-element</code>	Whether to ignore JAXBELEMENT elements - only needed to be set to false in very special use-cases.	false	Boolean
<code>camel.dataformat.jaxb.jaxb-provider-properties</code>	Refers to a custom java.util.Map to lookup in the registry containing custom JAXB provider properties to be used with the JAXB marshaller.		String
<code>camel.dataformat.jaxb.must-be-j-a-x-b-element</code>	Whether marhsalling must be java objects with JAXB annotations. And if not then it fails. This option can be set to false to relax that, such as when the data is already in XML format.	false	Boolean
<code>camel.dataformat.jaxb.namespace-prefix-ref</code>	When marshalling using JAXB or SOAP then the JAXB implementation will automatic assign namespace prefixes, such as ns2, ns3, ns4 etc. To control this mapping, Camel allows you to refer to a map which contains the desired mapping.		String

Name	Description	Default	Type
<code>camel.dataformat.jaxb.no-namespace-schema-location</code>	To define the location of the namespaceless schema.		String
<code>camel.dataformat.jaxb.object-factory</code>	Whether to allow using ObjectFactory classes to create the POJO classes during marshalling. This only applies to POJO classes that has not been annotated with JAXB and providing jaxb.index descriptor files.	false	Boolean
<code>camel.dataformat.jaxb.part-class</code>	Name of class used for fragment parsing. See more details at the fragment option.		String
<code>camel.dataformat.jaxb.part-namespace</code>	XML namespace to use for fragment parsing. See more details at the fragment option.		String
<code>camel.dataformat.jaxb.pretty-print</code>	To enable pretty printing output nicely formatted. Is by default false.	false	Boolean
<code>camel.dataformat.jaxb.schema</code>	To validate against an existing schema. You can use the prefix classpath:, file: or http: to specify how the resource should be resolved. You can separate multiple schema files by using the ',' character.		String
<code>camel.dataformat.jaxb.schema-location</code>	To define the location of the schema.		String
<code>camel.dataformat.jaxb.schema-severity-level</code>	Sets the schema severity level to use when validating against a schema. This level determines the minimum severity error that triggers JAXB to stop continue parsing. The default value of 0 (warning) means that any error (warning, error or fatal error) will trigger JAXB to stop. There are the following three levels: 0=warning, 1=error, 2=fatal error.	0	Integer
<code>camel.dataformat.jaxb.xml-stream-writer-wrapper</code>	To use a custom xml stream writer.		String

CHAPTER 61. JSON GSON

Gson is a Data Format which uses the [Gson Library](#).

```
from("activemq:My.Queue").
  marshal().json(JsonLibrary.Gson).
  to("mqseries:Another.Queue");
```

61.1. GSON OPTIONS

The JSON Gson dataformat supports 3 options, which are listed below.

Name	Default	Java Type	Description
prettyPrint		Boolean	To enable pretty printing output nicely formatted. Is by default false.
unmarshalType		String	Class name of the java type to use when unmarshalling.
contentTypeHeader		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.

61.2. DEPENDENCIES

To use Gson in your camel routes you need to add the dependency on **camel-gson** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-gson</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

61.3. SPRING BOOT AUTO-CONFIGURATION

When using json-gson with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-gson-starter</artifactId>
```

```

<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 4 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.json-gson.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.json-gson.enabled</code>	Whether to enable auto configuration of the json-gson data format. This is enabled by default.		Boolean
<code>camel.dataformat.json-gson.pretty-print</code>	To enable pretty printing output nicely formatted. Is by default false.	false	Boolean
<code>camel.dataformat.json-gson.unmarshal-type</code>	Class name of the java type to use when unmarshalling.		String

CHAPTER 62. JSON JACKSON

Jackson is a Data Format which uses the [Jackson Library](#)

```
from("activemq:My.Queue").
  marshal().json(JsonLibrary.Jackson).
  to("mqseries:Another.Queue");
```

62.1. JACKSON OPTIONS

The JSON Jackson dataformat supports 20 options, which are listed below.

Name	Default	Java Type	Description
<code>objectMapper</code>		String	Lookup and use the existing ObjectMapper with the given id when using Jackson.
<code>useDefaultObjectMapper</code>		Boolean	Whether to lookup and use default Jackson ObjectMapper from the registry.
<code>prettyPrint</code>		Boolean	To enable pretty printing output nicely formatted. Is by default false.
<code>unmarshalType</code>		String	Class name of the java type to use when unmarshalling.

Name	Default	Java Type	Description
<code>jsonView</code>		String	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.
<code>include</code>		String	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .
<code>allowJmsType</code>		Boolean	Used for JMS users to allow the <code>JMSType</code> header from the JMS spec to specify a FQN classname to use to unmarshal to.
<code>collectionType</code>		String	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than <code>java.util.Collection</code> based as default.

Name	Default	Java Type	Description
<code>useList</code>		Boolean	To unmarshal to a List of Map or a List of Pojo.
<code>moduleClassNames</code>		String	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.
<code>moduleRefs</code>		String	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.
<code>enableFeatures</code>		String	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.

Name	Default	Java Type	Description
disableFeatures		String	Set of features to disable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.
allowUnmarshalType		Boolean	If enabled then Jackson is allowed to attempt to use the <code>CamelJacksonUnmarshalType</code> header during the unmarshalling. This should only be enabled when desired to be used.
timezone		String	If set then Jackson will use the <code>Timezone</code> when marshalling/unmarshalling. This option will have no effect on the others <code>JsonDataFormat</code> , like <code>gson</code> , <code>fastjson</code> and <code>xstream</code> .

Name	Default	Java Type	Description
<code>autoDiscoverObjectMapper</code>		Boolean	If set to true then Jackson will lookup for an objectMapper into the registry.
<code>contentTypeHeader</code>		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.
<code>schemaResolver</code>		String	Optional schema resolver used to lookup schemas for the data in transit.
<code>autoDiscoverSchemaResolver</code>		Boolean	When not disabled, the SchemaResolver will be looked up into the registry.

Name	Default	Java Type	Description
namingStrategy		String	If set then Jackson will use the the defined Property Naming Strategy. Possible values are: LOWER_CAMEL_CASE, LOWER_DOT_CASE, LOWER_CASE, KEBAB_CASE, SNAKE_CASE and UPPER_CAMEL_CASE.

62.2. USING CUSTOM OBJECTMAPPER

You can configure **JacksonDataFormat** to use a custom **ObjectMapper** in case you need more control of the mapping configuration.

If you setup a single **ObjectMapper** in the registry, then Camel will automatic lookup and use this **ObjectMapper**. For example if you use Spring Boot, then Spring Boot can provide a default **ObjectMapper** for you if you have Spring MVC enabled. And this would allow Camel to detect that there is one bean of **ObjectMapper** class type in the Spring Boot bean registry and then use it. When this happens you should set a **INFO** logging from Camel.

62.3. USING JACKSON FOR AUTOMATIC TYPE CONVERSION

The **camel-jackson** module allows integrating Jackson as a [Type Converter](#). This works in a similar way to [JAXB](#) that integrates with Camel's type converter.

To use this **camel-jackson** must be enabled, which is done by setting the following options on the **CamelContext** global options, as shown:

```
@Bean
CamelContextConfiguration contextConfiguration() {
    return new CamelContextConfiguration() {
        @Override
        public void beforeApplicationStart(CamelContext context) {
            // Enable Jackson JSON type converter.
            context.getGlobalOptions().put(JacksonConstants.ENABLE_TYPE_CONVERTER, "true");
            // Allow Jackson JSON to convert to pojo types also
            // (by default Jackson only converts to String and other simple types)
            getContext().getGlobalOptions().put(JacksonConstants.TYPE_CONVERTER_TO_POJO,
"true");
        }
    }
}
```

```

@Override
public void afterApplicationStart(CamelContext camelContext) {

}
};
}

```

The **camel-jackson** type converter integrates with [JAXB](#) which means you can annotate POJO class with **JAXB** annotations that Jackson can use. You can also use Jackson's own annotations on your POJO classes.

62.4. DEPENDENCIES

To use Jackson in your camel routes you need to add the dependency on **camel-jackson** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jackson</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>

```

62.5. SPRING BOOT AUTO-CONFIGURATION

When using json-jackson with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jackson-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 21 options, which are listed below.

Name	Description	Default	Type
camel.dataformat.json-jackson.allow-jms-type	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.	false	Boolean
camel.dataformat.json-jackson.allow-unmarshall-type	If enabled then Jackson is allowed to attempt to use the CamelJacksonUnmarshalType header during the unmarshalling. This should only be enabled when desired to be used.	false	Boolean

Name	Description	Default	Type
<code>camel.dataformat.json-jackson.auto-discover-object-mapper</code>	If set to true then Jackson will lookup for an objectMapper into the registry.	false	Boolean
<code>camel.dataformat.json-jackson.auto-discover-schema-resolver</code>	When not disabled, the SchemaResolver will be looked up into the registry.	true	Boolean
<code>camel.dataformat.json-jackson.collection-type</code>	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.		String
<code>camel.dataformat.json-jackson.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example application/xml for data formats marshalling to XML, or application/json for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.json-jackson.disable-features</code>	Set of features to disable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.json-jackson.enable-features</code>	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.json-jackson.enabled</code>	Whether to enable auto configuration of the json-jackson data format. This is enabled by default.		Boolean
<code>camel.dataformat.json-jackson.include</code>	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .		String

Name	Description	Default	Type
<code>camel.dataformat.json-jackson.json-view</code>	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.		String
<code>camel.dataformat.json-jackson.module-class-names</code>	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.		String
<code>camel.dataformat.json-jackson.module-refs</code>	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.		String
<code>camel.dataformat.json-jackson.naming-strategy</code>	If set then Jackson will use the the defined Property Naming Strategy. Possible values are: <code>LOWER_CAMEL_CASE</code> , <code>LOWER_DOT_CASE</code> , <code>LOWER_CASE</code> , <code>KEBAB_CASE</code> , <code>SNAKE_CASE</code> and <code>UPPER_CAMEL_CASE</code> .		String
<code>camel.dataformat.json-jackson.object-mapper</code>	Lookup and use the existing <code>ObjectMapper</code> with the given id when using Jackson.		String
<code>camel.dataformat.json-jackson.pretty-print</code>	To enable pretty printing output nicely formatted. Is by default false.	false	Boolean
<code>camel.dataformat.json-jackson.schema-resolver</code>	Optional schema resolver used to lookup schemas for the data in transit.		String
<code>camel.dataformat.json-jackson.timezone</code>	If set then Jackson will use the <code>Timezone</code> when marshalling/unmarshalling. This option will have no effect on the others <code>Json DataFormat</code> , like <code>gson</code> , <code>fastjson</code> and <code>xstream</code> .		String
<code>camel.dataformat.json-jackson.unmarshal-type</code>	Class name of the java type to use when unmarshalling.		String

Name	Description	Default	Type
<code>camel.dataformat.json-jackson.use-default-object-mapper</code>	Whether to lookup and use default Jackson ObjectMapper from the registry.	true	Boolean
<code>camel.dataformat.json-jackson.use-list</code>	To unmarshal to a List of Map or a List of Pojo.	false	Boolean

CHAPTER 63. PROTOBUF JACKSON

Jackson Protobuf is a Data Format which uses the [Jackson library](#) with the [Protobuf extension](#) to unmarshal a Protobuf payload into Java objects or to marshal Java objects into a Protobuf payload.



NOTE

If you are familiar with Jackson, this Protobuf data format behaves in the same way as its JSON counterpart, and thus can be used with classes annotated for JSON serialization/deserialization.

```
from("kafka:topic").
  unmarshal().protobuf(ProtobufLibrary.Jackson, JsonNode.class).
  to("log:info");
```

63.1. CONFIGURING THE SCHEMARESOLVER

Since Protobuf serialization is schema-based, this data format requires that you provide a SchemaResolver object that is able to lookup the schema for each exchange that is going to be marshalled/unmarshalled.

You can add a single SchemaResolver to the registry and it will be looked up automatically. Or you can explicitly specify the reference to a custom SchemaResolver.

63.2. PROTOBUF JACKSON OPTIONS

The Protobuf Jackson dataformat supports 18 options, which are listed below.

Name	Default	Java Type	Description
<code>contentTypeHeader</code>		Boolean	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.
<code>objectMapper</code>		String	Lookup and use the existing ObjectMapper with the given id when using Jackson.
<code>useDefaultObjectMapper</code>		Boolean	Whether to lookup and use default Jackson ObjectMapper from the registry.
<code>unmarshalType</code>		String	Class name of the java type to use when unmarshalling.

Name	Default	Java Type	Description
<code>jsonView</code>		String	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.
<code>include</code>		String	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .
<code>allowJmsType</code>		Boolean	Used for JMS users to allow the <code>JMSType</code> header from the JMS spec to specify a FQN classname to use to unmarshal to.
<code>collectionType</code>		String	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than <code>java.util.Collection</code> based as default.
<code>useList</code>		Boolean	To unmarshal to a List of Map or a List of Pojo.
<code>moduleClassNames</code>		String	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.
<code>moduleRefs</code>		String	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.
<code>enableFeatures</code>		String	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches an enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.

Name	Default	Java Type	Description
<code>disableFeatures</code>		String	Set of features to disable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches an enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.
<code>allowUnmarshalType</code>		Boolean	If enabled then Jackson is allowed to attempt to use the <code>CamelJacksonUnmarshalType</code> header during the unmarshalling. This should only be enabled when desired to be used.
<code>timezone</code>		String	If set then Jackson will use the Timezone when marshalling/unmarshalling.
<code>autoDiscoverObjectMapper</code>		Boolean	If set to true then Jackson will lookup for an <code>ObjectMapper</code> into the registry.
<code>schemaResolver</code>		String	Optional schema resolver used to lookup schemas for the data in transit.
<code>autoDiscoverSchemaResolver</code>		Boolean	When not disabled, the <code>SchemaResolver</code> will be looked up into the registry.

63.3. USING CUSTOM PROTOBUFMAPPER

You can configure **`JacksonProtobufDataFormat`** to use a custom **`ProtobufMapper`** in case you need more control of the mapping configuration.

If you setup a single **`ProtobufMapper`** in the registry, then Camel will automatic lookup and use this **`ProtobufMapper`**.

63.4. DEPENDENCIES

To use Protobuf Jackson in your camel routes you need to add the dependency on **`camel-jackson-protobuf`** which implements this data format.

If you use maven, add the following to your **`pom.xml`**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-jackson-protobuf</artifactId>
```

```

<version>3.14.5.redhat-00018</version>
<!-- use the same version as your Camel core version -->
</dependency>

```

63.5. SPRING BOOT AUTO-CONFIGURATION

When using protobuf-jackson with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
<groupId>org.apache.camel.springboot</groupId>
<artifactId>camel-jackson-protobuf-starter</artifactId>
<version>3.14.5.redhat-00032</version>
<!-- use the same version as your Camel core version -->
</dependency>

```

The component supports 19 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.protobuf-jackson.allow-jms-type</code>	Used for JMS users to allow the JMSType header from the JMS spec to specify a FQN classname to use to unmarshal to.	false	Boolean
<code>camel.dataformat.protobuf-jackson.allow-unmarshall-type</code>	If enabled then Jackson is allowed to attempt to use the CamelJacksonUnmarshalType header during the unmarshalling. This should only be enabled when desired to be used.	false	Boolean
<code>camel.dataformat.protobuf-jackson.auto-discover-object-mapper</code>	If set to true then Jackson will lookup for an objectMapper into the registry.	false	Boolean
<code>camel.dataformat.protobuf-jackson.auto-discover-schema-resolver</code>	When not disabled, the SchemaResolver will be looked up into the registry.	true	Boolean
<code>camel.dataformat.protobuf-jackson.collection-type</code>	Refers to a custom collection type to lookup in the registry to use. This option should rarely be used, but allows to use different collection types than java.util.Collection based as default.		String

Name	Description	Default	Type
<code>camel.dataformat.protobuf-jackson.content-type-header</code>	Whether the data format should set the Content-Type header with the type from the data format. For example <code>application/xml</code> for data formats marshalling to XML, or <code>application/json</code> for data formats marshalling to JSON.	true	Boolean
<code>camel.dataformat.protobuf-jackson.disable-features</code>	Set of features to disable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.protobuf-jackson.enable-features</code>	Set of features to enable on the Jackson <code>com.fasterxml.jackson.databind.ObjectMapper</code> . The features should be a name that matches a enum from <code>com.fasterxml.jackson.databind.SerializationFeature</code> , <code>com.fasterxml.jackson.databind.DeserializationFeature</code> , or <code>com.fasterxml.jackson.databind.MapperFeature</code> . Multiple features can be separated by comma.		String
<code>camel.dataformat.protobuf-jackson.enabled</code>	Whether to enable auto configuration of the <code>protobuf-jackson</code> data format. This is enabled by default.		Boolean
<code>camel.dataformat.protobuf-jackson.include</code>	If you want to marshal a pojo to JSON, and the pojo has some fields with null values. And you want to skip these null values, you can set this option to <code>NON_NULL</code> .		String
<code>camel.dataformat.protobuf-jackson.json-view</code>	When marshalling a POJO to JSON you might want to exclude certain fields from the JSON output. With Jackson you can use JSON views to accomplish this. This option is to refer to the class which has <code>JsonView</code> annotations.		String
<code>camel.dataformat.protobuf-jackson.module-class-names</code>	To use custom Jackson modules <code>com.fasterxml.jackson.databind.Module</code> specified as a String with FQN class names. Multiple classes can be separated by comma.		String
<code>camel.dataformat.protobuf-jackson.module-refs</code>	To use custom Jackson modules referred from the Camel registry. Multiple modules can be separated by comma.		String

Name	Description	Default	Type
<code>camel.dataformat.protobuf-jackson.object-mapper</code>	Lookup and use the existing ObjectMapper with the given id when using Jackson.		String
<code>camel.dataformat.protobuf-jackson.schema-resolver</code>	Optional schema resolver used to lookup schemas for the data in transit.		String
<code>camel.dataformat.protobuf-jackson.timezone</code>	If set then Jackson will use the Timezone when marshalling/unmarshalling.		String
<code>camel.dataformat.protobuf-jackson.unmarshal-type</code>	Class name of the java type to use when unmarshalling.		String
<code>camel.dataformat.protobuf-jackson.use-default-object-mapper</code>	Whether to lookup and use default Jackson ObjectMapper from the registry.	true	Boolean
<code>camel.dataformat.protobuf-jackson.use-list</code>	To unmarshal to a List of Map or a List of Pojo.	false	Boolean

CHAPTER 64. SOAP

SOAP is a Data Format which uses JAXB2 and JAX-WS annotations to marshal and unmarshal SOAP payloads. It provides the basic features of Apache CXF without need for the CXF Stack.

Namespace prefix mapping

See [JAXB](#) for details how you can control namespace prefix mappings when marshalling using SOAP data format.

64.1. SOAP OPTIONS

The SOAP dataformat supports 6 options, which are listed below.

Name	Default	Java Type	Description
<code>contextPath</code>		String	Required Package name where your JAXB classes are located.
<code>encoding</code>		String	To overrule and use a specific encoding.
<code>elementNameStrategyRef</code>		String	Refers to an element strategy to lookup from the registry. An element name strategy is used for two purposes. The first is to find a xml element name for a given object and soap action when marshaling the object into a SOAP message. The second is to find an Exception class for a given soap fault name. The following three element strategy class name is provided out of the box. <code>QNameStrategy</code> - Uses a fixed qName that is configured on instantiation. Exception lookup is not supported <code>TypeNameStrategy</code> - Uses the name and namespace from the <code>XMLType</code> annotation of the given type. If no namespace is set then package-info is used. Exception lookup is not supported <code>ServiceInterfaceStrategy</code> - Uses information from a webservice interface to determine the type name and to find the exception class for a SOAP fault All three classes is located in the package name <code>org.apache.camel.dataformat.soap.name</code> If you have generated the web service stub code with <code>cxf-codegen</code> or a similar tool then you probably will want to use the <code>ServiceInterfaceStrategy</code> . In the case you have no annotated service interface you should use <code>QNameStrategy</code> or <code>TypeNameStrategy</code> .
<code>version</code>		String	SOAP version should either be 1.1 or 1.2. Is by default 1.1.
<code>namespacePrefixRef</code>		String	When marshalling using JAXB or SOAP then the JAXB implementation will automatic assign namespace prefixes, such as <code>ns2</code> , <code>ns3</code> , <code>ns4</code> etc. To control this mapping, Camel allows you to refer to a map which contains the desired mapping.

Name	Default	Java Type	Description
schema		String	To validate against an existing schema. You can use the prefix <code>classpath:</code> , <code>file:</code> or <code>http:</code> to specify how the resource should be resolved. You can separate multiple schema files by using the <code>'</code> character.

64.2. ELEMENTNAMESTRATEGY

An element name strategy is used for two purposes. The first is to find a xml element name for a given object and soap action when marshaling the object into a SOAP message. The second is to find an Exception class for a given soap fault name.

Strategy	Usage
QNameStrategy	Uses a fixed qName that is configured on instantiation. Exception lookup is not supported
TypeNameStrategy	Uses the name and namespace from the <code>@XMLType</code> annotation of the given type. If no namespace is set then <code>package-info</code> is used. Exception lookup is not supported
ServiceInterfaceStrategy	Uses information from a webservice interface to determine the type name and to find the exception class for a SOAP fault

If you have generated the web service stub code with `cxf-codegen` or a similar tool then you probably will want to use the `ServiceInterfaceStrategy`. In the case you have no annotated service interface you should use `QNameStrategy` or `TypeNameStrategy`.

64.3. USING THE JAVA DSL

The following example uses a named `DataFormat` of `soap` which is configured with the package `com.example.customerservice` to initialize the `JAXBContext`. The second parameter is the `ElementNameStrategy`. The route is able to marshal normal objects as well as exceptions. (Note the below just sends a SOAP Envelope to a queue. A web service provider would actually need to be listening to the queue for a SOAP call to actually occur, in which case it would be a one way SOAP request. If you need request reply then you should look at the next example.)

