Red Hat JBoss Enterprise Application Platform 7.4

Using JBoss EAP XP 4.0.0

For Use with JBoss EAP XP 4.0.0

Last Updated: 2024-02-08

Red Hat JBoss Enterprise Application Platform 7.4 Using JBoss EAP XP 4.0.0

For Use with JBoss EAP XP 4.0.0

Legal Notice

Copyright © 2024 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 document provides general information about using MicroProfile in JBoss EAP XP 4.0.0.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE 7

PROVIDING FEEDBACK ON JBOSS EAP DOCUMENTATION 8

CHAPTER 1. JBOSS EAP XP FOR THE LATEST MICROPROFILE CAPABILITIES 9
- 1.1. ABOUT JBOSS EAP XP 9
- 1.2. JBOSS EAP XP INSTALLATION 9
- 1.3. JBOSS EAP XP MANAGER 10
- 1.4. JBOSS EAP XP MANAGER 4.0 COMMANDS 10
- 1.5. INSTALLING JBOSS EAP XP 4.0.0 ON JBOSS EAP 7.4.X 13
- 1.6. UNINSTALLING JBOSS EAP XP 14
- 1.7. VIEWING THE STATUS OF JBOSS EAP XP 14
- 1.8. ROLLING BACK JBOSS EAP XP AND JBOSS EAP 7.4.X BASE PATCHES 15

CHAPTER 2. UNDERSTAND MICROPROFILE 16
- 2.1. MICROPROFILE CONFIG 16
  - 2.1.1. MicroProfile Config in JBoss EAP 16
  - 2.1.2. MicroProfile Config sources supported in MicroProfile Config 16
- 2.2. MICROPROFILE FAULT TOLERANCE 17
  - 2.2.1. About MicroProfile Fault Tolerance specification 17
  - 2.2.2. MicroProfile Fault Tolerance in JBoss EAP 17
- 2.3. MICROPROFILE HEALTH 18
  - 2.3.1. MicroProfile Health in JBoss EAP 18
- 2.4. MICROPROFILE JWT 19
  - 2.4.1. MicroProfile JWT integration in JBoss EAP 19
  - 2.4.2. Differences between a traditional deployment and an MicroProfile JWT deployment 20
  - 2.4.3. MicroProfile JWT activation in JBoss EAP 20
  - 2.4.4. Limitations of MicroProfile JWT in JBoss EAP 20
- 2.5. MICROPROFILE METRICS 20
  - 2.5.1. MicroProfile Metrics in JBoss EAP 21
- 2.6. MICROPROFILE OPENAPI 21
  - 2.6.1. MicroProfile OpenAPI in JBoss EAP 21
- 2.7. MICROPROFILE OPENTRACING 21
  - 2.7.1. MicroProfile OpenTracing 21
  - 2.7.2. MicroProfile OpenTracing in JBoss EAP 22
- 2.8. MICROPROFILE REST CLIENT 23
  - 2.8.1. MicroProfile REST client 23
  - 2.8.2. The resteasy.original.webapplicationexception.behavior MicroProfile Config property 23
    - Defining the resteasy.original.webapplicationexception.behavior MicroProfile Config property 24
- 2.9. MICROPROFILE REACTIVE MESSAGING 24
  - 2.9.1. MicroProfile reactive messaging 24
  - 2.9.2. MicroProfile reactive messaging connectors 25
    - The Apache Kafka connector and incorporated layers 25
  - 2.9.3. The Apache Kafka event streaming platform 25

CHAPTER 3. ADMINISTER MICROPROFILE IN JBOSS EAP 27
- 3.1. MICROPROFILE OPENTRACING ADMINISTRATION 27
  - 3.1.1. Enabling MicroProfile Open Tracing 27
  - 3.1.2. Removing the microprofile-opentracing-smallrye subsystem 27
  - 3.1.3. Installing Jaeger 27
- 3.2. MICROPROFILE CONFIG CONFIGURATION 28
  - 3.2.1. Adding properties in a ConfigSource management resource 28