```
SoapJaxbDataFormat soap = new SoapJaxbDataFormat("com.example.customerservice", new
ServiceInterfaceStrategy(CustomerService.class));
from("direct:start")
.marshall(soap)
.to("jms:myQueue");
```



NOTE

See also

As the SOAP dataformat inherits from the JAXB dataformat most settings apply here as well.

64.3.1. Using SOAP 1.2

Since Camel 2.11

```
SoapJaxbDataFormat soap = new SoapJaxbDataFormat("com.example.customerservice", new
ServiceInterfaceStrategy(CustomerService.class));
soap.setVersion("1.2");
from("direct:start")
.marshall(soap)
.to("jms:myQueue");
```

When using XML DSL there is a version attribute you can set on the <soapjaxb> element.

```
<!-- Defining a ServiceInterfaceStrategy for retrieving the element name when marshalling -->
<bean id="myNameStrategy"
class="org.apache.camel.dataformat.soap.name.ServiceInterfaceStrategy">
  <constructor-arg value="com.example.customerservice.CustomerService"/>
  <constructor-arg value="true"/>
</bean>
```

And in the Camel route

```
<route>
  <from uri="direct:start"/>
  <marshal>
    <soapjaxb contentPath="com.example.customerservice" version="1.2"
elementNameStrategyRef="myNameStrategy"/>
  </marshal>
  <to uri="jms:myQueue"/>
</route>
```

64.4. MULTI-PART MESSAGES

Multi-part SOAP messages are supported by the ServiceInterfaceStrategy. The ServiceInterfaceStrategy must be initialized with a service interface definition that is annotated in accordance with JAX-WS 2.2 and meets the requirements of the Document Bare style. The target method must meet the following criteria, as per the JAX-WS specification: 1) it must have at most one **in** or **in/out** non-header parameter, 2) if it has a return type other than **void** it must have no **in/out** or **out** non-header parameters, 3) if it has a return type of **void** it must have at most one **in/out** or **out** non-header parameter.

The ServiceInterfaceStrategy should be initialized with a boolean parameter that indicates whether the mapping strategy applies to the request parameters or response parameters.

```
ServiceInterfaceStrategy strat = new
ServiceInterfaceStrategy(com.example.customerservice.multipart.MultiPartCustomerService.class,
true);
```



```
SoapJaxbDataFormat soapDataFormat = new
SoapJaxbDataFormat("com.example.customerservice.multipart", strat);
```

64.4.1. Holder Object mapping

JAX-WS specifies the use of a type-parameterized **javax.xml.ws.Holder** object for **In/Out** and **Out** parameters. You may use an instance of the parameterized-type directly. The camel-soap DataFormat marshals Holder values in accordance with the JAXB mapping for the class of the **Holder's value. No mapping is provided for Holder** objects in an unmarshalled response.

64.5. EXAMPLES

64.5.1. Webservice client

The following route supports marshalling the request and unmarshalling a response or fault.

```
String WS_URI = "cxf://http://myserver/customerservice?
serviceClass=com.example.customerservice&dataFormat=RAW";
SoapJaxbDataFormat soapDF = new SoapJaxbDataFormat("com.example.customerservice", new
ServiceInterfaceStrategy(CustomerService.class));
from("direct:customerServiceClient")
    .onException(Exception.class)
    .handled(true)
    .unmarshal(soapDF)
    .end()
    .marshal(soapDF)
    .to(WS_URI)
    .unmarshal(soapDF);
```

The below snippet creates a proxy for the service interface and makes a SOAP call to the above route.

```
import org.apache.camel.Endpoint;
import org.apache.camel.component.bean.ProxyHelper;
...

Endpoint startEndpoint = context.getEndpoint("direct:customerServiceClient");
ClassLoader classLoader = Thread.currentThread().getContextClassLoader();
// CustomerService below is the service endpoint interface, *not* the javax.xml.ws.Service subclass
CustomerService proxy = ProxyHelper.createProxy(startEndpoint, classLoader,
CustomerService.class);
GetCustomersByNameResponse response = proxy.getCustomersByName(new
GetCustomersByName());
```

64.5.2. Webservice Server

Using the following route sets up a webservice server that listens on jms queue customerServiceQueue and processes requests using the class CustomerServiceImpl. The customerServiceImpl of course should implement the interface CustomerService. Instead of directly instantiating the server class it could be defined in a spring context as a regular bean.

```
SoapJaxbDataFormat soapDF = new SoapJaxbDataFormat("com.example.customerservice", new
ServiceInterfaceStrategy(CustomerService.class));
CustomerService serverBean = new CustomerServiceImpl();
```

```

from("jms://queue:customerServiceQueue")
  .onException(Exception.class)
  .handled(true)
  .marshal(soapDF)
  .end()
  .unmarshal(soapDF)
  .bean(serverBean)
  .marshal(soapDF);

```

64.6. DEPENDENCIES

To use the SOAP dataformat in your camel routes you need to add the following dependency to your pom.

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-soap</artifactId>
  <version>3.14.5.redhat-00018</version>
</dependency>

```

64.7. SPRING BOOT AUTO-CONFIGURATION

When using soapjaxb with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-soap-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 7 options, which are listed below.

Name	Description	Default	Type
camel.dataformat .soapjaxb.context -path	Package name where your JAXB classes are located.		String

Name	Description	Default	Type
camel.dataformat.soapjaxb.element-name-strategy-ref	Refers to an element strategy to lookup from the registry. An element name strategy is used for two purposes. The first is to find a xml element name for a given object and soap action when marshaling the object into a SOAP message. The second is to find an Exception class for a given soap fault name. The following three element strategy class name is provided out of the box. QNameStrategy - Uses a fixed qName that is configured on instantiation. Exception lookup is not supported TypeNameStrategy - Uses the name and namespace from the XMLType annotation of the given type. If no namespace is set then package-info is used. Exception lookup is not supported ServiceInterfaceStrategy - Uses information from a webservice interface to determine the type name and to find the exception class for a SOAP fault All three classes is located in the package name org.apache.camel.dataformat.soap.name If you have generated the web service stub code with cxf-codegen or a similar tool then you probably will want to use the ServiceInterfaceStrategy. In the case you have no annotated service interface you should use QNameStrategy or TypeNameStrategy.		String
camel.dataformat.soapjaxb.enabled	Whether to enable auto configuration of the soapjaxb data format. This is enabled by default.		Boolean
camel.dataformat.soapjaxb.encoding	To overrule and use a specific encoding.		String
camel.dataformat.soapjaxb.namespace-prefix-ref	When marshalling using JAXB or SOAP then the JAXB implementation will automatic assign namespace prefixes, such as ns2, ns3, ns4 etc. To control this mapping, Camel allows you to refer to a map which contains the desired mapping.		String
camel.dataformat.soapjaxb.schema	To validate against an existing schema. Your can use the prefix classpath;, file: or http: to specify how the resource should by resolved. You can separate multiple schema files by using the ',' character.		String
camel.dataformat.soapjaxb.version	SOAP version should either be 1.1 or 1.2. Is by default 1.1.	1.1	String

CHAPTER 65. ZIP FILE

The Zip File Data Format is a message compression and de-compression format. Messages can be marshalled (compressed) to Zip files containing a single entry, and Zip files containing a single entry can be unmarshalled (decompressed) to the original file contents. This data format supports ZIP64, as long as Java 7 or later is being used].

65.1. ZIPFILE OPTIONS

The Zip File dataformat supports 4 options, which are listed below.

Name	Default	Java Type	Description
<code>usingIterator</code>		Boolean	If the zip file has more then one entry, the setting this option to true, allows to work with the splitter EIP, to split the data using an iterator in a streaming mode.
<code>allowEmptyDirectory</code>		Boolean	If the zip file has more then one entry, setting this option to true, allows to get the iterator even if the directory is empty.
<code>preservePathElements</code>		Boolean	If the file name contains path elements, setting this option to true, allows the path to be maintained in the zip file.
<code>maxDecompressedSize</code>		Integer	Set the maximum decompressed size of a zip file (in bytes). The default value if not specified corresponds to 1 gigabyte. An IOException will be thrown if the decompressed size exceeds this amount. Set to -1 to disable setting a maximum decompressed size.

65.2. MARSHAL

In this example we marshal a regular text/XML payload to a compressed payload using Zip file compression, and send it to an ActiveMQ queue called MY_QUEUE.

```
from("direct:start")
  .marshal().zipFile()
  .to("activemq:queue:MY_QUEUE");
```

The name of the Zip entry inside the created Zip file is based on the incoming **CamelFileName** message header, which is the standard message header used by the file component. Additionally, the outgoing **CamelFileName** message header is automatically set to the value of the incoming **CamelFileName** message header, with the ".zip" suffix. So for example, if the following route finds a file named "test.txt" in the input directory, the output will be a Zip file named "test.txt.zip" containing a single Zip entry named "test.txt":

```
from("file:input/directory?antInclude=/*.txt")
  .marshal().zipFile()
  .to("file:output/directory");
```

If there is no incoming **CamelFileName** message header (for example, if the file component is not the

consumer), then the message ID is used by default, and since the message ID is normally a unique generated ID, you will end up with filenames like **ID-MACHINENAME-2443-1211718892437-1-0.zip**. If you want to override this behavior, then you can set the value of the **CamelFileName** header explicitly in your route:

```
from("direct:start")
    .setHeader(Exchange.FILE_NAME, constant("report.txt"))
    .marshal().zipFile()
    .to("file:output/directory");
```

This route would result in a Zip file named "report.txt.zip" in the output directory, containing a single Zip entry named "report.txt".

65.3. UNMARSHAL

In this example we unmarshal a Zip file payload from an ActiveMQ queue called MY_QUEUE to its original format, and forward it for processing to the **UnZippedMessageProcessor**.

```
from("activemq:queue:MY_QUEUE")
    .unmarshal().zipFile()
    .process(new UnZippedMessageProcessor());
```

If the zip file has more than one entry, the `usingIterator` option of `ZipFileDataFormat` to be true, and you can use `splitter` to do the further work.

```
ZipFileDataFormat zipFile = new ZipFileDataFormat();
zipFile.setUsingIterator(true);

from("file:src/test/resources/org/apache/camel/dataformat/zipfile/?delay=1000&noop=true")
    .unmarshal(zipFile)
    .split(body(Iterator.class)).streaming()
    .process(new UnZippedMessageProcessor())
    .end();
```

Or you can use the `ZipSplitter` as an expression for `splitter` directly like this

```
from("file:src/test/resources/org/apache/camel/dataformat/zipfile?delay=1000&noop=true")
    .split(new ZipSplitter()).streaming()
    .process(new UnZippedMessageProcessor())
    .end();
```

65.3.1. Aggregate



NOTE

This aggregation strategy requires eager completion check to work properly.

In this example we aggregate all text files found in the input directory into a single Zip file that is stored in the output directory.

```
from("file:input/directory?antInclude=/*.txt")
    .aggregate(constant(true), new ZipAggregationStrategy());
```

```
.completionFromBatchConsumer().eagerCheckCompletion()
.to("file:output/directory");
```

The outgoing **CamelFileName** message header is created using `java.io.File.createTempFile`, with the ".zip" suffix. If you want to override this behavior, then you can set the value of the **CamelFileName** header explicitly in your route:

```
from("file:input/directory?antInclude=*.txt")
  .aggregate(constant(true), new ZipAggregationStrategy())
  .completionFromBatchConsumer().eagerCheckCompletion()
  .setHeader(Exchange.FILE_NAME, constant("reports.zip"))
  .to("file:output/directory");
```

65.4. DEPENDENCIES

To use Zip files in your camel routes you need to add a dependency on **camel-zipfile** which implements this data format.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-zipfile</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

65.5. SPRING BOOT AUTO-CONFIGURATION

When using zipfile with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-zipfile-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```


The component supports 5 options, which are listed below.

Name	Description	Default	Type
<code>camel.dataformat.zipfile.allow-empty-directory</code>	If the zip file has more than one entry, setting this option to true, allows to get the iterator even if the directory is empty.	false	Boolean
<code>camel.dataformat.zipfile.enabled</code>	Whether to enable auto configuration of the zipfile data format. This is enabled by default.		Boolean

Name	Description	Default	Type
camel.dataformat.zipfile.max-decompressed-size	Set the maximum decompressed size of a zip file (in bytes). The default value if not specified corresponds to 1 gigabyte. An IOException will be thrown if the decompressed size exceeds this amount. Set to -1 to disable setting a maximum decompressed size.	1073741824	Long
camel.dataformat.zipfile.preserve-path-elements	If the file name contains path elements, setting this option to true, allows the path to be maintained in the zip file.	false	Boolean
camel.dataformat.zipfile.using-iterator	If the zip file has more than one entry, the setting this option to true, allows to work with the splitter EIP, to split the data using an iterator in a streaming mode.	false	Boolean

CHAPTER 66. CONSTANT

The Constant Expression Language is really just a way to use a constant value or object.



NOTE

This is a fixed constant value (or object) that is only set once during starting up the route, do not use this if you want dynamic values during routing.

66.1. CONSTANT OPTIONS

The Constant language supports 2 options, which are listed below.

Name	Default	Java Type	Description
<code>resultType</code>		String	Sets the class name of the constant type.
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

66.2. EXAMPLE

The **setHeader** EIP can utilize a constant expression like:

```
<route>
  <from uri="seda:a"/>
  <setHeader name="theHeader">
    <constant>the value</constant>
  </setHeader>
  <to uri="mock:b"/>
</route>
```

in this case, the message coming from the `seda:a` endpoint will have the header with key **theHeader** set its value as **the value** (string type).

And the same example using Java DSL:

```
from("seda:a")
  .setHeader("theHeader", constant("the value"))
  .to("mock:b");
```

66.2.1. Specifying type of value

The option **resultType** can be used to specify the type of the value, when the value is given as a **String** value, which happens when using XML or YAML DSL:

For example to set a header with **int** type you can do:

```
<route>
  <from uri="seda:a"/>
```



```

<setHeader name="zipCode">
  <constant resultType="int">90210</constant>
</setHeader>
<to uri="mock:b"/>
</route>

```

66.3. LOADING CONSTANT FROM EXTERNAL RESOURCE

You can externalize the constant and have Camel load it from a resource such as **"classpath:"**, **"file:"**, or **"http:"**.

This is done using the following syntax: **"resource:scheme:location"**, eg to refer to a file on the classpath you can do:

```

.setHeader("myHeader").constant("resource:classpath:constant.txt")

```

66.4. DEPENDENCIES

The Constant language is part of **camel-core**.