3.2.2. Configuring directories as ConfigSources 28
3.2.3. Obtaining ConfigSource from a ConfigSource class 29
3.2.4. Obtaining ConfigSource configuration from a ConfigSourceProvider class 29
3.3. MICROPROFILE FAULT TOLERANCE CONFIGURATION 30
3.3.1. Adding the MicroProfile Fault Tolerance extension 30
3.4. MICROPROFILE HEALTH CONFIGURATION 30
3.4.1. Examining health using the management CLI 31
3.4.2. Examining health using the management console 31
3.4.3. Examining health using the HTTP endpoint 31
3.4.4. Enabling authentication for MicroProfile Health 31
3.4.5. Readiness probes that determine server health and readiness 32
3.4.6. Global status when probes are not defined 33
3.5. MICROPROFILE JWT CONFIGURATION 34
3.5.1. Enabling microprofile-jwt-smallrye subsystem 34
3.6. MICROPROFILE METRICS ADMINISTRATION 34
3.6.1. Metrics available on the management interface 34
3.6.2. Examining metrics using the HTTP endpoint 35
3.6.3. Enabling Authentication for the MicroProfile Metrics HTTP Endpoint 35
3.6.4. Obtaining the request count for a web service 35
3.7. MICROPROFILE OPENAPI ADMINISTRATION 36
3.7.1. Enabling MicroProfile OpenAPI 36
3.7.2. Requesting an MicroProfile OpenAPI document using Accept HTTP header 37
3.7.3. Requesting an MicroProfile OpenAPI document using an HTTP parameter 37
3.7.4. Configuring JBoss EAP to serve a static OpenAPI document 38
3.7.5. Disabling microprofile-openapi-smallrye 38
3.8. MICROPROFILE REACTIVE MESSAGING ADMINISTRATION 39
3.8.1. Configuring the required MicroProfile reactive messaging extension and subsystem for JBoss EAP 39
3.9. STANDALONE SERVER CONFIGURATION 40
3.9.1. Standalone server configuration files 40
3.9.2. Updating standalone configurations with MicroProfile subsystems and extensions 41

CHAPTER 4. DEVELOP MICROPROFILE APPLICATIONS FOR JBOSS EAP 43
4.1. MAVEN AND THE JBOSS EAP MICROPROFILE MAVEN REPOSITORY 43
4.1.1. Downloading the JBoss EAP MicroProfile Maven repository patch as an archive file 43
4.1.2. Applying the JBoss EAP MicroProfile Maven repository patch on your local system 43
4.1.3. Supported JBoss EAP MicroProfile BOM 44
4.1.4. Using the JBoss EAP MicroProfile Maven repository 45
4.2. MICROPROFILE CONFIG DEVELOPMENT 46
4.2.1. Creating a Maven project for MicroProfile Config 46
4.2.2. Using MicroProfile Config property in an application 47
4.3. MICROPROFILE FAULT TOLERANCE APPLICATION DEVELOPMENT 49
4.3.1. Adding the MicroProfile Fault Tolerance extension 49
4.3.2. Configuring Maven project for MicroProfile Fault Tolerance 50
4.3.3. Creating a fault tolerant application 51
4.4. MICROPROFILE HEALTH DEVELOPMENT 54
4.4.1. The custom health check example 54
4.4.2. The @Liveness annotation example 55
4.4.3. The @Readiness annotation example 56
4.4.4. The @Startup annotation example 57
4.5. MICROPROFILE JWT APPLICATION DEVELOPMENT 57
4.5.1. Enabling microprofile-jwt-smallrye subsystem 57
4.5.2. Configuring Maven project for developing JWT applications 57
4.5.3. Creating an application with MicroProfile JWT 58

4.6. MICROPROFILE METRICS DEVELOPMENT 63
4.6.1. Creating an MicroProfile Metrics application 63
4.7. DEVELOPING AN MICROPROFILE OPENAPI APPLICATION 64
4.7.1. Enabling MicroProfile OpenAPI 64
4.7.2. Configuring Maven project for MicroProfile OpenAPI 65
4.7.3. Creating an MicroProfile OpenAPI application 67
4.7.4. Configuring JBoss EAP to serve a static OpenAPI document 71
4.8. MICROPROFILE REST CLIENT DEVELOPMENT 71
4.8.1. A comparison of MicroProfile REST client and Jakarta RESTful Web Services syntaxes 72
4.8.2. Programmatic registration of providers in MicroProfile REST client 73
4.8.3. Declarative registration of providers in MicroProfile REST client 73
4.8.4. Declarative specification of headers in MicroProfile REST client 73
4.8.5. ResponseExceptionMapper in MicroProfile REST client 74
4.8.6. Context dependency injection with MicroProfile REST client 74

CHAPTER 5. BUILD AND RUN MICROSERVICES APPLICATIONS ON THE OPENSHIFT IMAGE FOR JBOSS EAP XP 76
5.1. PREPARING OPENSHIFT FOR APPLICATION DEPLOYMENT 76
5.2. CONFIGURING AUTHENTICATION TO THE RED HAT CONTAINER REGISTRY 77
5.3. IMPORTING THE LATEST OPENSHIFT IMAGESTREAMS AND TEMPLATES FOR JBOSS EAP XP 77
5.4. DEPLOYING A JBOSS EAP XP SOURCE-TO-IMAGE (S2I) APPLICATION ON OPENSHIFT 78
5.5. COMPLETING POST-DEPLOYMENT TASKS FOR JBOSS EAP XP SOURCE-TO-IMAGE (S2I) APPLICATION 80

CHAPTER 6. CAPABILITY TRIMMING 82
6.1. AVAILABLE JBOSS EAP LAYERS 82
6.1.1. Base layers 82
datasources-web-server 82
jaxrs-server 83
cloud-server 83
6.1.2. Decorator layers 83
ejb-lite 84
Jakarta Enterprise Beans 84
ejb-local-cache 84
ejb-dist-cache 84
jdr 84
Jakarta Persistence 85
jpa-distributed 85
Jakarta Server Faces 85
microprofile-platform 85
observability 86
remote-activemq 86
sso 86
web-console 86
web-clustering 86
web-passivation 86
webservices 87

CHAPTER 7. ENABLE MICROPROFILE APPLICATION DEVELOPMENT FOR JBOSS EAP ON RED HAT CODEREADY STUDIO 88
7.1. CONFIGURING CODEREADY STUDIO TO USE MICROPROFILE CAPABILITIES 88
7.2. USING MICROPROFILE QUICKSTARTS FOR CODEREADY STUDIO 89

CHAPTER 8. THE BOOTABLE JAR 91

8.1. ABOUT THE BOOTABLE JAR 91
8.2. JBOSS EAP MAVEN PLUG-IN 91
8.3. BOOTABLE JAR ARGUMENTS 92
8.4. SPECIFYING GALLEON LAYERS FOR YOUR BOOTABLE JAR SERVER 94
8.5. USING A BOOTABLE JAR ON A JBOSS EAP BARE-METAL PLATFORM 96
8.6. CREATING A HOLLOW BOOTABLE JAR ON A JBOSS EAP BARE-METAL PLATFORM 99
8.7. CLI SCRIPTS EXECUTED AT BUILD TIME 101
8.8. EXECUTING CLI SCRIPT AT RUNTIME 102
8.9. USING A BOOTABLE JAR ON A JBOSS EAP OPENSHIFT PLATFORM 103
8.10. CONFIGURE THE BOOTABLE JAR FOR OPENSHIFT 106
8.11. USING A CONFIGMAP IN YOUR APPLICATION ON OPENSHIFT 107
8.12. CREATING A BOOTABLE JAR MAVEN PROJECT 109
8.13. ENABLING JSON LOGGING FOR YOUR BOOTABLE JAR 111
8.14. ENABLING WEB SESSION DATA STORAGE FOR MULTIPLE BOOTABLE JAR INSTANCES 116
8.15. ENABLING HTTP AUTHENTICATION FOR BOOTABLE JAR WITH A CLI SCRIPT 122
8.16. SECURING YOUR JBOSS EAP BOOTABLE JAR APPLICATION WITH RED HAT SINGLE SIGN-ON 126
8.17. PACKAGING A BOOTABLE JAR IN DEV MODE 132
8.18. UPGRADING SERVER ARTIFACTS 133
8.19. UPDATING EAP 7.4.GA DEPENDENCY 134
8.20. APPLYING THE JBOSS EAP PATCH TO YOUR BOOTABLE JAR 135