66.5. SPRING BOOT AUTO-CONFIGURATION

When using constant with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use the same version as your Camel core version -->
</dependency>

```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.user-name</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as name.namespace.svc.dnsDomain. When using dnssrv the service name is resolved with SRV query forsvc... When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable KUBERNETES_MASTER.		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DN SSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DN SSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String

Name	Description	Default	Type
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean

Name	Description	Default	Type
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map
<code>camel.hystrix.core-pool-size</code>	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
<code>camel.hystrix.enabled</code>	Enable the component.	true	Boolean
<code>camel.hystrix.execution-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer

Name	Description	Default	Type
camel.hystrix.execution-isolation-strategy	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	<code>THREAD</code>	String
camel.hystrix.execution-isolation-thread-interrupt-on-timeout	Whether the execution thread should attempt an interrupt (using <code>Future#cancel()</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	<code>true</code>	Boolean
camel.hystrix.execution-timeout-enabled	Whether the timeout mechanism is enabled for this command.	<code>true</code>	Boolean
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	<code>1000</code>	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	<code>true</code>	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	<code>10</code>	Integer
camel.hystrix.group-key	Sets the group key to use. The default value is <code>CamelHystrix</code> .	<code>CamelHystrix</code>	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long, TimeUnit)</code> .	<code>1</code>	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> . This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	<code>-1</code>	Integer

Name	Description	Default	Type
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer
camel.hystrix.metrics-rolling-percentile-window-in-milliseconds	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
camel.hystrix.metrics-rolling-statistical-window-buckets	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-statistical-window-in-milliseconds	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer

Name	Description	Default	Type
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.csimple.enabled</code>	Whether to enable auto configuration of the <code>csimple</code> language. This is enabled by default.		Boolean
<code>camel.language.csimple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the <code>exchangeProperty</code> language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean

Name	Description	Default	Type
camel.resilience4j.circuit-breaker-ref	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
camel.resilience4j.config-ref	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
camel.resilience4j.configurations	Define additional configuration definitions.		Map
camel.resilience4j.enabled	Enable the component.	true	Boolean
camel.resilience4j.failure-rate-threshold	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
camel.resilience4j.minimum-number-of-calls	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer
camel.resilience4j.permitted-number-of-calls-in-half-open-state	Configures the number of permitted calls when the <code>CircuitBreaker</code> is half open. The size must be greater than 0. Default size is 10.	10	Integer

Name	Description	Default	Type
camel.resilience4j.sliding-window-size	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
camel.resilience4j.sliding-window-type	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	COUNT_BASED	String
camel.resilience4j.slow-call-duration-threshold	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
camel.resilience4j.slow-call-rate-threshold	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float
camel.resilience4j.wait-duration-in-open-state	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer

Name	Description	Default	Type
<code>camel.resilience4j.writable-stack-trace-enabled</code>	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
<code>camel.rest.api-component</code>	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.rest.api-context-path</code>	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
<code>camel.rest.api-context-route-id</code>	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
<code>camel.rest.api-host</code>	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
<code>camel.rest.api-property</code>	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
<code>camel.rest.api-vendor-extension</code>	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean
<code>camel.rest.binding-mode</code>	Sets the binding mode to use. The default value is off.		RestBindingMode

Name	Description	Default	Type
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBElement is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map

Name	Description	Default	Type
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.host-name-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
camel.rest.skip-binding-on-error-code	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean
camel.rest.use-x-forward-headers	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean

Name	Description	Default	Type
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default JAXB will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 67. CSIMPLE

The CSimple language is **compiled** Simple language.

67.1. DIFFERENT BETWEEN CSIMPLE AND SIMPLE

The simple language is a dynamic expression language which is runtime parsed into a set of Camel Expressions or Predicates.

The csimple language is parsed into regular Java source code and compiled together with all the other source code, or compiled once during bootstrap via the **camel-csimple-joor** module.

The simple language is generally very lightweight and fast, however for some use-cases with dynamic method calls via OGNL paths, then the simple language does runtime introspection and reflection calls. This has an overhead on performance, and was one of the reasons why csimple was created.

The csimple language requires to be typesafe and method calls via OGNL paths requires to know the type during parsing. This means for csimple languages expressions you would need to provide the class type in the script, whereas simple introspects this at runtime.

In other words the simple language is using *duck typing* (if it looks like a duck, and quacks like a duck, then it is a duck) and csimple is using Java type (typesafety). If there is a type error then simple will report this at runtime, and with csimple there will be a Java compilation error.

67.1.1. Additional CSimple functions

The csimple language includes some additional functions to support common use-cases working with **Collection**, **Map** or array types. The following functions *bodyAsIndex*, *headerAsIndex*, and *exchangePropertyAsIndex* is used for these use-cases as they are typed.

Function	Type	Description
<code>bodyAsIndex(type, index)</code>	Type	To be used for collecting the body from an existing Collection , Map or array (lookup by the index) and then converting the body to the given type determined by its classname. The converted body can be null.
<code>mandatoryBodyAsIndex(type, index)</code>	Type	To be used for collecting the body from an existing Collection , Map or array (lookup by the index) and then converting the body to the given type determined by its classname. Expects the body to be not null.
<code>headerAsIndex(key, type, index)</code>	Type	To be used for collecting a header from an existing Collection , Map or array (lookup by the index) and then converting the header value to the given type determined by its classname. The converted header can be null.
<code>mandatoryHeaderAsIndex(key, type, index)</code>	Type	To be used for collecting a header from an existing Collection , Map or array (lookup by the index) and then converting the header value to the given type determined by its classname. Expects the header to be not null.

Function	Type	Description
<code>exchangePropertyAsIndex(key, type, index)</code>	Type	To be used for collecting an exchange property from an existing Collection, Map or array (lookup by the index) and then converting the exchange property to the given type determined by its classname. The converted exchange property can be null.
<code>mandatoryExchangePropertyAsIndex(key, type, index)</code>	Type	To be used for collecting an exchange property from an existing Collection, Map or array (lookup by the index) and then converting the exchange property to the given type determined by its classname. Expects the exchange property to be not null.

For example given the following simple expression:

```
Hello ${body[0].name}
```

This script has no type information, and the simple language will resolve this at runtime, by introspecting the message body and if it's a collection based then lookup the first element, and then invoke a method named **getName** via reflection.

In `csimple` (compiled) we want to pre compile this and therefore the end user must provide type information with the `bodyAsIndex` function:

```
Hello ${bodyAsIndex(com.foo.MyUser, 0).name}
```

67.2. COMPILATION

The `csimple` language is parsed into regular Java source code and compiled together with all the other source code, or it can be compiled once during bootstrap via the **camel-csimple-joor** module.

There are two ways to compile `csimple`

- using the **camel-csimple-maven-plugin** generating source code at built time.
- using **camel-csimple-joor** which does runtime in-memory compilation during bootstrap of Camel.

67.2.1. Using camel-csimple-maven-plugin

The **camel-csimple-maven-plugin** Maven plugin is used for discovering all the `csimple` scripts from the source code, and then automatic generate source code in the **src/generated/java** folder, which then gets compiled together with all the other sources.

The maven plugin will do source code scanning of **.java** and **.xml** files (Java and XML DSL). The scanner limits to detect certain code patterns, and it may miss discovering some `csimple` scripts if they are being used in unusual/rare ways.

The runtime compilation using **camel-csimple-joor** does not have this limitation.

The benefit is all the csimple scripts will be compiled using the regular Java compiler and therefore everything is included out of the box as **.class** files in the application JAR file, and no additional dependencies is required at runtime.

To use **camel-csimple-maven-plugin** you need to add it to your **pom.xml** file as shown:

```
<plugins>
  <!-- generate source code for csimple languages -->
  <plugin>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-csimple-maven-plugin</artifactId>
    <version>${camel.version}</version>
    <executions>
      <execution>
        <id>generate</id>
        <goals>
          <goal>generate</goal>
        </goals>
      </execution>
    </executions>
  </plugin>
  <!-- include source code generated to maven sources paths -->
  <plugin>
    <groupId>org.codehaus.mojo</groupId>
    <artifactId>build-helper-maven-plugin</artifactId>
    <version>3.1.0</version>
    <executions>
      <execution>
        <phase>generate-sources</phase>
        <goals>
          <goal>add-source</goal>
          <goal>add-resource</goal>
        </goals>
        <configuration>
          <sources>
            <source>src/generated/java</source>
          </sources>
          <resources>
            <resource>
              <directory>src/generated/resources</directory>
            </resource>
          </resources>
        </configuration>
      </execution>
    </executions>
  </plugin>
</plugins>
```

And then you must also add the **build-helper-maven-plugin** Maven plugin to include **src/generated** to the list of source folders for the Java compiler, to ensure the generated source code is compiled and included in the application JAR file.

See the **camel-example-csimple** example at [Camel Examples](#) which uses the maven plugin.

67.2.2. Using camel-csimple-joor

The jOOR library integrates with the Java compiler and performs runtime compilation of Java code.

The supported runtime when using **camel-simple-joor** is intended for Java standalone, Spring Boot, Camel Quarkus and other microservices runtimes. It is not supported in OSGi, Camel Karaf or any kind of Java Application Server runtime.

jOOR does not support runtime compilation with Spring Boot using *fat jar* packaging (<https://github.com/jOOQ/jOOR/issues/69>), it works with exploded classpath.

To use **camel-simple-joor** you simply just add it as dependency to the classpath:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-csimple-joor</artifactId>
  <version>3.14.5.redhat-00032</version>
</dependency>
```

There is no need for adding Maven plugins to the **pom.xml** file.

See the **camel-example-csimple-joor** example at [Camel Examples](#) which uses the jOOR compiler.

67.3. CSIMPLE LANGUAGE OPTIONS

The CSimple language supports 2 options, which are listed below.

Name	Default	Java Type	Description
resultType		String	Sets the class name of the result type (type from output).
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

67.4. LIMITATIONS

Currently, the csimple language does **not** support:

- nested functions (aka functions inside functions)
- the *null safe* operator (**?**).

For example the following scripts cannot compile:

```
Hello ${bean:greeter(${body}, ${header.counter})}
```

```
${bodyAs(MyUser)?.address?.zip} > 10000
```

67.5. AUTO IMPORTS

The csimple language will automatically import from:

```
import java.util.*;
import java.util.concurrent.*;
import java.util.stream.*;
import org.apache.camel.*;
import org.apache.camel.util.*;
```

67.6. CONFIGURATION FILE

You can configure the csimple language in the **camel-csimple.properties** file which is loaded from the root classpath.

For example you can add additional imports in the **camel-csimple.properties** file by adding:

```
import com.foo.MyUser;
import com.bar.*;
import static com.foo.MyHelper.*;
```

You can also add aliases (key=value) where an alias will be used as a shorthand replacement in the code.

```
echo()=${bodyAs(String)} ${bodyAs(String)}
```

Which allows to use `echo()` in the csimple language script such as:

```
from("direct:hello")
  .transform(csimple("Hello echo()"))
  .log("You said ${body}");
```

The `echo()` alias will be replaced with its value resulting in a script as:

```
.transform(csimple("Hello ${bodyAs(String)} ${bodyAs(String)}"))
```

67.7. SEE ALSO

See the [Simple](#) language as **csimple** has the same set of functions as simple language.

67.8. SPRING BOOT AUTO-CONFIGURATION

When using csimple with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.protocol</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.user-name</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as <code>name.namespace.svc.dnsDomain</code> . When using dnssrv the service name is resolved with SRV query for <code>....svc...</code> . When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable <code>KUBERNETES_MASTER</code> .		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String

Name	Description	Default	Type
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
camel.hystrix.core-pool-size	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
camel.hystrix.enabled	Enable the component.	true	Boolean
camel.hystrix.execution-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
camel.hystrix.execution-isolation-strategy	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
camel.hystrix.execution-isolation-thread-interrupt-on-timeout	Whether the execution thread should attempt an interrupt (using <code>Future#cancel</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
camel.hystrix.execution-timeout-enabled	Whether the timeout mechanism is enabled for this command.	true	Boolean
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer

Name	Description	Default	Type
camel.hystrix.group-key	Sets the group key to use. The default value is CamelHystrix.	Camel Hystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long,TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer

Name	Description	Default	Type
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.csimple.enabled</code>	Whether to enable auto configuration of the csimple language. This is enabled by default.		Boolean
<code>camel.language.csimple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean

Name	Description	Default	Type
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer

Name	Description	Default	Type
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
<code>camel.resilience4j.sliding-window-type</code>	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	<code>COUNT_BASED</code>	String
<code>camel.resilience4j.slow-call-duration-threshold</code>	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.slow-call-rate-threshold</code>	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float

Name	Description	Default	Type
camel.resilience4j.wait-duration-in-open-state	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
camel.resilience4j.writable-stack-trace-enabled	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
camel.rest.api-component	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
camel.rest.api-context-path	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
camel.rest.api-context-route-id	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
camel.rest.api-host	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
camel.rest.api-property	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
camel.rest.api-vendor-extension	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean

Name	Description	Default	Type
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBElement is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map

Name	Description	Default	Type
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String

Name	Description	Default	Type
<code>camel.rest.skip-binding-on-error-code</code>	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 68. EXCHANGEPROPERTY

The ExchangeProperty Expression Language allows you to extract values of named exchange properties.

68.1. EXCHANGE PROPERTY OPTIONS

The ExchangeProperty language supports 1 options, which are listed below.

Name	Default	Java Type	Description
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

68.2. EXAMPLE

The **recipientList** EIP can utilize a exchangeProperty like:

```
<route>
  <from uri="direct:a" />
  <recipientList>
    <exchangeProperty>myProperty</exchangeProperty>
  </recipientList>
</route>
```

In this case, the list of recipients are contained in the property 'myProperty'.

And the same example in Java DSL:

```
from("direct:a").recipientList(exchangeProperty("myProperty"));
```

68.3. DEPENDENCIES

The ExchangeProperty language is part of **camel-core**.

68.4. SPRING BOOT AUTO-CONFIGURATION

When using exchangeProperty with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.user-name</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as <code>name.namespace.svc.dnsDomain</code> . When using dnssrv the service name is resolved with SRV query for <code>....svc...</code> . When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable <code>KUBERNETES_MASTER</code> .		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String

Name	Description	Default	Type
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
camel.hystrix.core-pool-size	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
camel.hystrix.enabled	Enable the component.	true	Boolean
camel.hystrix.execution-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
camel.hystrix.execution-isolation-strategy	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
camel.hystrix.execution-isolation-thread-interrupt-on-timeout	Whether the execution thread should attempt an interrupt (using <code>Future#cancel()</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
camel.hystrix.execution-timeout-enabled	Whether the timeout mechanism is enabled for this command.	true	Boolean
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer

Name	Description	Default	Type
camel.hystrix.group-key	Sets the group key to use. The default value is CamelHystrix.	Camel Hystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long,TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> . This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String

Name	Description	Default	Type
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer

Name	Description	Default	Type
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
<code>camel.resilience4j.sliding-window-type</code>	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	<code>COUNT_BASED</code>	String
<code>camel.resilience4j.slow-call-duration-threshold</code>	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.slow-call-rate-threshold</code>	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float

Name	Description	Default	Type
<code>camel.resilience4j.wait-duration-in-open-state</code>	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.writable-stack-trace-enabled</code>	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
<code>camel.rest.api-component</code>	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.rest.api-context-path</code>	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
<code>camel.rest.api-context-route-id</code>	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
<code>camel.rest.api-host</code>	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
<code>camel.rest.api-property</code>	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
<code>camel.rest.api-vendor-extension</code>	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean

Name	Description	Default	Type
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBElement is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean

Name	Description	Default	Type
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
camel.rest.skip-binding-on-error-code	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean

Name	Description	Default	Type
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 69. FILE

The File Expression Language is an extension to the language, adding file related capabilities. These capabilities are related to common use cases working with file path and names. The goal is to allow expressions to be used with the

components for setting dynamic file patterns for both consumer and producer.



NOTE

The file language is merged with language which means you can use all the file syntax directly within the simple language.

69.1. FILE LANGUAGE OPTIONS

The File language supports 2 options, which are listed below.

Name	Default	Java Type	Description
<code>resultType</code>		String	Sets the class name of the result type (type from output).
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

69.2. SYNTAX

This language is an **extension** to the language so the syntax applies also. So the table below only lists the additional file related functions.

All the file tokens use the same expression name as the method on the **java.io.File** object, for instance **file:absolute** refers to the **java.io.File.getAbsolute()** method. Notice that not all expressions are supported by the current Exchange. For instance the component supports some options, whereas the File component supports all of them.

Expression	Type	File Consumer	File Producer	FTP Consumer	FTP Producer	Description
<code>file:name</code>	String	yes	no	yes	no	refers to the file name (is relative to the starting directory, see note below)
<code>file:name.ext</code>	String	yes	no	yes	no	refers to the file extension only

Expression	Type	File Consumer	File Producer	FTP Consumer	FTP Producer	Description
file:name.ext.single	String	yes	no	yes	no	refers to the file extension. If the file extension has multiple dots, then this expression strips and only returns the last part.
file:name.noext	String	yes	no	yes	no	refers to the file name with no extension (is relative to the starting directory, see note below)
file:name.noext.single	String	yes	no	yes	no	refers to the file name with no extension (is relative to the starting directory, see note below). If the file extension has multiple dots, then this expression strips only the last part, and keep the others.
file:onlyname	String	yes	no	yes	no	refers to the file name only with no leading paths.
file:onlyname.noext	String	yes	no	yes	no	refers to the file name only with no extension and with no leading paths.
file:onlyname.noext.single	String	yes	no	yes	no	refers to the file name only with no extension and with no leading paths. If the file extension has multiple dots, then this expression strips only the last part, and keep the others.
file:ext	String	yes	no	yes	no	refers to the file extension only

Expression	Type	File Consumer	File Producer	FTP Consumer	FTP Producer	Description
file:parent	String	yes	no	yes	no	refers to the file parent
file:path	String	yes	no	yes	no	refers to the file path
file:absolute	Boolean	yes	no	no	no	refers to whether the file is regarded as absolute or relative
file:absolute.path	String	yes	no	no	no	refers to the absolute file path
file:length	Long	yes	no	yes	no	refers to the file length returned as a Long type
file:size	Long	yes	no	yes	no	refers to the file length returned as a Long type
file:modified	Date	yes	no	yes	no	Refers to the file last modified returned as a Date type
date_command:pattern_	String	yes	yes	yes	yes	for date formatting using the java.text.SimpleDateFormat patterns. Is an extension to the language. Additional command is: file (consumers only) for the last modified timestamp of the file. Notice: all the commands from the language can also be used.

69.3. FILE TOKEN EXAMPLE

69.3.1. Relative paths

We have a **java.io.File** handle for the file **hello.txt** in the following **relative** directory: **.filelanguage/test**. And we configure our endpoint to use this starting directory **.filelanguage**. The file tokens will return as:

Expression	Returns
file:name	test\hello.txt
file:name.ext	txt
file:name.noext	test\hello
file:onlyname	hello.txt
file:onlyname.noext	hello
file:ext	txt
file:parent	filelanguage\test
file:path	filelanguage\test\hello.txt
file:absolute	false
file:absolute.path	\workspace\camel\camel-core\target\filelanguage\test\hello.txt

69.3.2. Absolute paths

We have a **java.io.File** handle for the file **hello.txt** in the following **absolute** directory: **\workspace\camel\camel-core\target\filelanguage\test**. And we configure out endpoint to use the absolute starting directory **\workspace\camel\camel-core\target\filelanguage**. The file tokens will return as:

Expression	Returns
file:name	test\hello.txt
file:name.ext	txt
file:name.noext	test\hello
file:onlyname	hello.txt
file:onlyname.noext	hello
file:ext	txt
file:parent	\workspace\camel\camel-core\target\filelanguage\test

Expression	Returns
file:path	\workspace\camel\camel-core\target\filelanguage\test\hello.txt
file:absolute	true
file:absolute.path	\workspace\camel\camel-core\target\filelanguage\test\hello.txt

69.4. SAMPLES

You can enter a fixed file name such as **myfile.txt**:

```
fileName="myfile.txt"
```

Let's assume we use the file consumer to read files and want to move the read files to back up folder with the current date as a sub folder. This can be done using an expression like:

```
fileName="backup/${date:now:yyyyMMdd}/${file:name.noext}.bak"
```

relative folder names are also supported so suppose the backup folder should be a sibling folder then you can append .. as shown:

```
fileName="../backup/${date:now:yyyyMMdd}/${file:name.noext}.bak"
```

As this is an extension to the language we have access to all the goodies from this language also, so in this use case we want to use the `in.header.type` as a parameter in the dynamic expression:

```
fileName="../backup/${date:now:yyyyMMdd}/type-${in.header.type}/backup-of-${file:name.noext}.bak"
```

If you have a custom date you want to use in the expression then Camel supports retrieving dates from the message header:

```
fileName="orders/order-${in.header.customerId}-${date:in.header.orderDate:yyyyMMdd}.xml"
```

And finally we can also use a bean expression to invoke a POJO class that generates some String output (or convertible to String) to be used:

```
fileName="uniquefile-${bean:myguidgenerator.generateid}.txt"
```

Of course all this can be combined in one expression where you can use the `and` the language in one combined expression. This is pretty powerful for those common file path patterns.

69.5. DEPENDENCIES

The File language is part of **camel-core**.