CHAPTER 9. OPENID CONNECT IN JBOSS EAP 137
9.1. OPENID CONNECT CONFIGURATION IN JBOSS EAP 137
Deployment configuration 137
Subsystem configuration 138
9.2. ENABLING THE ELYTRON-OIDC-CLIENT SUBSYSTEM 139
9.3. SECURING APPLICATIONS USING OPENID CONNECT WITH RED HAT SINGLE SIGN-ON 140
9.3.1. Configuring Red Hat Single Sign-On as an OpenID provider 140
9.3.2. Configuring a Maven project for creating a secure application 142
9.3.3. Creating a secure application that uses OpenID Connect 144
9.3.4. Restricting access to applications based on user roles 147
9.3.5. Creating and assigning user roles in Red Hat Single Sign-On 148
9.4. DEVELOPING JBOSS EAP BOOTABLE JAR APPLICATION WITH OPENID CONNECT 149
9.4.1. Configuring Red Hat Single Sign-On as an OpenID provider 150
9.4.2. Configuring a Maven project for a bootable jar OIDC application 151
9.4.3. Creating a bootable jar application that uses OpenID Connect 154
9.4.4. Restricting access based on user roles in bootable jar OIDC applications 158
9.4.5. Creating and assigning user roles in Red Hat Single Sign-On 159

CHAPTER 10. OBSERVABILITY IN JBOSS EAP 161
10.1. OPENTELEMETRY IN JBOSS EAP 161
10.2. OPENTELEMETRY CONFIGURATION IN JBOSS EAP 161
10.3. OPENTELEMETRY TRACING IN JBOSS EAP 162
10.4. ENABLING OPENTELEMETRY TRACING IN JBOSS EAP 163
10.5. CONFIGURING THE OPENTELEMETRY SUBSYSTEM 163
10.6. USING JAEGER TO OBSERVE THE OPENTELEMETRY TRACES FOR AN APPLICATION 164
10.7. OPENTELEMETRY TRACING APPLICATION DEVELOPMENT 165
10.7.1. Configuring a Maven project for OpenTelemetry tracing 165
10.7.2. Creating applications that create custom spans 168

CHAPTER 11. REFERENCE 172
11.1. MICROPROFILE CONFIG REFERENCE 172
11.1.1. Default MicroProfile Config attributes 172
11.1.2. MicroProfile Config SmallRye ConfigSources 172

11.2. MICROPROFILE FAULT TOLERANCE REFERENCE 172
11.2.1. MicroProfile Fault Tolerance configuration properties 172
11.3. MICROPROFILE JWT REFERENCE 173
11.3.1. MicroProfile Config JWT standard properties 173
11.4. MICROPROFILE OPENAPI REFERENCE 173
11.4.1. MicroProfile OpenAPI configuration properties 173
11.5. MICROPROFILE REACTIVE MESSAGING REFERENCE 175
11.5.1. MicroProfile reactive messaging connectors for integrating with external messaging systems 175
11.5.2. Example of the data exchange between reactive messaging streams and user-initialized code 176
11.5.3. The Apache Kafka user API 177
Example of how to write and read a message key 178
Example of Kafka mapping in a microprofile-config.properties file 178
11.5.4. Example MicroProfile Config properties file for the Kafka connector 179
Mandatory MicroProfile Reactive Messaging prefixes 180
11.6. OPENID CONNECT REFERENCE 181
11.6.1. elytron-oidc-client subsystem attributes 181
11.7. OPENTELEMETRY REFERENCE 193
11.7.1. OpenTelemetry subsystem attributes 193

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 .

PROVIDING FEEDBACK ON JBOSS EAP DOCUMENTATION

To report an error or to improve our documentation, log in to your Red Hat Jira account and submit an issue. If you do not have a Red Hat Jira account, then you will be prompted to create an account.

Procedure

1. Click the following link to create a ticket.
2. Please include the **Document URL**, the **section number** and **describe the issue**.
3. Enter a brief description of the issue in the **Summary**.
4. Provide a detailed description of the issue or enhancement in the **Description**. Include a URL to where the issue occurs in the documentation.
5. Clicking **Submit** creates and routes the issue to the appropriate documentation team.

CHAPTER 1. JBOSS EAP XP FOR THE LATEST MICROPROFILE CAPABILITIES

1.1. ABOUT JBOSS EAP XP

The MicroProfile Expansion Pack (JBoss EAP XP) is available as a patch stream, which is provided using JBoss EAP XP manager.

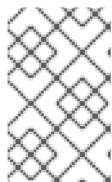

NOTE

JBoss EAP XP is subject to a separate support and life cycle policy. For more details, see the JBoss Enterprise Application Platform expansion pack Support and Life Cycle Policies page.

The JBoss EAP XP patch provides the following MicroProfile 4.1 components:

- MicroProfile Config
- MicroProfile Fault Tolerance
- MicroProfile Health
- MicroProfile JWT
- MicroProfile Metrics
- MicroProfile OpenAPI
- MicroProfile OpenTracing
- MicroProfile REST Client
- MicroProfile Reactive Messaging

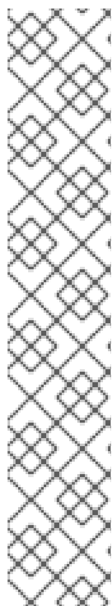

NOTE

The MicroProfile Reactive Messaging subsystem supports Red Hat AMQ Streams. This feature implements the MicroProfile Reactive Messaging 2.0.1 API and Red Hat provides the feature as a technology preview for JBoss EAP XP 4.0.0.

Red Hat tested Red Hat AMQ Streams 2021.Q4 on JBoss EAP. However, check the Red Hat JBoss Enterprise Application Platform supported configurations page for information about the latest Red Hat AMQ Streams version that has been tested on JBoss EAP XP 4.0.0.

Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see Technology Preview Features Support Scope .

1.2. JBOSS EAP XP INSTALLATION

When you install JBoss EAP XP, make sure that the JBoss EAP XP patch is compatible with your version of JBoss EAP. The JBoss EAP XP 4.0.x patch is compatible with the JBoss EAP 7.4 release.

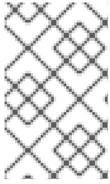

NOTE

You can install JBoss EAP XP either through the XP manager and EAP archive or using the JBoss EAP XP OpenShift Container images. You cannot install JBoss EAP XP on top of EAP RPMs.

Additional Resources

- For more information about installing the latest JBoss EAP XP patch on the latest JBoss EAP release, see Installing JBoss EAP XP 4.0.0 on JBoss EAP 7.4.x .

1.3. JBOSS EAP XP MANAGER

JBoss EAP XP manager is an executable **jar** file that you can download from the **Product Downloads** page. Use JBoss EAP XP manager to apply the JBoss EAP XP patches from the JBoss EAP XP patch stream. The patches contain the MicroProfile 4.1 implementations and the bug fixes for these MicroProfile 4.1 implementations.

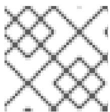

NOTE

You can not manage the JBoss EAP XP patches using the management console.

If you run JBoss EAP XP manager without any arguments, or with the **help** command, you get a list of all the available commands with a description of what they do.

Run the manager with the **help** command to get more information about the arguments available.

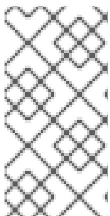

NOTE

Most of the JBoss EAP XP manager commands take a **--jboss-home** argument to point to the JBoss EAP XP server to manage the JBoss EAP XP patch stream. Specify the path to the server in the **JBOSS_HOME** environment variable if you want to omit this. **--jboss-home** takes precedence over the environment variable.

1.4. JBOSS EAP XP MANAGER 4.0 COMMANDS

JBoss EAP XP manager 4.0 provides different commands for managing JBoss EAP XP patch streams.

The following commands are provided:

patch-apply

Use this command to apply patches to your JBoss EAP installation.
The **patch-apply** command is similar to the **patch apply** management CLI command. The **patch-apply** command accepts only those arguments that are required for applying patches using the tool. It uses the default values for other **patch apply** management CLI command arguments.

You can use the **patch-apply** command to apply patches to any patch stream that is enabled on the server. You can also use the command to apply both the base server patches as well as the XP patches.