69.6. SPRING BOOT AUTO-CONFIGURATION

When using file with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
camel.cloud.consul.service-discovery.acl-token	Sets the ACL token to be used with Consul.		String
camel.cloud.consul.service-discovery.block-seconds	The seconds to wait for a watch event, default 10 seconds.	10	Integer
camel.cloud.consul.service-discovery.configurations	Define additional configuration definitions.		Map
camel.cloud.consul.service-discovery.connect-timeout-millis	Connect timeout for OkHttpClient.		Long
camel.cloud.consul.service-discovery.datacenter	The data center.		String
camel.cloud.consul.service-discovery.enabled	Enable the component.	true	Boolean
camel.cloud.consul.service-discovery.password	Sets the password to be used for basic authentication.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.username</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as name.namespace.svc.dnsDomain. When using dnssrv the service name is resolved with SRV query for svc... When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable KUBERNETES_MASTER.		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DNSSRV lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String

Name	Description	Default	Type
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean

Name	Description	Default	Type
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map
<code>camel.hystrix.core-pool-size</code>	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
<code>camel.hystrix.enabled</code>	Enable the component.	true	Boolean
<code>camel.hystrix.execution-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
<code>camel.hystrix.execution-isolation-strategy</code>	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
<code>camel.hystrix.execution-isolation-thread-interrupt-on-timeout</code>	Whether the execution thread should attempt an interrupt (using <code>Future#cancel</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
<code>camel.hystrix.execution-timeout-enabled</code>	Whether the timeout mechanism is enabled for this command.	true	Boolean

Name	Description	Default	Type
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer
camel.hystrix.group-key	Sets the group key to use. The default value is <code>CamelHystrix</code> .	CamelHystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long, TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> . This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-bucket-size</code>	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-percentile-enabled</code>	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
<code>camel.hystrix.metrics-rolling-percentile-window-buckets</code>	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean

Name	Description	Default	Type
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.csimple.enabled</code>	Whether to enable auto configuration of the <code>csimple</code> language. This is enabled by default.		Boolean
<code>camel.language.csimple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the <code>exchangeProperty</code> language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the CircuitBreaker can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the CircuitBreaker will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
<code>camel.resilience4j.sliding-window-type</code>	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	<code>COUNT_BASED</code>	String

Name	Description	Default	Type
<code>camel.resilience4j.slow-call-duration-threshold</code>	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.slow-call-rate-threshold</code>	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float
<code>camel.resilience4j.wait-duration-in-open-state</code>	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.writable-stack-trace-enabled</code>	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
<code>camel.rest.api-component</code>	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.rest.api-context-path</code>	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
<code>camel.rest.api-context-route-id</code>	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String

Name	Description	Default	Type
camel.rest.api-host	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
camel.rest.api-property	Allows to configure as many additional properties for the api documentation (swagger). For example set property api.title to my cool stuff.		Map
camel.rest.api-vendor-extension	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String

Name	Description	Default	Type
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBELEMENT is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String

Name	Description	Default	Type
<code>camel.rest.producer-api-doc</code>	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
<code>camel.rest.producer-component</code>	Sets the name of the Camel component to use as the REST producer.		String
<code>camel.rest.scheme</code>	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
<code>camel.rest.skip-binding-on-error-code</code>	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from PatternHelper#matchPattern(String,String).		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 70. HEADER

The Header Expression Language allows you to extract values of named headers.

70.1. HEADER OPTIONS

The Header language supports 1 options, which are listed below.

Name	Default	Java Type	Description
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

70.2. EXAMPLE USAGE

The **recipientList** EIP can utilize a header:

```
<route>
  <from uri="direct:a" />
  <recipientList>
    <header>myHeader</header>
  </recipientList>
</route>
```

In this case, the list of recipients are contained in the header 'myHeader'.

And the same example in Java DSL:

```
from("direct:a").recipientList(header("myHeader"));
```

70.3. DEPENDENCIES

The Header language is part of **camel-core**.

70.4. SPRING BOOT AUTO-CONFIGURATION

When using header with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.user-name</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as <code>name.namespace.svc.dnsDomain</code> . When using dnssrv the service name is resolved with SRV query for <code>....svc...</code> . When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable <code>KUBERNETES_MASTER</code> .		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String

Name	Description	Default	Type
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
camel.hystrix.core-pool-size	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
camel.hystrix.enabled	Enable the component.	true	Boolean
camel.hystrix.execution-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
camel.hystrix.execution-isolation-strategy	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
camel.hystrix.execution-isolation-thread-interrupt-on-timeout	Whether the execution thread should attempt an interrupt (using <code>Future#cancel</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
camel.hystrix.execution-timeout-enabled	Whether the timeout mechanism is enabled for this command.	true	Boolean
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer

Name	Description	Default	Type
camel.hystrix.group-key	Sets the group key to use. The default value is CamelHystrix.	Camel Hystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long,TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer

Name	Description	Default	Type
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
<code>camel.resilience4j.sliding-window-type</code>	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	<code>COUNT_BASED</code>	String
<code>camel.resilience4j.slow-call-duration-threshold</code>	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.slow-call-rate-threshold</code>	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float

Name	Description	Default	Type
camel.resilience4j.wait-duration-in-open-state	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
camel.resilience4j.writable-stack-trace-enabled	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
camel.rest.api-component	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
camel.rest.api-context-path	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
camel.rest.api-context-route-id	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
camel.rest.api-host	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
camel.rest.api-property	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
camel.rest.api-vendor-extension	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode

Name	Description	Default	Type
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBELEMENT is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean

Name	Description	Default	Type
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
camel.rest.skip-binding-on-error-code	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean

Name	Description	Default	Type
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 71. JSONPATH

Camel supports [JSONPath](#) to allow using [Expression](#) or [Predicate](#) on JSON messages.

71.1. JSONPATH OPTIONS

The JSONPath language supports 8 options, which are listed below.

Name	Default	Java Type	Description
<code>resultType</code>		String	Sets the class name of the result type (type from output).
<code>suppressExceptions</code>		Boolean	Whether to suppress exceptions such as <code>PathNotFoundException</code> .
<code>allowSimple</code>		Boolean	Whether to allow inlined Simple exceptions in the JSONPath expression.
<code>allowEasyPredicate</code>		Boolean	Whether to allow using the easy predicate parser to pre-parse predicates.
<code>writeAsString</code>		Boolean	Whether to write the output of each row/element as a JSON String value instead of a Map/POJO value.
<code>headerName</code>		String	Name of header to use as input, instead of the message body.
<code>option</code>		Enum	To configure additional options on JSONPath. Multiple values can be separated by comma. Enum values: <ul style="list-style-type: none"> ● <code>DEFAULT_PATH_LEAF_TO_NULL</code> ● <code>ALWAYS_RETURN_LIST</code> ● <code>AS_PATH_LIST</code> ● <code>SUPPRESS_EXCEPTIONS</code> ● <code>REQUIRE_PROPERTIES</code>
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

71.2. EXAMPLES

For example, you can use JSONPath in a [Predicate](#) with the [Content Based Router EIP](#) .

```
from("queue:books.new")
  .choice()
```

```
.when().jsonpath("$.store.book[?(@.price < 10)"])\n  .to("jms:queue:book.cheap")\n.when().jsonpath("$.store.book[?(@.price < 30)"])\n  .to("jms:queue:book.average")\n.otherwise()\n  .to("jms:queue:book.expensive");
```

And in XML DSL:

```
<route>\n  <from uri="direct:start"/>\n  <choice>\n    <when>\n      <jsonpath>$.store.book[?(@.price &lt; 10)]</jsonpath>\n      <to uri="mock:cheap"/>\n    </when>\n    <when>\n      <jsonpath>$.store.book[?(@.price &lt; 30)]</jsonpath>\n      <to uri="mock:average"/>\n    </when>\n    <otherwise>\n      <to uri="mock:expensive"/>\n    </otherwise>\n  </choice>\n</route>
```

71.3. JSONPATH SYNTAX

Using the JSONPath syntax takes some time to learn, even for basic predicates. So for example to find out all the cheap books you have to do:

```
$.store.book[?(@.price < 20)]
```

71.3.1. Easy JSONPath Syntax

However, what if you could just write it as:

```
store.book.price < 20
```

And you can omit the path if you just want to look at nodes with a price key:

```
price < 20
```

To support this there is a **EasyPredicateParser** which kicks-in if you have defined the predicate using a basic style. That means the predicate must not start with the **\$** sign, and only include one operator.

The easy syntax is:

```
left OP right
```

You can use Camel simple language in the right operator, eg:

```
store.book.price < ${header.limit}
```

See the [JSONPath](#) project page for more syntax examples.

71.4. SUPPORTED MESSAGE BODY TYPES

Camel JSonPath supports message body using the following types:

Type	Comment
File	Reading from files
String	Plain strings
Map	Message bodies as java.util.Map types
List	Message bodies as java.util.List types
POJO	Optional If Jackson is on the classpath, then camel-jsonpath is able to use Jackson to read the message body as POJO and convert to java.util.Map which is supported by JSonPath. For example, you can add camel-jackson as dependency to include Jackson.
InputStream	If none of the above types matches, then Camel will attempt to read the message body as a java.io.InputStream .

If a message body is of unsupported type then an exception is thrown by default, however you can configure JSonPath to suppress exceptions (see below)

71.5. SUPPRESSING EXCEPTIONS

By default, jsonpath will throw an exception if the json payload does not have a valid path accordingly to the configured jsonpath expression. In some use-cases you may want to ignore this in case the json payload contains optional data. Therefore, you can set the option **suppressExceptions** to **true** to ignore this as shown:

```
from("direct:start")
  .choice()
    // use true to suppress exceptions
    .when().jsonpath("person.middlename", true)
      .to("mock:middle")
    .otherwise()
      .to("mock:other");
```

And in XML DSL:

```
<route>
  <from uri="direct:start"/>
  <choice>
    <when>
```

```

    <jsonpath suppressExceptions="true">person.middlename</jsonpath>
    <to uri="mock:middle"/>
  </when>
  <otherwise>
    <to uri="mock:other"/>
  </otherwise>
</choice>
</route>

```

This option is also available on the `@JsonPath` annotation.

71.6. INLINE SIMPLE EXPRESSIONS

It's possible to inlined [Simple](#) language in the JSONPath expression using the simple syntax `#{xxx}`.

An example is shown below:

```

from("direct:start")
  .choice()
  .when().jsonpath("$.store.book[?(@.price < #{header.cheap})]")
    .to("mock:cheap")
  .when().jsonpath("$.store.book[?(@.price < #{header.average})]")
    .to("mock:average")
  .otherwise()
    .to("mock:expensive");

```

And in XML DSL:

```

<route>
  <from uri="direct:start"/>
  <choice>
    <when>
      <jsonpath>$.store.book[?(@.price &lt; #{header.cheap})]</jsonpath>
      <to uri="mock:cheap"/>
    </when>
    <when>
      <jsonpath>$.store.book[?(@.price &lt; #{header.average})]</jsonpath>
      <to uri="mock:average"/>
    </when>
    <otherwise>
      <to uri="mock:expensive"/>
    </otherwise>
  </choice>
</route>

```

You can turn off support for inlined Simple expression by setting the option `allowSimple` to `false` as shown:

```

.when().jsonpath("$.store.book[?(@.price < 10)]", false, false)

```

And in XML DSL:

```

<jsonpath allowSimple="false">$.store.book[?(@.price &lt; 10)]</jsonpath>

```

71.7. JSONPATH INJECTION

You can use [Bean Integration](#) to invoke a method on a bean and use various languages such as JSONPath (via the `@JsonPath` annotation) to extract a value from the message and bind it to a method parameter, as shown below:

```
public class Foo {

    @Consume("activemq:queue:books.new")
    public void doSomething(@JsonPath("$.store.book[*].author") String author, @Body String json) {
        // process the inbound message here
    }
}
```

71.8. ENCODING DETECTION

The encoding of the JSON document is detected automatically, if the document is encoded in unicode (UTF-8, UTF-16LE, UTF-16BE, UTF-32LE, UTF-32BE) as specified in RFC-4627. If the encoding is a non-unicode encoding, you can either make sure that you enter the document in String format to JSONPath, or you can specify the encoding in the header `CamelJsonPathJsonEncoding` which is defined as a constant in: `JsonpathConstants.HEADER_JSON_ENCODING`.

71.9. SPLIT JSON DATA INTO SUB ROWS AS JSON

You can use JSONPath to split a JSON document, such as:

```
from("direct:start")
    .split().jsonpath("$.store.book[*]")
    .to("log:book");
```

Then each book is logged, however the message body is a **Map** instance. Sometimes you may want to output this as plain String JSON value instead, which can be done with the `writeAsString` option as shown:

```
from("direct:start")
    .split().jsonpathWriteAsString("$.store.book[*]")
    .to("log:book");
```

Then each book is logged as a String JSON value.

71.10. USING HEADER AS INPUT

By default, JSONPath uses the message body as the input source. However, you can also use a header as input by specifying the `headerName` option.

For example to count the number of books from a JSON document that was stored in a header named **books** you can do:

```
from("direct:start")
    .setHeader("numberOfBooks")
    .jsonpath("$.store.book.length()", false, int.class, "books")
    .to("mock:result");
```

In the **jsonpath** expression above we specify the name of the header as **books**, and we also told that we wanted the result to be converted as an integer by **int.class**.

The same example in XML DSL would be:

```
<route>
  <from uri="direct:start"/>
  <setHeader name="numberOfBooks">
    <jsonpath headerName="books" resultType="int">$.store.book.length()</jsonpath>
  </setHeader>
  <to uri="mock:result"/>
</route>
```

71.11. SPRING BOOT AUTO-CONFIGURATION

When using jsonpath with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-jsonpath-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 8 options, which are listed below.

Name	Description	Default	Type
<code>camel.language.jsonpath.allow-easy-predicate</code>	Whether to allow using the easy predicate parser to pre-parse predicates.	true	Boolean
<code>camel.language.jsonpath.allow-simple</code>	Whether to allow inlined Simple exceptions in the JSONPath expression.	true	Boolean
<code>camel.language.jsonpath.enabled</code>	Whether to enable auto configuration of the jsonpath language. This is enabled by default.		Boolean
<code>camel.language.jsonpath.header-name</code>	Name of header to use as input, instead of the message body.		String
<code>camel.language.jsonpath.option</code>	To configure additional options on JSONPath. Multiple values can be separated by comma.		String
<code>camel.language.jsonpath.suppress-exceptions</code>	Whether to suppress exceptions such as PathNotFoundException.	false	Boolean

Name	Description	Default	Type
<code>camel.language.js onpath.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.js onpath.write-as- string</code>	Whether to write the output of each row/element as a JSON String value instead of a Map/POJO value.	false	Boolean

CHAPTER 72. REF

The Ref Expression Language is really just a way to lookup a custom **Expression** or **Predicate** from the [Registry](#).

This is particular useable in XML DSLs.

72.1. REF LANGUAGE OPTIONS

The Ref language supports 1 options, which are listed below.

Name	Default	Java Type	Description
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

72.2. EXAMPLE USAGE

The Splitter EIP in XML DSL can utilize a custom expression using `<ref>` like:

```
<bean id="myExpression" class="com.mycompany.MyCustomExpression"/>
<route>
  <from uri="seda:a"/>
  <split>
    <ref>myExpression</ref>
    <to uri="mock:b"/>
  </split>
</route>
```

in this case, the message coming from the seda:a endpoint will be splitted using a custom **Expression** which has the id **myExpression** in the [Registry](#).

And the same example using Java DSL:

```
from("seda:a").split().ref("myExpression").to("seda:b");
```

72.3. DEPENDENCIES

The Ref language is part of **camel-core**.

72.4. SPRING BOOT AUTO-CONFIGURATION

When using ref with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
```

```

<version>3.14.5.redhat-00032</version>
<!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.protocol</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.username</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as <code>name.namespace.svc.dnsDomain</code> . When using dnssrv the service name is resolved with SRV query for <code>....svc...</code> . When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable <code>KUBERNETES_MASTER</code> .		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String

Name	Description	Default	Type
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.hystrix.core-pool-size</code>	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
<code>camel.hystrix.enabled</code>	Enable the component.	true	Boolean
<code>camel.hystrix.execution-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
<code>camel.hystrix.execution-isolation-strategy</code>	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
<code>camel.hystrix.execution-isolation-thread-interrupt-on-timeout</code>	Whether the execution thread should attempt an interrupt (using <code>Future#cancel</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
<code>camel.hystrix.execution-timeout-enabled</code>	Whether the timeout mechanism is enabled for this command.	true	Boolean
<code>camel.hystrix.execution-timeout-in-milliseconds</code>	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
<code>camel.hystrix.fallback-enabled</code>	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
<code>camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer

Name	Description	Default	Type
camel.hystrix.group-key	Sets the group key to use. The default value is CamelHystrix.	Camel Hystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long,TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String

Name	Description	Default	Type
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer

Name	Description	Default	Type
camel.resilience4j.permitted-number-of-calls-in-half-open-state	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
camel.resilience4j.sliding-window-size	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
camel.resilience4j.sliding-window-type	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	COUNT_BASED	String
camel.resilience4j.slow-call-duration-threshold	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
camel.resilience4j.slow-call-rate-threshold	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float

Name	Description	Default	Type
<code>camel.resilience4j.wait-duration-in-open-state</code>	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.writable-stack-trace-enabled</code>	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
<code>camel.rest.api-component</code>	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.rest.api-context-path</code>	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
<code>camel.rest.api-context-route-id</code>	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
<code>camel.rest.api-host</code>	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
<code>camel.rest.api-property</code>	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
<code>camel.rest.api-vendor-extension</code>	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean

Name	Description	Default	Type
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBElement is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map

Name	Description	Default	Type
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String

Name	Description	Default	Type
<code>camel.rest.skip-binding-on-error-code</code>	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 73. XQUERY

Camel supports [XQuery](#) to allow an [Expression](#) or [Predicate](#) to be used in the [DSL](#).

For example, you could use XQuery to create a predicate in a [Message Filter](#) or as an expression for a [Recipient List](#).

73.1. XQUERY LANGUAGE OPTIONS

The XQuery language supports 4 options, which are listed below.

Name	Default	Java Type	Description
<code>type</code>		String	Sets the class name of the result type (type from output) The default result type is NodeSet.
<code>headerName</code>		String	Name of header to use as input, instead of the message body.
<code>configurationRef</code>		String	Reference to a saxon configuration instance in the registry to use for xquery (requires camel-saxon). This may be needed to add custom functions to a saxon configuration, so these custom functions can be used in xquery expressions.
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

73.2. VARIABLES

The message body will be set as the **contextItem**. And the following variables are available as well:

Variable	Type	Description
<code>exchange</code>	Exchange	The current Exchange
<code>in.body</code>	Object	The message body
<code>out.body</code>	Object	deprecated The OUT message body (if any)
<code>in.headers.*</code>	Object	You can access the value of exchange.in.headers with key foo by using the variable which name is in.headers.foo
<code>out.headers.*</code>	Object	deprecated You can access the value of exchange.out.headers with key foo by using the variable which name is out.headers.foo variable

Variable	Type	Description
key name	Object	Any exchange.properties and exchange.in.headers and any additional parameters set using setParameters(Map) . These parameters are added with they own key name, for instance if there is an IN header with the key name foo then its added as foo .

73.3. EXAMPLE

```
from("queue:foo")
  .filter().xquery("//foo")
  .to("queue:bar")
```

You can also use functions inside your query, in which case you need an explicit type conversion, or you will get an **org.w3c.dom.DOMException: HIERARCHY_REQUEST_ERR**). You need to pass in the expected output type of the function. For example the concat function returns a **String** which is done as shown:

```
from("direct:start")
  .recipientList().xquery("concat('mock:foo.', /person/@city)", String.class);
```

And in XML DSL:

```
<route>
  <from uri="direct:start"/>
  <recipientList>
    <xquery type="java.lang.String">concat('mock:foo.', /person/@city</xquery>
  </recipientList>
</route>
```

73.3.1. Using namespaces

If you have a standard set of namespaces you wish to work with and wish to share them across many XQuery expressions you can use the **org.apache.camel.support.builder.Namespaces** when using Java DSL as shown:

```
Namespaces ns = new Namespaces("c", "http://acme.com/cheese");

from("direct:start")
  .filter().xquery("/c:person[@name='James']", ns)
  .to("mock:result");
```

Notice how the namespaces are provided to **xquery** with the **ns** variable that are passed in as the 2nd parameter.

Each namespace is a key=value pair, where the prefix is the key. In the XQuery expression then the namespace is used by its prefix, eg:

```
/c:person[@name='James']
```

The namespace builder supports adding multiple namespaces as shown:

```
Namespaces ns = new Namespaces("c", "http://acme.com/cheese")
    .add("w", "http://acme.com/wine")
    .add("b", "http://acme.com/beer");
```

When using namespaces in XML DSL then its different, as you setup the namespaces in the XML root tag (or one of the **camelContext**, **routes**, **route** tags).