Example of using the **patch-apply** command:

```
$ java -jar jboss-eap-xp-manager.jar patch-apply --jboss-home=/PATH/TO/EAP --
patch=/PATH/TO/PATCH/jboss-eap-7.3.4-patch.zip
```

When you apply an XP patch, JBoss EAP XP manager 4.0 performs validation to prevent patch and patch stream mismatch. The following example illustrates incorrect combinations:

- Trying to install JBoss EAP XP 3.0 patch on a server with XP 4.0 patch stream set up causes the following error:

```
java.lang.IllegalStateException: The JBoss EAP XP patch stream in the patch 'jboss-eap-
xp-3.0' does not match the currently enabled JBoss EAP XP patch stream [jboss-eap-xp-
4.0]
 at
org.jboss.eap.util.xp.patch.stream.manager.ManagerPatchApplyAction.doExecute(Manager
PatchApplyAction.java:33)
 at
org.jboss.eap.util.xp.patch.stream.manager.ManagerAction.execute(ManagerAction.java:40)

 at org.jboss.eap.util.xp.patch.stream.manager.ManagerMain.main(ManagerMain.java:50)
```

- Trying to install JBoss EAP XP 4.0.0 patch on a server that is not set up for JBoss EAP XP 4.0.0 patch stream causes the following error:

```
java.lang.IllegalStateException: You are attempting to install a patch for the 'jboss-eap-xp-
4.0' JBoss EAP XP Patch Stream. However this patch stream is not yet set up in the
JBoss EAP server. Run the 'setup' command to enable the patch stream.
 at
org.jboss.eap.util.xp.patch.stream.manager.ManagerPatchApplyAction.doExecute(Manager
PatchApplyAction.java:29)
 at
org.jboss.eap.util.xp.patch.stream.manager.ManagerAction.execute(ManagerAction.java:40)

 at org.jboss.eap.util.xp.patch.stream.manager.ManagerMain.main(ManagerMain.java:50)
```

In both the cases, no changes are made to the server.

remove

Use this command to remove the JBoss EAP XP patch stream setup from the JBoss EAP server.

Example of using the **remove** command

```
$ java -jar jboss-eap-xp-manager.jar remove --jboss-home=/PATH/TO/EAP
```

setup

Use this command to set up a clean JBoss EAP server for the JBoss EAP XP patch stream.
When you use the **setup** command, JBoss EAP XP manager performs the following actions:

- Enables the JBoss EAP XP 4.0.0 patch stream.
- Applies patches specified using **--base-patch** and **--xp-patch** attributes.

- Copies the **standalone-microprofile.xml** and **standalone-microprofile-ha.xml** configuration files into the server configuration directory.
  If older configuration files are already installed, the new files are saved as timestamped copies in the target configuration directory, such as **standalone-microprofile-yyyyMMdd-HHmmss.xml**.

  You can set the target directory using the **--jboss-config-directory** argument.

Example of using the setup command

```
$ java -jar jboss-eap-xp-manager.jar setup --jboss-home=/PATH/TO/EAP
```

status

Use this command to find the current status of your JBoss EAP XP server. The status command returns the following information:

- The status of the JBoss EAP XP stream.
- Any support policy changes due to being in the current state.
- The major version of JBoss EAP XP.
- Enabled patch streams and their cumulative patch IDs.
- The available JBoss EAP XP manager commands to change the state.

Example of using the status command

```
$ java -jar jboss-eap-xp-manager.jar status --jboss-home=/PATH/TO/EAP
```

upgrade

Use this command to upgrade an old JBoss EAP XP patch stream to the latest patch stream in the JBoss EAP server.
When you use the **upgrade** command, JBoss EAP XP manager performs the following actions:

- Creates a backup of the files enabling the old patch stream in the server.
- Enables the JBoss EAP XP 4.0 patch stream.
- Applies patches specified using **--base-patch** and **--xp-patch** attributes.
- Copies the **standalone-microprofile.xml** and **standalone-microprofile-ha.xml** configuration files into the server configuration directory. If older configuration files are already installed, the new files are saved as timestamped copies in the target configuration directory, such as **standalone-microprofile-yyyyMMdd-HHmmss.xml**.
- If something goes wrong, JBoss EAP XP manager attempts to restore the previous patch stream from the backup it created.
  You can set the target directory using the **--jboss-config-directory** argument

Example of using the upgrade command:

```
$ java -jar jboss-eap-xp-manager.jar upgrade --jboss-home=/PATH/TO/EAP
```

1.5. INSTALLING JBOSS EAP XP 4.0.0 ON JBOSS EAP 7.4.X

Install JBoss EAP XP 4.0.0 on the JBoss EAP 7.4 base server.