In the XML example below we use Spring XML where the namespace is declared in the root tag **beans**, in the line with **xmlns:foo="http://example.com/person"**:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:foo="http://example.com/person"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-
    spring.xsd">

  <camelContext id="camel" xmlns="http://activemq.apache.org/camel/schema/spring">
    <route>
      <from uri="activemq:MyQueue"/>
      <filter>
        <xquery>/foo:person[@name='James']</xquery>
        <to uri="mqseries:SomeOtherQueue"/>
      </filter>
    </route>
  </camelContext>
</beans>
```

This namespace uses **foo** as prefix, so the **<xquery>** expression uses **/foo:** to use this namespace.

73.4. USING XQUERY AS TRANSFORMATION

We can do a message translation using transform or setBody in the route, as shown below:

```
from("direct:start").
  transform().xquery("/people/person");
```

Notice that xquery will use DOMResult by default, so if we want to grab the value of the person node, using **text()** we need to tell XQuery to use String as result type, as shown:

```
from("direct:start").
  transform().xquery("/people/person/text()", String.class);
```

If you want to use Camel variables like headers, you have to explicitly declare them in the XQuery expression.

```
<transform>
  <xquery>
    declare variable $in.headers.foo external;
```

```

    element item {$in.headers.foo}
  </xquery>
</transform>

```

73.5. LOADING SCRIPT FROM EXTERNAL RESOURCE

You can externalize the script and have Camel load it from a resource such as **"classpath:"**, **"file:"**, or **"http:"**. This is done using the following syntax: **"resource:scheme:location"**, e.g. to refer to a file on the classpath you can do:

```
.setHeader("myHeader").xquery("resource:classpath:myxquery.txt", String.class)
```

73.6. LEARNING XQUERY

XQuery is a very powerful language for querying, searching, sorting and returning XML. For help learning XQuery try these tutorials

- Mike Kay's [XQuery Primer](#)
- The W3Schools [XQuery Tutorial](#)

73.7. DEPENDENCIES

To use XQuery in your camel routes you need to add the a dependency on **camel-saxon** which implements the XQuery language.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```

<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-saxon</artifactId>
  <version>3.14.5.redhat-00018</version>
</dependency>

```

73.8. SPRING BOOT AUTO-CONFIGURATION

When using xquery with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```

<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-saxon-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>

```

The component supports 11 options, which are listed below.

Name	Description	Default	Type
<code>camel.component.xquery.autowired-enabled</code>	Whether autowiring is enabled. This is used for automatic autowiring options (the option must be marked as autowired) by looking up in the registry to find if there is a single instance of matching type, which then gets configured on the component. This can be used for automatic configuring JDBC data sources, JMS connection factories, AWS Clients, etc.	true	Boolean
<code>camel.component.xquery.bridge-error-handler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	Boolean
<code>camel.component.xquery.configuration</code>	To use a custom Saxon configuration. The option is a <code>net.sf.saxon.Configuration</code> type.		Configuration
<code>camel.component.xquery.configuration-properties</code>	To set custom Saxon configuration properties.		Map
<code>camel.component.xquery.enabled</code>	Whether to enable auto configuration of the xquery component. This is enabled by default.		Boolean
<code>camel.component.xquery.lazy-start-producer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	Boolean
<code>camel.component.xquery.module-uri-resolver</code>	To use the custom ModuleURIResolver. The option is a <code>net.sf.saxon.lib.ModuleURIResolver</code> type.		ModuleURIResolver

Name	Description	Default	Type
camel.language.xquery.configuration-ref	Reference to a saxon configuration instance in the registry to use for xquery (requires camel-saxon). This may be needed to add custom functions to a saxon configuration, so these custom functions can be used in xquery expressions.		String
camel.language.xquery.enabled	Whether to enable auto configuration of the xquery language. This is enabled by default.		Boolean
camel.language.xquery.trim	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
camel.language.xquery.type	Sets the class name of the result type (type from output) The default result type is NodeSet.		String

CHAPTER 74. SIMPLE

The Simple Expression Language was a really simple language when it was created, but has since grown more powerful. It is primarily intended for being a very small and simple language for evaluating **Expression** or **Predicate** without requiring any new dependencies or knowledge of other scripting languages such as Groovy.

The simple language is designed with intend to cover almost all the common use cases when little need for scripting in your Camel routes.

However, for much more complex use cases then a more powerful language is recommended such as:

- [Groovy](#)
- [MVEL](#)
- [OGNL](#)



NOTE

The simple language requires **camel-bean** JAR as classpath dependency if the simple language uses OGNL expressions, such as calling a method named **myMethod** on the message body: **`${body.myMethod()}`**. At runtime the simple language will then us its built-in OGNL support which requires the **camel-bean** component.

The simple language uses **`${body}`** placeholders for complex expressions or functions.



NOTE

See also the [CSimple](#) language which is compiled.



NOTE

Alternative syntax

You can also use the alternative syntax which uses **`$simple{ }`** as placeholders. This can be used in situations to avoid clashes when using for example Spring property placeholder together with Camel.

74.1. SIMPLE LANGUAGE OPTIONS

The Simple language supports 2 options, which are listed below.

Name	Default	Java Type	Description
<code>resultType</code>		String	Sets the class name of the result type (type from output).
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

74.2. VARIABLES

Variable	Type	Description
camelId	String	the CamelContext name
camelContext. OGNL	Object	the CamelContext invoked using a Camel OGNL expression.
exchange	Exchange	the Exchange
exchange. OGNL	Object	the Exchange invoked using a Camel OGNL expression.
exchangeId	String	the exchange id
id	String	the message id
messageTimestamp	String	the message timestamp (millis since epoch) that this message originates from. Some systems like JMS, Kafka, AWS have a timestamp on the event/message, that Camel received. This method returns the timestamp, if a timestamp exists. The message timestamp and exchange created are not the same. An exchange always have a created timestamp which is the local timestamp when Camel created the exchange. The message timestamp is only available in some Camel components when the consumer is able to extract the timestamp from the source event. If the message has no timestamp then 0 is returned.
body	Object	the body
body. OGNL	Object	the body invoked using a Camel OGNL expression.
bodyAs(<i>type</i>)	Type	Converts the body to the given type determined by its classname. The converted body can be null.
bodyAs(<i>type</i>). OGNL	Object	Converts the body to the given type determined by its classname and then invoke methods using a Camel OGNL expression. The converted body can be null.
bodyOneLine	String	Converts the body to a String and removes all line-breaks so the string is in one line.
mandatoryBodyAs(<i>type</i>)	Type	Converts the body to the given type determined by its classname, and expects the body to be not null.
mandatoryBodyAs(<i>type</i>). OGNL	Object	Converts the body to the given type determined by its classname and then invoke methods using a Camel OGNL expression.
header.foo	Object	refer to the foo header

Variable	Type	Description
header[foo]	Object	refer to the foo header
headers.foo	Object	refer to the foo header
headers:foo	Object	refer to the foo header
headers[foo]	Object	refer to the foo header
header.foo[bar]	Object	regard foo header as a map and perform lookup on the map with bar as key
header.foo. OGNL	Object	refer to the foo header and invoke its value using a Camel OGNL expression.
headerAs(<i>key,type</i>)	Type	converts the header to the given type determined by its classname
headers	Map	refer to the headers
exchangeProperty.foo	Object	refer to the foo property on the exchange
exchangeProperty[foo]	Object	refer to the foo property on the exchange
exchangeProperty.foo. OGNL	Object	refer to the foo property on the exchange and invoke its value using a Camel OGNL expression.
sys.foo	String	refer to the JVM system property
sysenv.foo	String	refer to the system environment variable
env.foo	String	refer to the system environment variable
exception	Object	refer to the exception object on the exchange, is null if no exception set on exchange. Will fallback and grab caught exceptions (Exchange.EXCEPTION_CAUGHT) if the Exchange has any.
exception. OGNL	Object	refer to the exchange exception invoked using a Camel OGNL expression object
exception.message	String	refer to the exception.message on the exchange, is null if no exception set on exchange. Will fallback and grab caught exceptions (Exchange.EXCEPTION_CAUGHT) if the Exchange has any.

Variable	Type	Description
exception.stacktrace	String	refer to the exception.stracktrace on the exchange, is null if no exception set on exchange. Will fallback and grab caught exceptions (Exchange.EXCEPTION_CAUGHT) if the Exchange has any.
date:_command_	Date	evaluates to a Date object. Supported commands are: now for current timestamp, exchangeCreated for the timestamp when the current exchange was created, header.xxx to use the Long/Date object header with the key xxx. exchangeProperty.xxx to use the Long/Date object in the exchange property with the key xxx. file for the last modified timestamp of the file (available with a File consumer). Command accepts offsets such as: now-24h or header.xxx+1h or even now+1h30m-100 .
date:_command:pattern_	String	Date formatting using java.text.SimpleDateFormat patterns.
date-with-timezone:_command:timezone:pattern_	String	Date formatting using java.text.SimpleDateFormat timezones and patterns.
bean:_bean expression_	Object	Invoking a bean expression using the language. Specifying a method name you must use dot as separator. We also support the ?method=methodname syntax that is used by the component. Camel will by default lookup a bean by the given name. However if you need to refer to a bean class (such as calling a static method) then you can prefix with type, such as bean:type:fqnClassName .
properties:key:default	String	Lookup a property with the given key. If the key does not exists or has no value, then an optional default value can be specified.
routeld	String	Returns the id of the current route the Exchange is being routed.
stepId	String	Returns the id of the current step the Exchange is being routed.
threadName	String	Returns the name of the current thread. Can be used for logging purpose.
hostname	String	Returns the local hostname (may be empty if not possible to resolve).
ref:xxx	Object	To lookup a bean from the Registry with the given id.

Variable	Type	Description
type:name.field	Object	To refer to a type or field by its FQN name. To refer to a field you can append <code>.FIELD_NAME</code> . For example, you can refer to the constant field from Exchange as: org.apache.camel.Exchange.FILE_NAME
null	null	represents a null
random(value)	Integer	returns a random Integer between 0 (included) and <i>value</i> (excluded)
random(min,max)	Integer	returns a random Integer between <i>min</i> (included) and <i>max</i> (excluded)
collate(group)	List	The collate function iterates the message body and groups the data into sub lists of specified size. This can be used with the Splitter EIP to split a message body and group/batch the splitted sub message into a group of N sub lists. This method works similar to the collate method in Groovy.
skip(number)	Iterator	The skip function iterates the message body and skips the first number of items. This can be used with the Splitter EIP to split a message body and skip the first N number of items.
messageHistory	String	The message history of the current exchange how it has been routed. This is similar to the route stack-trace message history the error handler logs in case of an unhandled exception.
messageHistory(false)	String	As messageHistory but without the exchange details (only includes the route stack-trace). This can be used if you do not want to log sensitive data from the message itself.

74.3. OGNL EXPRESSION SUPPORT

When using **OGNL** then **camel-bean** JAR is required to be on the classpath.

Camel's OGNL support is for invoking methods only. You cannot access fields. Camel support accessing the length field of Java arrays.

The [Simple](#) and Bean language now supports a Camel OGNL notation for invoking beans in a chain like fashion. Suppose the Message IN body contains a POJO which has a **getAddress()** method.

Then you can use Camel OGNL notation to access the address object:

```
simple("${body.address}")
simple("${body.address.street}")
simple("${body.address.zip}")
```

Camel understands the shorthand names for getters, but you can invoke any method or use the real name such as:

```
simple("${body.address}")
simple("${body.getAddress.getStreet}")
simple("${body.address.getZip}")
simple("${body.doSomething}")
```

You can also use the null safe operator (**?.**) to avoid NPE if for example the body does NOT have an address

```
simple("${body?.address?.street}")
```

It is also possible to index in **Map** or **List** types, so you can do:

```
simple("${body[foo].name}")
```

To assume the body is **Map** based and lookup the value with **foo** as key, and invoke the **getName** method on that value.

If the key has space, then you **must** enclose the key with quotes, for example 'foo bar':

```
simple("${body['foo bar'].name}")
```

You can access the **Map** or **List** objects directly using their key name (with or without dots) :

```
simple("${body[foo]}")
simple("${body[this.is.foo]}")
```

Suppose there was no value with the key **foo** then you can use the null safe operator to avoid the NPE as shown:

```
simple("${body[foo]?.name}")
```

You can also access **List** types, for example to get lines from the address you can do:

```
simple("${body.address.lines[0]}")
simple("${body.address.lines[1]}")
simple("${body.address.lines[2]}")
```

There is a special **last** keyword which can be used to get the last value from a list.

```
simple("${body.address.lines[last]}")
```

And to get the 2nd last you can subtract a number, so we can use **last-1** to indicate this:

```
simple("${body.address.lines[last-1]}")
```

And the 3rd last is of course:

```
simple("${body.address.lines[last-2]}")
```

And you can call the size method on the list with

```
simple("${body.address.lines.size}")
```

Camel supports the length field for Java arrays as well, eg:

```
String[] lines = new String[]{"foo", "bar", "cat"};
exchange.getIn().setBody(lines);
```

```
simple("There are ${body.length} lines")
```

And yes you can combine this with the operator support as shown below:

```
simple("${body.address.zip} > 1000")
```

74.4. OPERATOR SUPPORT

The parser is limited to only support a single operator.

To enable it the left value must be enclosed in `$$\{ }`. The syntax is:

```
`${leftValue} OP rightValue
```

Where the **rightValue** can be a String literal enclosed in `' '`, **null**, a constant value or another expression enclosed in ``${ }`.



NOTE

There **must** be spaces around the operator.

Camel will automatically type convert the rightValue type to the leftValue type, so it is able to eg. convert a string into a numeric, so you can use `>` comparison for numeric values.

The following operators are supported:

Operator	Description
<code>==</code>	equals
<code>=~</code>	equals ignore case (will ignore case when comparing String values)
<code>></code>	greater than
<code>>=</code>	greater than or equals
<code><</code>	less than
<code>←</code>	less than or equals

Operator	Description
!=	not equals
!~=	not equals ignore case (will ignore case when comparing String values)
contains	For testing if contains in a string based value
!contains	For testing if not contains in a string based value
~~	For testing if contains by ignoring case sensitivity in a string based value
!~~	For testing if not contains by ignoring case sensitivity in a string based value
regex	For matching against a given regular expression pattern defined as a String value
!regex	For not matching against a given regular expression pattern defined as a String value
in	For matching if in a set of values, each element must be separated by comma. If you want to include an empty value, then it must be defined using double comma, eg '„bronze,silver,gold', which is a set of four values with an empty value and then the three medals.
!in	For matching if not in a set of values, each element must be separated by comma. If you want to include an empty value, then it must be defined using double comma, eg '„bronze,silver,gold', which is a set of four values with an empty value and then the three medals.
is	For matching if the left hand side type is an instance of the value.
!is	For matching if the left hand side type is not an instance of the value.
range	For matching if the left hand side is within a range of values defined as numbers: from..to..
!range	For matching if the left hand side is not within a range of values defined as numbers: from..to. .
startsWith	For testing if the left hand side string starts with the right hand string.
starts with	Same as the startsWith operator.
endsWith	For testing if the left hand side string ends with the right hand string.
ends with	Same as the endsWith operator.

And the following unary operators can be used:

Operator	Description
++	To increment a number by one. The left hand side must be a function, otherwise parsed as literal.
-	To decrement a number by one. The left hand side must be a function, otherwise parsed as literal.
\n	To use newline character.
\t	To use tab character.
\r	To use carriage return character.
\}	To use the } character as text. This may be needed when building a JSon structure with the simple language.

And the following logical operators can be used to group expressions:

Operator	Description
&&	The logical and operator is used to group two expressions.
	The logical or operator is used to group two expressions.

The syntax for AND is:

```
${leftValue} OP rightValue && ${leftValue} OP rightValue
```

And the syntax for OR is:

```
${leftValue} OP rightValue || ${leftValue} OP rightValue
```

Some examples:

```
// exact equals match
simple("${header.foo} == 'foo')
```

```
// ignore case when comparing, so if the header has value FOO this will match
simple("${header.foo} =~ 'foo')
```

```
// here Camel will type convert '100' into the type of header.bar and if it is an Integer '100' will also be
converter to an Integer
simple("${header.bar} == '100')
```

```
simple("${header.bar} == 100")

// 100 will be converted to the type of header.bar so we can do > comparison
simple("${header.bar} > 100")
```

74.4.1. Comparing with different types

When you compare with different types such as String and int, then you have to take a bit care. Camel will use the type from the left hand side as 1st priority. And fallback to the right hand side type if both values couldn't be compared based on that type.

This means you can flip the values to enforce a specific type. Suppose the bar value above is a String. Then you can flip the equation:

```
simple("100 < ${header.bar}")
```

which then ensures the int type is used as 1st priority.

This may change in the future if the Camel team improves the binary comparison operations to prefer numeric types to String based. It's most often the String type which causes problem when comparing with numbers.

```
// testing for null
simple("${header.baz} == null")

// testing for not null
simple("${header.baz} != null")
```

And a bit more advanced example where the right value is another expression

```
simple("${header.date} == ${date:now:yyyyMMdd}")

simple("${header.type} == ${bean:orderService?method=getOrderType}")
```

And an example with contains, testing if the title contains the word Camel

```
simple("${header.title} contains 'Camel'")
```

And an example with regex, testing if the number header is a 4 digit value:

```
simple("${header.number} regex '\\d{4}'")
```

And finally an example if the header equals any of the values in the list. Each element must be separated by comma, and no space around.

This also works for numbers etc, as Camel will convert each element into the type of the left hand side.

```
simple("${header.type} in 'gold,silver'")
```

And for all the last 3 we also support the negate test using not:

```
simple("${header.type} !in 'gold,silver'")
```

And you can test if the type is a certain instance, eg for instance a String

-

```
simple("${header.type} is 'java.lang.String'")
```

We have added a shorthand for all **java.lang** types so you can write it as:

```
simple("${header.type} is 'String'")
```

Ranges are also supported. The range interval requires numbers and both from and end are inclusive. For instance to test whether a value is between 100 and 199:

```
simple("${header.number} range 100..199")
```

Notice we use `..` in the range without spaces. It is based on the same syntax as Groovy.

```
simple("${header.number} range '100..199'")
```

As the XML DSL does not have all the power as the Java DSL with all its various builder methods, you have to resort to use some other languages for testing with simple operators. Now you can do this with the simple language. In the sample below we want to test if the header is a widget order:

```
<from uri="seda:orders">
  <filter>
    <simple>${header.type} == 'widget'</simple>
    <to uri="bean:orderService?method=handleWidget"/>
  </filter>
</from>
```

74.4.2. Using and / or

If you have two expressions you can combine them with the **&&** or **||** operator.

For instance:

```
simple("${header.title} contains 'Camel' && ${header.type} == 'gold'")
```

And of course the **||** is also supported. The sample would be:

```
simple("${header.title} contains 'Camel' || ${header.type} == 'gold'")
```

74.5. EXAMPLES

In the XML DSL sample below we filter based on a header value:

```
<from uri="seda:orders">
  <filter>
    <simple>${header.foo}</simple>
    <to uri="mock:fooOrders"/>
  </filter>
</from>
```

The Simple language can be used for the predicate test above in the Message Filter pattern, where we test if the in message has a **foo** header (a header with the key **foo** exists). If the expression evaluates to **true** then the message is routed to the **mock:fooOrders** endpoint, otherwise the message is dropped.

The same example in Java DSL:

```
from("seda:orders")
  .filter().simple("${header.foo}")
  .to("seda:fooOrders");
```

You can also use the simple language for simple text concatenations such as:

```
from("direct:hello")
  .transform().simple("Hello ${header.user} how are you?")
  .to("mock:reply");
```

Notice that we must use `$$` placeholders in the expression now to allow Camel to parse it correctly.

And this sample uses the date command to output current date.

```
from("direct:hello")
  .transform().simple("The today is ${date:now:yyyyMMdd} and it is a great day.")
  .to("mock:reply");
```

And in the sample below we invoke the bean language to invoke a method on a bean to be included in the returned string:

```
from("direct:order")
  .transform().simple("OrderId: ${bean:orderIdGenerator}")
  .to("mock:reply");
```

Where **orderIdGenerator** is the id of the bean registered in the Registry. If using Spring then it is the Spring bean id.

If we want to declare which method to invoke on the order id generator bean we must prepend **.methodName** such as below where we invoke the **generateId** method.

```
from("direct:order")
  .transform().simple("OrderId: ${bean:orderIdGenerator.generateId}")
  .to("mock:reply");
```

We can use the **?method=methodName** option that we are familiar with the Bean component itself:

```
from("direct:order")
  .transform().simple("OrderId: ${bean:orderIdGenerator?method=generateId}")
  .to("mock:reply");
```

You can also convert the body to a given type, for example to ensure that it is a String you can do:

```
<transform>
  <simple>Hello ${bodyAs(String)} how are you?</simple>
</transform>
```

There are a few types which have a shorthand notation, so we can use **String** instead of **java.lang.String**. These are: **byte[]**, **String**, **Integer**, **Long**. All other types must use their FQN name, e.g. **org.w3c.dom.Document**.

It is also possible to lookup a value from a header **Map**:

```
<transform>
  <simple>The gold value is ${header.type[gold]}</simple>
</transform>
```

In the code above we lookup the header with name **type** and regard it as a **java.util.Map** and we then lookup with the key **gold** and return the value. If the header is not convertible to Map an exception is thrown. If the header with name **type** does not exist **null** is returned.

You can nest functions, such as shown below:

```
<setHeader name="myHeader">
  <simple>${properties:${header.someKey}}</simple>
</setHeader>
```

74.6. SETTING RESULT TYPE

You can now provide a result type to the **Simple** expression, which means the result of the evaluation will be converted to the desired type. This is most usable to define types such as booleans, integers, etc.

For example to set a header as a boolean type you can do:

```
.setHeader("cool", simple("true", Boolean.class))
```

And in XML DSL

```
<setHeader name="cool">
  <!-- use resultType to indicate that the type should be a java.lang.Boolean -->
  <simple resultType="java.lang.Boolean">true</simple>
</setHeader>
```

74.7. USING NEW LINES OR TABS IN XML DSLS

It is easier to specify new lines or tabs in XML DSLs as you can escape the value now

```
<transform>
  <simple>The following text\nis on a new line</simple>
</transform>
```

74.8. LEADING AND TRAILING WHITESPACE HANDLING

The **trim** attribute of the expression can be used to control whether the leading and trailing whitespace characters are removed or preserved. The default value is **true**, which removes the whitespace characters.

```
<setBody>
  <simple trim="false">You get some trailing whitespace characters. </simple>
</setBody>
```

74.9. LOADING SCRIPT FROM EXTERNAL RESOURCE

You can externalize the script and have Camel load it from a resource such as **"classpath:"**, **"file:"**, or **"http:"**. This is done using the following syntax: **"resource:scheme:location"**, e.g. to refer to a file on the classpath you can do:

```
.setHeader("myHeader").simple("resource:classpath:mysimple.txt")
```

74.10. SPRING BOOT AUTO-CONFIGURATION

When using simple with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	<code>true</code>	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	<code>on-demand</code>	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.user-name</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.kubernetes.service-discovery.lookup</code>	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as name.namespace.svc.dnsDomain. When using dnssrv the service name is resolved with SRV query for svc... When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
<code>camel.cloud.kubernetes.service-discovery.master-url</code>	Sets the URL to the master when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.namespace</code>	Sets the namespace to use. Will by default use namespace from the ENV variable KUBERNETES_MASTER.		String
<code>camel.cloud.kubernetes.service-discovery.oauth-token</code>	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.password</code>	Sets the password for authentication when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-name</code>	Sets the Port Name to use for DNS/DN SSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.port-protocol</code>	Sets the Port Protocol to use for DNS/DN SSRV lookup.		String
<code>camel.cloud.kubernetes.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer

Name	Description	Default	Type
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map
<code>camel.hystrix.core-pool-size</code>	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
<code>camel.hystrix.enabled</code>	Enable the component.	true	Boolean
<code>camel.hystrix.execution-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
<code>camel.hystrix.execution-isolation-strategy</code>	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String

Name	Description	Default	Type
<code>camel.hystrix.execution-isolation-thread-interrupt-on-timeout</code>	Whether the execution thread should attempt an interrupt (using <code>Future#cancel</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
<code>camel.hystrix.execution-timeout-enabled</code>	Whether the timeout mechanism is enabled for this command.	true	Boolean
<code>camel.hystrix.execution-timeout-in-milliseconds</code>	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
<code>camel.hystrix.fallback-enabled</code>	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
<code>camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests</code>	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer
<code>camel.hystrix.group-key</code>	Sets the group key to use. The default value is <code>CamelHystrix</code> .	CamelHystrix	String
<code>camel.hystrix.keep-alive-time</code>	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long, TimeUnit)</code> .	1	Integer
<code>camel.hystrix.max-queue-size</code>	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> . This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
<code>camel.hystrix.maximum-size</code>	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-health-snapshot-interval-in-milliseconds</code>	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
<code>camel.hystrix.metrics-rolling-percentile-bucket-size</code>	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-percentile-enabled</code>	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
<code>camel.hystrix.metrics-rolling-percentile-window-buckets</code>	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer

Name	Description	Default	Type
<code>camel.hystrix.request-log-enabled</code>	Whether HystrixCommand execution and events should be logged to HystrixRequestLog.	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as groupKey has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into HystrixRollingNumber inside each HystrixThreadPoolMetrics instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into HystrixRollingNumber inside each HystrixThreadPoolMetrics instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.csimple.enabled</code>	Whether to enable auto configuration of the csimple language. This is enabled by default.		Boolean
<code>camel.language.csimple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean

Name	Description	Default	Type
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the CircuitBreaker can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the CircuitBreaker will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer

Name	Description	Default	Type
camel.resilience4j.sliding-window-type	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If slidingWindowType is COUNT_BASED, the last slidingWindowSize calls are recorded and aggregated. If slidingWindowType is TIME_BASED, the calls of the last slidingWindowSize seconds are recorded and aggregated. Default slidingWindowType is COUNT_BASED.	COUNT_BASED	String
camel.resilience4j.slow-call-duration-threshold	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
camel.resilience4j.slow-call-rate-threshold	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than slowCallDurationThreshold Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than slowCallDurationThreshold.		Float
camel.resilience4j.wait-duration-in-open-state	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
camel.resilience4j.writable-stack-trace-enabled	Enables writable stack traces. When set to false, Exception.getStackTrace returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
camel.rest.api-component	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a org.apache.camel.spi.RestApiProcessorFactory is registered in the registry. If either one is found, then that is being used.		String

Name	Description	Default	Type
camel.rest.api-context-path	Sets a leading API context-path the REST API services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path.		String
camel.rest.api-context-route-id	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
camel.rest.api-host	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
camel.rest.api-property	Allows to configure as many additional properties for the api documentation (swagger). For example set property api.title to my cool stuff.		Map
camel.rest.api-vendor-extension	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean
camel.rest.binding-mode	Sets the binding mode to use. The default value is off.		RestBindingMode
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String

Name	Description	Default	Type
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBElement is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String

Name	Description	Default	Type
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
camel.rest.skip-binding-on-error-code	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean
camel.rest.use-x-forward-headers	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
camel.rest.xml-data-format	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String

Name	Description	Default	Type
camel.rest.api-context-id-pattern	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
camel.rest.api-context-listing	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 75. TOKENIZE

The tokenizer language is a built-in language in **camel-core**, which is most often used with the [Split](#) EIP to split a message using a token-based strategy.

The tokenizer language is intended to tokenize text documents using a specified delimiter pattern. It can also be used to tokenize XML documents with some limited capability. For a truly XML-aware tokenization, the use of the [XML Tokenize](#) language is recommended as it offers a faster, more efficient tokenization specifically for XML documents.

75.1. TOKENIZE OPTIONS

The Tokenize language supports 11 options, which are listed below.

Name	Default	Java Type	Description
token		String	Required The (start) token to use as tokenizer, for example you can use the new line token. You can use simple language as the token to support dynamic tokens.
endToken		String	The end token to use as tokenizer if using start/end token pairs. You can use simple language as the token to support dynamic tokens.
inheritNamespace TagName		String	To inherit namespaces from a root/parent tag name when using XML You can use simple language as the tag name to support dynamic names.
headerName		String	Name of header to tokenize instead of using the message body.
regex		Boolean	If the token is a regular expression pattern. The default value is false.
xml		Boolean	Whether the input is XML messages. This option must be set to true if working with XML payloads.
includeTokens		Boolean	Whether to include the tokens in the parts when using pairs The default value is false.
group		String	To group N parts together, for example to split big files into chunks of 1000 lines. You can use simple language as the group to support dynamic group sizes.
groupDelimiter		String	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.
skipFirst		Boolean	To skip the very first element.

Name	Default	Java Type	Description
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

75.2. EXAMPLE

The following example shows how to take a request from the `direct:a` endpoint then split it into pieces using an [Expression](#), then forward each piece to `direct:b`:

```
<route>
  <from uri="direct:a"/>
  <split>
    <tokenize token="\n"/>
    <to uri="direct:b"/>
  </split>
</route>
```

And in Java DSL:

```
from("direct:a")
  .split(body().tokenize("\n"))
  .to("direct:b");
```

75.3. SEE ALSO

For more examples see [Split](#) EIP.

75.4. SPRING BOOT AUTO-CONFIGURATION

When using `tokenize` with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-core-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 147 options, which are listed below.

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.acl-token</code>	Sets the ACL token to be used with Consul.		String
<code>camel.cloud.consul.service-discovery.block-seconds</code>	The seconds to wait for a watch event, default 10 seconds.	10	Integer
<code>camel.cloud.consul.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.consul.service-discovery.connect-timeout-millis</code>	Connect timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.datacenter</code>	The data center.		String
<code>camel.cloud.consul.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.consul.service-discovery.password</code>	Sets the password to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.consul.service-discovery.read-timeout-millis</code>	Read timeout for OkHttpClient.		Long
<code>camel.cloud.consul.service-discovery.url</code>	The Consul agent URL.		String

Name	Description	Default	Type
<code>camel.cloud.consul.service-discovery.username</code>	Sets the username to be used for basic authentication.		String
<code>camel.cloud.consul.service-discovery.write-timeout-millis</code>	Write timeout for OkHttpClient.		Long
<code>camel.cloud.dns.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.dns.service-discovery.domain</code>	The domain name;.		String
<code>camel.cloud.dns.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.dns.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.dns.service-discovery.proto</code>	The transport protocol of the desired service.	<code>_tcp</code>	String
<code>camel.cloud.etcd.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.etcd.service-discovery.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.etcd.service-discovery.password</code>	The password to use for basic authentication.		String

Name	Description	Default	Type
<code>camel.cloud.etcd.service-discovery.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.etcd.service-discovery.service-path</code>	The path to look for for service discovery.	<code>/services/</code>	String
<code>camel.cloud.etcd.service-discovery.timeout</code>	To set the maximum time an action could take to complete.		Long
<code>camel.cloud.etcd.service-discovery.type</code>	To set the discovery type, valid values are on-demand and watch.	on-demand	String
<code>camel.cloud.etcd.service-discovery.uris</code>	The URIs the client can connect to.		String
<code>camel.cloud.etcd.service-discovery.username</code>	The user name to use for basic authentication.		String
<code>camel.cloud.kubernetes.service-discovery.api-version</code>	Sets the API version when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-data</code>	Sets the Certificate Authority data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.ca-cert-file</code>	Sets the Certificate Authority data that are loaded from the file when using client lookup.		String

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.client-cert-data</code>	Sets the Client Certificate data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-cert-file</code>	Sets the Client Certificate data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-algo</code>	Sets the Client Keystore algorithm, such as RSA when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-data</code>	Sets the Client Keystore data when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-file</code>	Sets the Client Keystore data that are loaded from the file when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.client-key-passphrase</code>	Sets the Client Keystore passphrase when using client lookup.		String
<code>camel.cloud.kubernetes.service-discovery.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.kubernetes.service-discovery.dns-domain</code>	Sets the DNS domain to use for DNS lookup.		String
<code>camel.cloud.kubernetes.service-discovery.enabled</code>	Enable the component.	true	Boolean

Name	Description	Default	Type
camel.cloud.kubernetes.service-discovery.lookup	How to perform service lookup. Possible values: client, dns, environment. When using client, then the client queries the kubernetes master to obtain a list of active pods that provides the service, and then random (or round robin) select a pod. When using dns the service name is resolved as name.namespace.svc.dnsDomain. When using dnssrv the service name is resolved with SRV query for svc... When using environment then environment variables are used to lookup the service. By default environment is used.	environment	String
camel.cloud.kubernetes.service-discovery.master-url	Sets the URL to the master when using client lookup.		String
camel.cloud.kubernetes.service-discovery.namespace	Sets the namespace to use. Will by default use namespace from the ENV variable KUBERNETES_MASTER.		String
camel.cloud.kubernetes.service-discovery.oauth-token	Sets the OAUTH token for authentication (instead of username/password) when using client lookup.		String
camel.cloud.kubernetes.service-discovery.password	Sets the password for authentication when using client lookup.		String
camel.cloud.kubernetes.service-discovery.port-name	Sets the Port Name to use for DNS/DNSSRV lookup.		String
camel.cloud.kubernetes.service-discovery.port-protocol	Sets the Port Protocol to use for DNS/DNSSRV lookup.		String
camel.cloud.kubernetes.service-discovery.properties	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in com.netflix.client.config.CommonClientConfigKey.		Map

Name	Description	Default	Type
<code>camel.cloud.kubernetes.service-discovery.trust-certs</code>	Sets whether to turn on trust certificate check when using client lookup.	false	Boolean
<code>camel.cloud.kubernetes.service-discovery.username</code>	Sets the username for authentication when using client lookup.		String
<code>camel.cloud.ribbon.load-balancer.client-name</code>	Sets the Ribbon client name.		String
<code>camel.cloud.ribbon.load-balancer.configurations</code>	Define additional configuration definitions.		Map
<code>camel.cloud.ribbon.load-balancer.enabled</code>	Enable the component.	true	Boolean
<code>camel.cloud.ribbon.load-balancer.namespace</code>	The namespace.		String
<code>camel.cloud.ribbon.load-balancer.password</code>	The password.		String
<code>camel.cloud.ribbon.load-balancer.properties</code>	Set client properties to use. These properties are specific to what service call implementation are in use. For example if using ribbon, then the client properties are define in <code>com.netflix.client.config.CommonClientConfigKey</code> .		Map
<code>camel.cloud.ribbon.load-balancer.username</code>	The username.		String

Name	Description	Default	Type
<code>camel.hystrix.allow-maximum-size-to-diverge-from-core-size</code>	Allows the configuration for <code>maximumSize</code> to take effect. That value can then be equal to, or higher, than <code>coreSize</code> .	false	Boolean
<code>camel.hystrix.circuit-breaker-enabled</code>	Whether to use a <code>HystrixCircuitBreaker</code> or not. If false no circuit-breaker logic will be used and all requests permitted. This is similar in effect to <code>circuitBreakerForceClosed()</code> except that continues tracking metrics and knowing whether it should be open/closed, this property results in not even instantiating a circuit-breaker.	true	Boolean
<code>camel.hystrix.circuit-breaker-error-threshold-percentage</code>	Error percentage threshold (as whole number such as 50) at which point the circuit breaker will trip open and reject requests. It will stay tripped for the duration defined in <code>circuitBreakerSleepWindowInMilliseconds</code> ; The error percentage this is compared against comes from <code>HystrixCommandMetrics.getHealthCounts()</code> .	50	Integer
<code>camel.hystrix.circuit-breaker-force-closed</code>	If true the <code>HystrixCircuitBreaker#allowRequest()</code> will always return true to allow requests regardless of the error percentage from <code>HystrixCommandMetrics.getHealthCounts()</code> . The <code>circuitBreakerForceOpen()</code> property takes precedence so if it set to true this property does nothing.	false	Boolean
<code>camel.hystrix.circuit-breaker-force-open</code>	If true the <code>HystrixCircuitBreaker.allowRequest()</code> will always return false, causing the circuit to be open (tripped) and reject all requests. This property takes precedence over <code>circuitBreakerForceClosed()</code> ;	false	Boolean
<code>camel.hystrix.circuit-breaker-request-volume-threshold</code>	Minimum number of requests in the <code>metricsRollingStatisticalWindowInMilliseconds()</code> that must exist before the <code>HystrixCircuitBreaker</code> will trip. If below this number the circuit will not trip regardless of error percentage.	20	Integer
<code>camel.hystrix.circuit-breaker-sleep-window-in-milliseconds</code>	The time in milliseconds after a <code>HystrixCircuitBreaker</code> trips open that it should wait before trying requests again.	5000	Integer
<code>camel.hystrix.configurations</code>	Define additional configuration definitions.		Map

Name	Description	Default	Type
camel.hystrix.core-pool-size	Core thread-pool size that gets passed to <code>java.util.concurrent.ThreadPoolExecutor#setCorePoolSize(int)</code> .	10	Integer
camel.hystrix.enabled	Enable the component.	true	Boolean
camel.hystrix.execution-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.run()</code> . Requests beyond the concurrent limit will be rejected. Applicable only when <code>executionIsolationStrategy == SEMAPHORE</code> .	20	Integer
camel.hystrix.execution-isolation-strategy	What isolation strategy <code>HystrixCommand.run()</code> will be executed with. If <code>THREAD</code> then it will be executed on a separate thread and concurrent requests limited by the number of threads in the thread-pool. If <code>SEMAPHORE</code> then it will be executed on the calling thread and concurrent requests limited by the semaphore count.	THREAD	String
camel.hystrix.execution-isolation-thread-interrupt-on-timeout	Whether the execution thread should attempt an interrupt (using <code>Future#cancel()</code>) when a thread times out. Applicable only when <code>executionIsolationStrategy() == THREAD</code> .	true	Boolean
camel.hystrix.execution-timeout-enabled	Whether the timeout mechanism is enabled for this command.	true	Boolean
camel.hystrix.execution-timeout-in-milliseconds	Time in milliseconds at which point the command will timeout and halt execution. If <code>executionIsolationThreadInterruptOnTimeout == true</code> and the command is thread-isolated, the executing thread will be interrupted. If the command is semaphore-isolated and a <code>HystrixObservableCommand</code> , that command will get unsubscribed.	1000	Integer
camel.hystrix.fallback-enabled	Whether <code>HystrixCommand.getFallback()</code> should be attempted when failure occurs.	true	Boolean
camel.hystrix.fallback-isolation-semaphore-max-concurrent-requests	Number of concurrent requests permitted to <code>HystrixCommand.getFallback()</code> . Requests beyond the concurrent limit will fail-fast and not attempt retrieving a fallback.	10	Integer

Name	Description	Default	Type
camel.hystrix.group-key	Sets the group key to use. The default value is CamelHystrix.	Camel Hystrix	String
camel.hystrix.keep-alive-time	Keep-alive time in minutes that gets passed to <code>ThreadPoolExecutor#setKeepAliveTime(long,TimeUnit)</code> .	1	Integer
camel.hystrix.max-queue-size	Max queue size that gets passed to <code>BlockingQueue</code> in <code>HystrixConcurrencyStrategy.getBlockingQueue(int)</code> This should only affect the instantiation of a threadpool - it is not eligible to change a queue size on the fly. For that, use <code>queueSizeRejectionThreshold()</code> .	-1	Integer
camel.hystrix.maximum-size	Maximum thread-pool size that gets passed to <code>ThreadPoolExecutor#setMaximumPoolSize(int)</code> . This is the maximum amount of concurrency that can be supported without starting to reject <code>HystrixCommands</code> . Please note that this setting only takes effect if you also set <code>allowMaximumSizeToDivergeFromCoreSize</code> .	10	Integer
camel.hystrix.metrics-health-snapshot-interval-in-milliseconds	Time in milliseconds to wait between allowing health snapshots to be taken that calculate success and error percentages and affect <code>HystrixCircuitBreaker.isOpen()</code> status. On high-volume circuits the continual calculation of error percentage can become CPU intensive thus this controls how often it is calculated.	500	Integer
camel.hystrix.metrics-rolling-percentile-bucket-size	Maximum number of values stored in each bucket of the rolling percentile. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
camel.hystrix.metrics-rolling-percentile-enabled	Whether percentile metrics should be captured using <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	true	Boolean
camel.hystrix.metrics-rolling-percentile-window-buckets	Number of buckets the rolling percentile window is broken into. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	6	Integer

Name	Description	Default	Type
<code>camel.hystrix.metrics-rolling-percentile-window-in-milliseconds</code>	Duration of percentile rolling window in milliseconds. This is passed into <code>HystrixRollingPercentile</code> inside <code>HystrixCommandMetrics</code> .	10000	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside <code>HystrixCommandMetrics</code> .	10	Integer
<code>camel.hystrix.metrics-rolling-statistical-window-in-milliseconds</code>	This property sets the duration of the statistical rolling window, in milliseconds. This is how long metrics are kept for the thread pool. The window is divided into buckets and rolls by those increments.	10000	Integer
<code>camel.hystrix.queue-size-rejection-threshold</code>	Queue size rejection threshold is an artificial max size at which rejections will occur even if <code>maxQueueSize</code> has not been reached. This is done because the <code>maxQueueSize</code> of a <code>BlockingQueue</code> can not be dynamically changed and we want to support dynamically changing the queue size that affects rejections. This is used by <code>HystrixCommand</code> when queuing a thread for execution.	5	Integer
<code>camel.hystrix.request-log-enabled</code>	Whether <code>HystrixCommand</code> execution and events should be logged to <code>HystrixRequestLog</code> .	true	Boolean
<code>camel.hystrix.thread-pool-key</code>	Sets the thread pool key to use. Will by default use the same value as <code>groupKey</code> has been configured to use.	Camel Hystrix	String
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-buckets</code>	Number of buckets the rolling statistical window is broken into. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10	Integer
<code>camel.hystrix.thread-pool-rolling-number-statistical-window-in-milliseconds</code>	Duration of statistical rolling window in milliseconds. This is passed into <code>HystrixRollingNumber</code> inside each <code>HystrixThreadPoolMetrics</code> instance.	10000	Integer
<code>camel.language.constant.enabled</code>	Whether to enable auto configuration of the constant language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.constant.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.exchangeproperty.enabled</code>	Whether to enable auto configuration of the exchangeProperty language. This is enabled by default.		Boolean
<code>camel.language.exchangeproperty.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.file.enabled</code>	Whether to enable auto configuration of the file language. This is enabled by default.		Boolean
<code>camel.language.file.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.header.enabled</code>	Whether to enable auto configuration of the header language. This is enabled by default.		Boolean
<code>camel.language.header.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.ref.enabled</code>	Whether to enable auto configuration of the ref language. This is enabled by default.		Boolean
<code>camel.language.ref.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.simple.enabled</code>	Whether to enable auto configuration of the simple language. This is enabled by default.		Boolean
<code>camel.language.simple.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.language.tokenize.enabled</code>	Whether to enable auto configuration of the tokenize language. This is enabled by default.		Boolean

Name	Description	Default	Type
<code>camel.language.tokenize.group-delimiter</code>	Sets the delimiter to use when grouping. If this has not been set then token will be used as the delimiter.		String
<code>camel.language.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean
<code>camel.resilience4j.automatic-transition-from-open-to-half-open-enabled</code>	Enables automatic transition from OPEN to HALF_OPEN state once the <code>waitDurationInOpenState</code> has passed.	false	Boolean
<code>camel.resilience4j.circuit-breaker-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreaker</code> instance to lookup and use from the registry. When using this, then any other circuit breaker options are not in use.		String
<code>camel.resilience4j.config-ref</code>	Refers to an existing <code>io.github.resilience4j.circuitbreaker.CircuitBreakerConfig</code> instance to lookup and use from the registry.		String
<code>camel.resilience4j.configurations</code>	Define additional configuration definitions.		Map
<code>camel.resilience4j.enabled</code>	Enable the component.	true	Boolean
<code>camel.resilience4j.failure-rate-threshold</code>	Configures the failure rate threshold in percentage. If the failure rate is equal or greater than the threshold the <code>CircuitBreaker</code> transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 50 percentage.		Float
<code>camel.resilience4j.minimum-number-of-calls</code>	Configures the minimum number of calls which are required (per sliding window period) before the <code>CircuitBreaker</code> can calculate the error rate. For example, if <code>minimumNumberOfCalls</code> is 10, then at least 10 calls must be recorded, before the failure rate can be calculated. If only 9 calls have been recorded the <code>CircuitBreaker</code> will not transition to open even if all 9 calls have failed. Default <code>minimumNumberOfCalls</code> is 100.	100	Integer

Name	Description	Default	Type
<code>camel.resilience4j.permitted-number-of-calls-in-half-open-state</code>	Configures the number of permitted calls when the CircuitBreaker is half open. The size must be greater than 0. Default size is 10.	10	Integer
<code>camel.resilience4j.sliding-window-size</code>	Configures the size of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. <code>slidingWindowSize</code> configures the size of the sliding window. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. The <code>slidingWindowSize</code> must be greater than 0. The <code>minimumNumberOfCalls</code> must be greater than 0. If the <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the <code>minimumNumberOfCalls</code> cannot be greater than <code>slidingWindowSize</code> . If the <code>slidingWindowType</code> is <code>TIME_BASED</code> , you can pick whatever you want. Default <code>slidingWindowSize</code> is 100.	100	Integer
<code>camel.resilience4j.sliding-window-type</code>	Configures the type of the sliding window which is used to record the outcome of calls when the CircuitBreaker is closed. Sliding window can either be count-based or time-based. If <code>slidingWindowType</code> is <code>COUNT_BASED</code> , the last <code>slidingWindowSize</code> calls are recorded and aggregated. If <code>slidingWindowType</code> is <code>TIME_BASED</code> , the calls of the last <code>slidingWindowSize</code> seconds are recorded and aggregated. Default <code>slidingWindowType</code> is <code>COUNT_BASED</code> .	<code>COUNT_BASED</code>	String
<code>camel.resilience4j.slow-call-duration-threshold</code>	Configures the duration threshold (seconds) above which calls are considered as slow and increase the slow calls percentage. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.slow-call-rate-threshold</code>	Configures a threshold in percentage. The CircuitBreaker considers a call as slow when the call duration is greater than <code>slowCallDurationThreshold</code> Duration. When the percentage of slow calls is equal or greater the threshold, the CircuitBreaker transitions to open and starts short-circuiting calls. The threshold must be greater than 0 and not greater than 100. Default value is 100 percentage which means that all recorded calls must be slower than <code>slowCallDurationThreshold</code> .		Float

Name	Description	Default	Type
<code>camel.resilience4j.wait-duration-in-open-state</code>	Configures the wait duration (in seconds) which specifies how long the CircuitBreaker should stay open, before it switches to half open. Default value is 60 seconds.	60	Integer
<code>camel.resilience4j.writable-stack-trace-enabled</code>	Enables writable stack traces. When set to false, <code>Exception.getStackTrace</code> returns a zero length array. This may be used to reduce log spam when the circuit breaker is open as the cause of the exceptions is already known (the circuit breaker is short-circuiting calls).	true	Boolean
<code>camel.rest.api-component</code>	The name of the Camel component to use as the REST API (such as swagger) If no API Component has been explicit configured, then Camel will lookup if there is a Camel component responsible for servicing and generating the REST API documentation, or if a <code>org.apache.camel.spi.RestApiProcessorFactory</code> is registered in the registry. If either one is found, then that is being used.		String
<code>camel.rest.api-context-path</code>	Sets a leading API context-path the REST API services will be using. This can be used when using components such as <code>camel-servlet</code> where the deployed web application is deployed using a context-path.		String
<code>camel.rest.api-context-route-id</code>	Sets the route id to use for the route that services the REST API. The route will by default use an auto assigned route id.		String
<code>camel.rest.api-host</code>	To use an specific hostname for the API documentation (eg swagger) This can be used to override the generated host with this configured hostname.		String
<code>camel.rest.api-property</code>	Allows to configure as many additional properties for the api documentation (swagger). For example set property <code>api.title</code> to my cool stuff.		Map
<code>camel.rest.api-vendor-extension</code>	Whether vendor extension is enabled in the Rest APIs. If enabled then Camel will include additional information as vendor extension (eg keys starting with x-) such as route ids, class names etc. Not all 3rd party API gateways and tools supports vendor-extensions when importing your API docs.	false	Boolean
<code>camel.rest.binding-mode</code>	Sets the binding mode to use. The default value is off.		RestBindingMode

Name	Description	Default	Type
camel.rest.client-request-validation	Whether to enable validation of the client request to check whether the Content-Type and Accept headers from the client is supported by the Rest-DSL configuration of its consumes/produces settings. This can be turned on, to enable this check. In case of validation error, then HTTP Status codes 415 or 406 is returned. The default value is false.	false	Boolean
camel.rest.component	The Camel Rest component to use for the REST transport (consumer), such as netty-http, jetty, servlet, undertow. If no component has been explicit configured, then Camel will lookup if there is a Camel component that integrates with the Rest DSL, or if a org.apache.camel.spi.RestConsumerFactory is registered in the registry. If either one is found, then that is being used.		String
camel.rest.component-property	Allows to configure as many additional properties for the rest component in use.		Map
camel.rest.consumer-property	Allows to configure as many additional properties for the rest consumer in use.		Map
camel.rest.context-path	Sets a leading context-path the REST services will be using. This can be used when using components such as camel-servlet where the deployed web application is deployed using a context-path. Or for components such as camel-jetty or camel-netty-http that includes a HTTP server.		String
camel.rest.cors-headers	Allows to configure custom CORS headers.		Map
camel.rest.data-format-property	Allows to configure as many additional properties for the data formats in use. For example set property prettyPrint to true to have json outputted in pretty mode. The properties can be prefixed to denote the option is only for either JSON or XML and for either the IN or the OUT. The prefixes are: json.in. json.out. xml.in. xml.out. For example a key with value xml.out.mustBeJAXBELEMENT is only for the XML data format for the outgoing. A key without a prefix is a common key for all situations.		Map
camel.rest.enable-cors	Whether to enable CORS headers in the HTTP response. The default value is false.	false	Boolean

Name	Description	Default	Type
camel.rest.endpoint-property	Allows to configure as many additional properties for the rest endpoint in use.		Map
camel.rest.host	The hostname to use for exposing the REST service.		String
camel.rest.hostname-resolver	If no hostname has been explicit configured, then this resolver is used to compute the hostname the REST service will be using.		RestHostNameResolver
camel.rest.json-data-format	Name of specific json data format to use. By default json-jackson will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
camel.rest.port	The port number to use for exposing the REST service. Notice if you use servlet component then the port number configured here does not apply, as the port number in use is the actual port number the servlet component is using. eg if using Apache Tomcat its the tomcat http port, if using Apache Karaf its the HTTP service in Karaf that uses port 8181 by default etc. Though in those situations setting the port number here, allows tooling and JMX to know the port number, so its recommended to set the port number to the number that the servlet engine uses.		String
camel.rest.producer-api-doc	Sets the location of the api document (swagger api) the REST producer will use to validate the REST uri and query parameters are valid accordingly to the api document. This requires adding camel-swagger-java to the classpath, and any miss configuration will let Camel fail on startup and report the error(s). The location of the api document is loaded from classpath by default, but you can use file: or http: to refer to resources to load from file or http url.		String
camel.rest.producer-component	Sets the name of the Camel component to use as the REST producer.		String
camel.rest.scheme	The scheme to use for exposing the REST service. Usually http or https is supported. The default value is http.		String
camel.rest.skip-binding-on-error-code	Whether to skip binding on output if there is a custom HTTP error code header. This allows to build custom error messages that do not bind to json / xml etc, as success messages otherwise will do.	false	Boolean

Name	Description	Default	Type
<code>camel.rest.use-x-forward-headers</code>	Whether to use X-Forward headers for Host and related setting. The default value is true.	true	Boolean
<code>camel.rest.xml-data-format</code>	Name of specific XML data format to use. By default jaxb will be used. Important: This option is only for setting a custom name of the data format, not to refer to an existing data format instance.		String
<code>camel.rest.api-context-id-pattern</code>	Deprecated Sets an CamelContext id pattern to only allow Rest APIs from rest services within CamelContext's which name matches the pattern. The pattern name refers to the CamelContext name, to match on the current CamelContext only. For any other value, the pattern uses the rules from <code>PatternHelper#matchPattern(String,String)</code> .		String
<code>camel.rest.api-context-listing</code>	Deprecated Sets whether listing of all available CamelContext's with REST services in the JVM is enabled. If enabled it allows to discover these contexts, if false then only the current CamelContext is in use.	false	Boolean

CHAPTER 76. XML TOKENIZE

The XML Tokenize language is a built-in language in **camel-xml-jaxp**, which is a truly XML-aware tokenizer that can be used with the Split EIP as the conventional [Tokenize](#) to efficiently and effectively tokenize XML documents..

XML Tokenize is capable of not only recognizing XML namespaces and hierarchical structures of the document but also more efficiently tokenizing XML documents than the conventional [Tokenize](#) language.

Additional dependency

In order to use this component, an additional dependency is required as follows:

```
<dependency>
  <groupId>org.codehaus.woodstox</groupId>
  <artifactId>woodstox-core-asl</artifactId>
  <version>4.4.1</version>
</dependency>
```

or

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-stax-starter</artifactId>
</dependency>
```

76.1. XML TOKENIZER OPTIONS

The XML Tokenize language supports 4 options, which are listed below.

Name	Default	Java Type	Description
headerName		String	Name of header to tokenize instead of using the message body.
mode		Enum	<p>The extraction mode. The available extraction modes are: i - injecting the contextual namespace bindings into the extracted token (default) w - wrapping the extracted token in its ancestor context u - unwrapping the extracted token to its child content t - extracting the text content of the specified element.</p> <p>Enum values:</p> <ul style="list-style-type: none"> ● i ● w ● u ● t

Name	Default	Java Type	Description
group		Integer	To group N parts together.
trim		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

76.2. EXAMPLE

See [Split EIP](#) which has examples using the XML Tokenize language.

76.3. SPRING BOOT AUTO-CONFIGURATION

When using xtokenize with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-xml-jaxp-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 3 options, which are listed below.

Name	Description	Default	Type
<code>camel.language.xml.tokenize.enabled</code>	Whether to enable auto configuration of the xtokenize language. This is enabled by default.		Boolean
<code>camel.language.xml.tokenize.mode</code>	The extraction mode. The available extraction modes are: i - injecting the contextual namespace bindings into the extracted token (default) w - wrapping the extracted token in its ancestor context u - unwrapping the extracted token to its child content t - extracting the text content of the specified element.		String
<code>camel.language.xml.tokenize.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean

CHAPTER 77. XPATH

Camel supports [XPath](#) to allow an [Expression](#) or [Predicate](#) to be used in the [DSL](#).

For example, you could use XPath to create a predicate in a [Message Filter](#) or as an expression for a [Recipient List](#).

77.1. XPATH LANGUAGE OPTIONS

The XPath language supports 10 options, which are listed below.

Name	Default	Java Type	Description
<code>documentType</code>		String	Name of class for document type The default value is <code>org.w3c.dom.Document</code> .
<code>resultType</code>		Enum	Sets the class name of the result type (type from output) The default result type is <code>NodeSet</code> . Enum values: <ul style="list-style-type: none"> ● NUMBER ● STRING ● BOOLEAN ● NODESET ● NODE
<code>saxon</code>		Boolean	Whether to use Saxon.
<code>factoryRef</code>		String	References to a custom <code>XPathFactory</code> to lookup in the registry.
<code>objectModel</code>		String	The XPath object model to use.
<code>logNamespaces</code>		Boolean	Whether to log namespaces which can assist during troubleshooting.
<code>headerName</code>		String	Name of header to use as input, instead of the message body.

Name	Default	Java Type	Description
<code>threadSafety</code>		Boolean	Whether to enable thread-safety for the returned result of the xpath expression. This applies to when using NODESET as the result type, and the returned set has multiple elements. In this situation there can be thread-safety issues if you process the NODESET concurrently such as from a Camel Splitter EIP in parallel processing mode. This option prevents concurrency issues by doing defensive copies of the nodes. It is recommended to turn this option on if you are using camel-saxon or Saxon in your application. Saxon has thread-safety issues which can be prevented by turning this option on.
<code>preCompile</code>		Boolean	Whether to enable pre-compiling the xpath expression during initialization phase. pre-compile is enabled by default. This can be used to turn off, for example in cases the compilation phase is desired at the starting phase, such as if the application is ahead of time compiled (for example with camel-quarkus) which would then load the xpath factory of the built operating system, and not a JVM runtime.
<code>trim</code>		Boolean	Whether to trim the value to remove leading and trailing whitespaces and line breaks.

77.2. NAMESPACES

You can easily use namespaces with XPath expressions using the **Namespaces** helper class.

77.3. VARIABLES

Variables in XPath is defined in different namespaces. The default namespace is <http://camel.apache.org/schema/spring>.

Namespace URI	Local part	Type	Description
http://camel.apache.org/xml/in/	in	Message	the message
http://camel.apache.org/xml/out/	out	Message	deprecated the output message (do not use)
http://camel.apache.org/xml/function/	functions	Object	Additional functions

Namespace URI	Local part	Type	Description
http://camel.apache.org/xml/variables/environment-variables	env	Object	OS environment variables
http://camel.apache.org/xml/variables/system-properties	system	Object	Java System properties
http://camel.apache.org/xml/variables/exchange-property		Object	the exchange property

Camel will resolve variables according to either:

- namespace given
- no namespace given

77.3.1. Namespace given

If the namespace is given then Camel is instructed exactly what to return. However, when resolving Camel will try to resolve a header with the given local part first, and return it. If the local part has the value **body** then the body is returned instead.

77.3.2. No namespace given

If there is no namespace given then Camel resolves only based on the local part. Camel will try to resolve a variable in the following steps:

- from **variables** that has been set using the **variable(name, value)** fluent builder
- from **message.in.header** if there is a header with the given key
- from **exchange.properties** if there is a property with the given key

77.4. FUNCTIONS

Camel adds the following XPath functions that can be used to access the exchange:

Function	Argument	Type	Description
in:body	none	Object	Will return the message body.
in:header	the header name	Object	Will return the message header.
out:body	none	Object	deprecated Will return the out message body.
out:header	the header name	Object	deprecated Will return the out message header.

Function	Argument	Type	Description
function:properties	key for property	String	To use a .
function:simple	simple expression	Object	To evaluate a language.



NOTE

function:properties and **function:simple** is not supported when the return type is a **NodeSet**, such as when using with a [Split](#) EIP.

Here's an example showing some of these functions in use.