Use JBoss EAP XP manager 4.0.0 to manage JBoss EAP XP 4.0.0 patch streams.

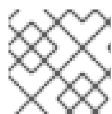

NOTE

JBoss EAP XP 4.0.0 is certified with JBoss EAP 7.4.x.

Prerequisites

- You have downloaded the following files from the **Product Downloads** page:
    - The **jboss-eap-xp-4.0.0-manager.jar** file (JBoss EAP XP manager 4.0)
    - JBoss EAP 7.4 server archive file
    - JBoss EAP XP 4.0.0 patch

Procedure

1. Extract the downloaded JBoss EAP 7.4 server archive file to the path of your JBoss EAP installation.

2. Set up JBoss EAP XP manager 4.0.0 to manage the JBoss EAP XP 4.0 patch stream by using the following command:

```
$ java -jar jboss-eap-xp-manager.jar setup --jboss-home=<path_to_eap>
```

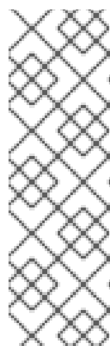

NOTE

You can apply the JBoss EAP XP 4.0.0 patch at the same time. Include the path to the JBoss EAP XP 4.0.0 patch by using the **--xp-patch** argument.

Example:

```
$ java -jar jboss-eap-xp-manager.jar setup --jboss-home=<path_to_eap> --xp-patch=<path_to_patch>jboss-eap-xp-4.0.0-patch.zip
```

The server is now ready to manage the JBoss EAP XP 4.0.0 patch stream.

3. Optional: If you have not applied the JBoss EAP XP 4.0.0 patch to your JBoss EAP server by using the **--xp-patch** argument, apply the JBoss EAP XP 4.0.0 patch by using the JBoss EAP XP manager 4.0.0 **patch-apply** command:

```
$ java -jar jboss-eap-xp-manager.jar patch-apply --jboss-home=<path_to_eap> --patch=<path_to_patch>jboss-eap-xp-4.0.0-patch.zip
```

The **patch-apply** command is similar to the **patch apply** management CLI command. You can also use the **patch apply** management CLI command to apply the patch.

The JBoss EAP server is now ready to manage the JBoss EAP XP 4.0.0 patch stream as you patched the JBoss EAP server with the JBoss EAP XP 4.0.0 patch.

Additional Resources

- JBoss EAP XP manager 4.0 commands

1.6. UNINSTALLING JBOSS EAP XP

Uninstalling JBoss EAP XP removes all the files related to enabling the JBoss EAP XP 4.0.0 patch stream and the MicroProfile 4.1 functionality. The uninstallation process does not affect anything in the base server patch stream or functionality.

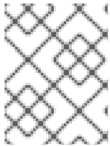

NOTE

The uninstallation process does not remove any configuration files, including the ones you added to the JBoss EAP XP patches when you enabled the JBoss EAP XP patch stream.

Procedure

- Uninstall JBoss EAP XP 4.0.0 by issuing the following command:

```
$ java -jar jboss-eap-xp-manager.jar remove --jboss-home=/PATH/TO/EAP
```

To install MicroProfile 4.1 functionality again, run the **setup** command again to enable the patch stream, and then apply JBoss EAP XP patches to add the MicroProfile 4.1 modules.

1.7. VIEWING THE STATUS OF JBOSS EAP XP

You can view the following information with the **status** command:

- The status of the JBoss EAP XP stream.
- Any support policy changes due to being in the current state.
- The major version of JBoss EAP XP.
- Enabled patch streams and their cumulative patch ids.
- The available JBoss EAP XP manager commands to change the state.

JBoss EAP XP can be in one of the following states:

Not set up

JBoss EAP is clean and does not have JBoss EAP XP set up.

Set up

JBoss EAP has JBoss EAP XP set up. The version of the XP patch stream is not displays as the user can use CLI to determine it.