77.4.1. Functions example

If you prefer to configure your routes in your Spring XML file then you can use XPath expressions as follows

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-
    spring.xsd">

  <camelContext id="camel" xmlns="http://activemq.apache.org/camel/schema/spring"
    xmlns:foo="http://example.com/person">
    <route>
      <from uri="activemq:MyQueue"/>
      <filter>
        <xpath>/foo:person[@name='James']</xpath>
        <to uri="mqseries:SomeOtherQueue"/>
      </filter>
    </route>
  </camelContext>
</beans>
```

Notice how we can reuse the namespace prefixes, **foo** in this case, in the XPath expression for easier namespace based XPath expressions.

77.5. STREAM BASED MESSAGE BODIES

If the message body is stream based, which means the input it receives is submitted to Camel as a stream. That means you will only be able to read the content of the stream **once**. So often when you use [XPath](#) as Message Filter or Content Based Router then you need to access the data multiple times, and you should use Stream Caching or convert the message body to a **String** prior which is safe to be re-read multiple times.

```
from("queue:foo").
  filter().xpath("//foo").
  to("queue:bar")
```

```
from("queue:foo").
  choice().xpath("//foo").to("queue:bar").
  otherwise().to("queue:others");
```

77.6. SETTING RESULT TYPE

The XPath expression will return a result type using native XML objects such as **org.w3c.dom.NodeList**. However, many times you want a result type to be a **String**. To do this you have to instruct the XPath which result type to use.

In Java DSL:

```
xpath("/foo:person/@id", String.class)
```

In XML DSL you use the **resultType** attribute to provide the fully qualified classname.

```
<xpath resultType="java.lang.String">/foo:person/@id</xpath>
```



NOTE

Classes from **java.lang** can omit the FQN name, so you can use **resultType="String"**

Using **@XPath** annotation:

```
@XPath(value = "concat('foo-',//order/name/)", resultType = String.class) String name)
```

Where we use the xpath function concat to prefix the order name with **foo-**. In this case we have to specify that we want a **String** as result type, so the concat function works.

77.7. USING XPATH ON HEADERS

Some users may have XML stored in a header. To apply an XPath to a header's value you can do this by defining the 'headerName' attribute.

```
<xpath headerName="invoiceDetails">/invoice/@orderType = 'premium'</xpath>
```

And in Java DSL you specify the headerName as the 2nd parameter as shown:

```
xpath("/invoice/@orderType = 'premium'", "invoiceDetails")
```

77.8. EXAMPLE

Here is a simple example using an XPath expression as a predicate in a [Message Filter](#):

```
from("direct:start")
  .filter().xpath("/person[@name='James']")
  .to("mock:result");
```

And in XML

```

<route>
  <from uri="direct:start"/>
  <filter>
    <xpath>/person[@name='James']</xpath>
    <to uri="mock:result"/>
  </filter>
</route>

```

77.9. USING NAMESPACES

If you have a standard set of namespaces you wish to work with and wish to share them across many XPath expressions you can use the **org.apache.camel.support.builder.Namespaces** when using Java DSL as shown:

```

Namespaces ns = new Namespaces("c", "http://acme.com/cheese");

from("direct:start")
  .filter(xpath("/c:person[@name='James']", ns))
  .to("mock:result");

```

Notice how the namespaces are provided to **xpath** with the **ns** variable that are passed in as the 2nd parameter.

Each namespace is a key=value pair, where the prefix is the key. In the XPath expression then the namespace is used by its prefix, eg:

```

/c:person[@name='James']

```

The namespace builder supports adding multiple namespaces as shown:

```

Namespaces ns = new Namespaces("c", "http://acme.com/cheese")
  .add("w", "http://acme.com/wine")
  .add("b", "http://acme.com/beer");

```

When using namespaces in XML DSL then its different, as you setup the namespaces in the XML root tag (or one of the **camelContext**, **routes**, **route** tags).

In the XML example below we use Spring XML where the namespace is declared in the root tag **beans**, in the line with **xmlns:foo="http://example.com/person"**:

```

<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:foo="http://example.com/person"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd
  ">

  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="direct:start"/>
      <filter>

```

```

    <xpath logNamespaces="true">/foo:person[@name='James']</xpath>
    <to uri="mock:result"/>
  </filter>
</route>
</camelContext>


</beans>

```

This namespace uses **foo** as prefix, so the **<xpath>** expression uses **/foo:** to use this namespace.

77.10. USING @XPath ANNOTATION FOR BEAN INTEGRATION

You can use [Bean Integration](#) to invoke a method on a bean and use various languages such as **@XPath** to extract a value from the message and bind it to a method parameter.



NOTE

The default **@XPath** annotation has SOAP and XML namespaces available.

```

public class Foo {

    @Consume(uri = "activemq:my.queue")
    public void doSomething(@XPath("/person/@name") String name, String xml) {
        // process the inbound message here
    }
}

```

77.11. USING XPathBuilder WITHOUT AN EXCHANGE

You can now use the **org.apache.camel.language.xpath.XPathBuilder** without the need for an **Exchange**. This comes handy if you want to use it as a helper to do custom XPath evaluations.

It requires that you pass in a **CamelContext** since a lot of the moving parts inside the **XPathBuilder** requires access to the Camel [Type Converter](#) and hence why **CamelContext** is needed.

For example, you can do something like this:

```

boolean matches = XPathBuilder.xpath("/foo/bar/@xyz").matches(context, "<foo><bar xyz='cheese'>
</foo>");

```

This will match the given predicate.

You can also evaluate as shown in the following three examples:

```

String name = XPathBuilder.xpath("foo/bar").evaluate(context, "<foo><bar>cheese</bar></foo>",
String.class);
Integer number = XPathBuilder.xpath("foo/bar").evaluate(context, "<foo><bar>123</bar></foo>",
Integer.class);
Boolean bool = XPathBuilder.xpath("foo/bar").evaluate(context, "<foo><bar>>true</bar></foo>",
Boolean.class);

```

Evaluating with a **String** result is a common requirement and make this simpler:

■

```
String name = XPathBuilder.xpath("foo/bar").evaluate(context, "<foo><bar>cheese</bar></foo>");
```

77.12. USING SAXON WITH XPATHBUILDER

You need to add **camel-saxon** as dependency to your project.

It's now easier to use [Saxon](#) with the XPathBuilder which can be done in several ways as shown below

- Using a custom XPathFactory
- Using ObjectModel

77.12.1. Setting a custom XPathFactory using System Property

Camel now supports reading the [JVM system property](#) **javax.xml.xpath.XPathFactory** that can be used to set a custom XPathFactory to use.

This unit test shows how this can be done to use Saxon instead:

Camel will log at **INFO** level if it uses a non default XPathFactory such as:

```
XPathBuilder INFO Using system property
javax.xml.xpath.XPathFactory:http://saxon.sf.net/jaxp/xpath/om with value:
net.sf.saxon.xpath.XPathFactoryImpl when creating XPathFactory
```

To use Apache Xerces you can configure the system property

```
-Djavax.xml.xpath.XPathFactory=org.apache.xpath.jaxp.XPathFactoryImpl
```

77.12.2. Enabling Saxon from XML DSL

Similarly to Java DSL, to enable Saxon from XML DSL you have three options:

Referring to a custom factory:

```
<xpath factoryRef="saxonFactory" resultType="java.lang.String">current-dateTime()</xpath>
```

And declare a bean with the factory:

```
<bean id="saxonFactory" class="net.sf.saxon.xpath.XPathFactoryImpl"/>
```

Specifying the object model:

```
<xpath objectModel="http://saxon.sf.net/jaxp/xpath/om" resultType="java.lang.String">current-dateTime()</xpath>
```

And the recommended approach is to set **saxon=true** as shown:

```
<xpath saxon="true" resultType="java.lang.String">current-dateTime()</xpath>
```

77.13. NAMESPACE AUDITING TO AID DEBUGGING

Many XPath-related issues that users frequently face are linked to the usage of namespaces. You may have some misalignment between the namespaces present in your message, and those that your XPath expression is aware of or referencing. XPath predicates or expressions that are unable to locate the XML elements and attributes due to namespaces issues may simply look like *they are not working*, when in reality all there is to it is a lack of namespace definition.

Namespaces in XML are completely necessary, and while we would love to simplify their usage by implementing some magic or voodoo to wire namespaces automatically, truth is that any action down this path would disagree with the standards and would greatly hinder interoperability.

Therefore, the utmost we can do is assist you in debugging such issues by adding two new features to the XPath Expression Language and are thus accessible from both predicates and expressions.

77.13.1. Logging the Namespace Context of your XPath expression/predicate

Every time a new XPath expression is created in the internal pool, Camel will log the namespace context of the expression under the **org.apache.camel.language.xpath.XPathBuilder** logger. Since Camel represents Namespace Contexts in a hierarchical fashion (parent-child relationships), the entire tree is output in a recursive manner with the following format:

```
[me: {prefix -> namespace}, {prefix -> namespace}], [parent: [me: {prefix -> namespace}, {prefix -> namespace}], [parent: [me: {prefix -> namespace}]]]
```

Any of these options can be used to activate this logging:

- Enable TRACE logging on the **org.apache.camel.language.xpath.XPathBuilder** logger, or some parent logger such as **org.apache.camel** or the root logger
- Enable the **logNamespaces** option as indicated in the following section, in which case the logging will occur on the INFO level

77.13.2. Auditing namespaces

Camel is able to discover and dump all namespaces present on every incoming message before evaluating an XPath expression, providing all the richness of information you need to help you analyse and pinpoint possible namespace issues.

To achieve this, it in turn internally uses another specially tailored XPath expression to extract all namespace mappings that appear in the message, displaying the prefix and the full namespace URI(s) for each individual mapping.

Some points to take into account:

- The implicit XML namespace (**xmlns:xml="http://www.w3.org/XML/1998/namespace"**) is suppressed from the output because it adds no value
- Default namespaces are listed under the **DEFAULT** keyword in the output
- Keep in mind that namespaces can be remapped under different scopes. Think of a top-level 'a' prefix which in inner elements can be assigned a different namespace, or the default namespace changing in inner scopes. For each discovered prefix, all associated URIs are listed.

You can enable this option in Java DSL and XML DSL:

Java DSL:

■

```
XPathBuilder.xpath("/foo:person/@id", String.class).logNamespaces()
```

XML DSL:

```
<xpath logNamespaces="true" resultType="String">/foo:person/@id</xpath>
```

The result of the auditing will be appeared at the INFO level under the **org.apache.camel.language.xpath.XPathBuilder** logger and will look like the following:

```
2012-01-16 13:23:45,878 [stSaxonWithFlag] INFO XPathBuilder - Namespaces discovered in
message:
{xmlns:a=[http://apache.org/camel], DEFAULT=[http://apache.org/default],
xmlns:b=[http://apache.org/camelA, http://apache.org/camelB]}
```

77.14. LOADING SCRIPT FROM EXTERNAL RESOURCE

You can externalize the script and have Camel load it from a resource such as **"classpath:"**, **"file:"**, or **"http:"**. This is done using the following syntax: **"resource:scheme:location"**, eg to refer to a file on the classpath you can do:

```
.setHeader("myHeader").xpath("resource:classpath:myxpath.txt", String.class)
```

77.15. DEPENDENCIES

To use XPath in your camel routes you need to add the dependency on **camel-xpath** which implements the XPath language.

If you use maven, add the following to your **pom.xml**, substituting the version number for the latest version (see the download page for the latest version).

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-xpath</artifactId>
  <version>3.14.5.redhat-00032</version>
</dependency>
```

77.16. SPRING BOOT AUTO-CONFIGURATION

When using xpath with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-xpath-starter</artifactId>
  <version>3.14.5.redhat-00018</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 9 options, which are listed below.

Name	Description	Default	Type
<code>camel.language.xpath.document-type</code>	Name of class for document type The default value is <code>org.w3c.dom.Document</code> .		String
<code>camel.language.xpath.enabled</code>	Whether to enable auto configuration of the xpath language. This is enabled by default.		Boolean
<code>camel.language.xpath.factory-ref</code>	References to a custom XPathFactory to lookup in the registry.		String
<code>camel.language.xpath.log-namespaces</code>	Whether to log namespaces which can assist during troubleshooting.	false	Boolean
<code>camel.language.xpath.object-model</code>	The XPath object model to use.		String
<code>camel.language.xpath.pre-compile</code>	Whether to enable pre-compiling the xpath expression during initialization phase. pre-compile is enabled by default. This can be used to turn off, for example in cases the compilation phase is desired at the starting phase, such as if the application is ahead of time compiled (for example with camel-quarkus) which would then load the xpath factory of the built operating system, and not a JVM runtime.	true	Boolean
<code>camel.language.xpath.saxon</code>	Whether to use Saxon.	false	Boolean
<code>camel.language.xpath.thread-safety</code>	Whether to enable thread-safety for the returned result of the xpath expression. This applies to when using NODESET as the result type, and the returned set has multiple elements. In this situation there can be thread-safety issues if you process the NODESET concurrently such as from a Camel Splitter EIP in parallel processing mode. This option prevents concurrency issues by doing defensive copies of the nodes. It is recommended to turn this option on if you are using camel-saxon or Saxon in your application. Saxon has thread-safety issues which can be prevented by turning this option on.	false	Boolean
<code>camel.language.xpath.trim</code>	Whether to trim the value to remove leading and trailing whitespaces and line breaks.	true	Boolean

CHAPTER 78. OPENAPI JAVA

The Rest DSL can be integrated with the **camel-openapi-java** module which is used for exposing the REST services and their APIs using [OpenApi](#).

Maven users will need to add the following dependency to their **pom.xml** for this component:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-openapi-java</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- use the same version as your Camel core version -->
</dependency>
```

The camel-openapi-java module can be used from the REST components (without the need for servlet)

78.1. USING OPENAPI IN REST-DSL

You can enable the OpenApi api from the rest-dsl by configuring the **apiContextPath** dsl as shown below:

```
public class UserRouteBuilder extends RouteBuilder {
  @Override
  public void configure() throws Exception {
    // configure we want to use servlet as the component for the rest DSL
    // and we enable json binding mode
    restConfiguration().component("netty-http").bindingMode(RestBindingMode.json)
    // and output using pretty print
    .dataFormatProperty("prettyPrint", "true")
    // setup context path and port number that netty will use
    .contextPath("/").port(8080)
    // add OpenApi api-doc out of the box
    .apiContextPath("/api-doc")
    .apiProperty("api.title", "User API").apiProperty("api.version", "1.2.3")
    // and enable CORS
    .apiProperty("cors", "true");

    // this user REST service is json only
    rest("/user").description("User rest service")
    .consumes("application/json").produces("application/json")
    .get("/{id}").description("Find user by id").outType(User.class)
    .param().name("id").type(path).description("The id of the user to
get").dataType("int").endParam()
    .to("bean:userService?method=getUser(${header.id})")
    .put().description("Updates or create a user").type(User.class)
    .param().name("body").type(body).description("The user to update or create").endParam()
    .to("bean:userService?method=updateUser")
    .get("/findAll").description("Find all users").outType(User[].class)
    .to("bean:userService?method=listUsers");
  }
}
```

78.2. OPTIONS

The OpenApi module can be configured using the following options. To configure using a servlet you use the `init-param` as shown above. When configuring directly in the `rest-dsl`, you use the appropriate method, such as `enableCORS`, `host`, `contextPath`, `dsl`. The options with `api.xxx` is configured using `apiProperty` `dsl`.

Option	Type	Description
<code>cors</code>	Boolean	Whether to enable CORS. Notice this only enables CORS for the api browser, and not the actual access to the REST services. Is default false.
<code>openapi.version</code>	String	OpenApi spec version. Is default 3.0.
<code>host</code>	String	To setup the hostname. If not configured camel-openapi-java will calculate the name as localhost based.
<code>schemes</code>	String	The protocol schemes to use. Multiple values can be separated by comma such as "http,https". The default value is "http".
<code>base.path</code>	String	Required: To setup the base path where the REST services is available. The path is relative (eg do not start with http/https) and camel-openapi-java will calculate the absolute base path at runtime, which will be protocol://host:port/context-path/base.path
<code>api.path</code>	String	To setup the path where the API is available (eg /api-docs). The path is relative (eg do not start with http/https) and camel-openapi-java will calculate the absolute base path at runtime, which will be protocol://host:port/context-path/api.path So using relative paths is much easier. See above for an example.
<code>api.version</code>	String	The version of the api. Is default 0.0.0.
<code>api.title</code>	String	The title of the application.
<code>api.description</code>	String	A short description of the application.
<code>api.termsOfService</code>	String	A URL to the Terms of Service of the API.
<code>api.contact.name</code>	String	Name of person or organization to contact
<code>api.contact.email</code>	String	An email to be used for API-related correspondence.
<code>api.contact.url</code>	String	A URL to a website for more contact information.
<code>api.license.name</code>	String	The license name used for the API.
<code>api.license.url</code>	String	A URL to the license used for the API.

78.3. ADDING SECURITY DEFINITIONS IN API DOC

The Rest DSL now supports declaring OpenApi **securityDefinitions** in the generated API document. For example as shown below:

```
rest("/user").tag("dude").description("User rest service")
    // setup security definitions
    .securityDefinitions()
        .oauth2("petstore_auth").authorizationUrl("http://petstore.swagger.io/oauth/dialog").end()
        .apiKey("api_key").withHeader("myHeader").end()
    .end()
    .consumes("application/json").produces("application/json")
```

Here we have setup two security definitions

- OAuth2 - with implicit authorization with the provided url
- Api Key - using an api key that comes from HTTP header named *myHeader*

Then you need to specify on the rest operations which security to use by referring to their key (petstore_auth or api_key).

```
.get("/{id}/{date}").description("Find user by id and date").outType(User.class)
    .security("api_key")
...
.put().description("Updates or create a user").type(User.class)
    .security("petstore_auth", "write:pets,read:pets")
```

Here the get operation is using the Api Key security and the put operation is using OAuth security with permitted scopes of read and write pets.

78.4. JSON OR YAML

The camel-openapi-java module supports both JSON and Yaml out of the box. You can specify in the request url what you want returned by using /openapi.json or /openapi.yaml for either one. If none is specified then the HTTP Accept header is used to detect if json or yaml can be accepted. If either both is accepted or none was set as accepted then json is returned as the default format.

78.5. USEXFORWARDHEADERS AND API URL RESOLUTION

The OpenApi specification allows you to specify the host, port & path that is serving the API. In OpenApi V2 this is done via the **host** field and in OpenAPI V3 it is part of the **servers** field.

By default, the value for these fields is determined by **X-Forwarded** headers, **X-Forwarded-Host** & **X-Forwarded-Proto**.

This can be overridden by disabling the lookup of **X-Forwarded** headers and by specifying your own host, port & scheme on the REST configuration.

```
restConfiguration().component("netty-http")
    .useXForwardHeaders(false)
    .apiProperty("schemes", "https");
    .host("localhost")
    .port(8080);
```

78.6. EXAMPLES

In the Apache Camel distribution we ship the **camel-example-openapi-cdi** and **camel-example-spring-boot-rest-openapi-simple** which demonstrates using this OpenApi component.

78.7. SPRING BOOT AUTO-CONFIGURATION

When using openapi-java with Spring Boot make sure to use the following Maven dependency to have support for auto configuration:

```
<dependency>
  <groupId>org.apache.camel.springboot</groupId>
  <artifactId>camel-openapi-java-starter</artifactId>
  <version>3.14.5.redhat-00032</version>
  <!-- Use your Camel Spring Boot version -->
</dependency>
```

The component supports 1 options, which are listed below.

Name	Description	Default	Type
camel.openapi.enabled	Enables Camel Rest DSL to automatic register its OpenAPI (eg swagger doc) in Spring Boot which allows tooling such as SpringDoc to integrate with Camel.	true	Boolean