Inconsistent

The files relating to the JBoss EAP XP are in an inconsistent state. This is an error condition and should not happen normally. If you encounter this error, remove the JBoss EAP XP manager as described in the Uninstalling JBoss EAP XP topic and install JBoss EAP XP again using the **setup** command.

Procedure

- View the status of JBoss EAP XP by issuing the following command:

```
$ java -jar jboss-eap-xp-manager.jar status --jboss-home=<path_to_eap>
```

Additional Resources

- Uninstalling JBoss EAP XP
- Installing JBoss EAP XP 4.0.0 on JBoss EAP 7.4.x

1.8. ROLLING BACK JBOSS EAP XP AND JBOSS EAP 7.4.X BASE PATCHES

You can roll back a previously applied JBoss EAP XP patch or JBoss EAP 7.4.x base patch using the management CLI.

Additional resources

- For more information about rolling back a JBoss EAP XP patch or a JBoss EAP 7.4.x base patch, see Rolling back a patch using the management CLI .

CHAPTER 2. UNDERSTAND MICROPROFILE

2.1. MICROPROFILE CONFIG

2.1.1. MicroProfile Config in JBoss EAP

Configuration data can change dynamically and applications need to be able to access the latest configuration information without restarting the server.

MicroProfile Config provides portable externalization of configuration data. This means, you can configure applications and microservices to run in multiple environments without modification or repackaging.

MicroProfile Config functionality is implemented in JBoss EAP using the SmallRye Config component and is provided by the **microprofile-config-smallrye** subsystem.

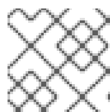

NOTE

MicroProfile Config is only supported in JBoss EAP XP. It is not supported in JBoss EAP.

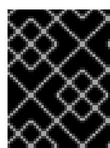

IMPORTANT

If you are adding your own Config implementations, you need to use the methods in the latest version of the Config interface.

Additional Resources

- MicroProfile Config
- SmallRye Config
- Config implementations

2.1.2. MicroProfile Config sources supported in MicroProfile Config

MicroProfile Config configuration properties can come from different locations and can be in different formats. These properties are provided by ConfigSources. ConfigSources are implementations of the **org.eclipse.microprofile.config.spi.ConfigSource** interface.

The MicroProfile Config specification provides the following default **ConfigSource** implementations for retrieving configuration values:

- **System.getProperties()**.
- **System.getenv()**.
- All **META-INF/microprofile-config.properties** files on the class path.

The **microprofile-config-smallrye** subsystem supports additional types of **ConfigSource** resources for retrieving configuration values. You can also retrieve the configuration values from the following resources:

- Properties in a **microprofile-config-smallrye/config-source** management resource

- Files in a directory
- **ConfigSource** class
- **ConfigSourceProvider** class

Additional Resources

- **org.jboss.resteasy.microprofile.config.BaseServletConfigSource**

2.2. MICROPROFILE FAULT TOLERANCE

2.2.1. About MicroProfile Fault Tolerance specification

The MicroProfile Fault Tolerance specification defines strategies to deal with errors inherent in distributed microservices.

The MicroProfile Fault Tolerance specification defines the following strategies to handle errors:

Timeout

Define the amount of time within which an execution must finish. Defining a timeout prevents waiting for an execution indefinitely.

Retry

Define the criteria for retrying a failed execution.

Fallback

Provide an alternative in the case of a failed execution.

CircuitBreaker

Define the number of failed execution attempts before temporarily stopping. You can define the length of the delay before resuming execution.

Bulkhead

Isolate failures in part of the system so that the rest of the system can still function.

Asynchronous

Execute client request in a separate thread.

Additional Resources

- MicroProfile Fault Tolerance specification

2.2.2. MicroProfile Fault Tolerance in JBoss EAP

The **microprofile-fault-tolerance-smallrye** subsystem provides support for MicroProfile Fault Tolerance in JBoss EAP. The subsystem is available only in the JBoss EAP XP stream.

The **microprofile-fault-tolerance-smallrye** subsystem provides the following annotations for interceptor bindings:

- **@Timeout**
- **@Retry**
- **@Fallback**

- **@CircuitBreaker**
- **@Bulkhead**
- **@Asynchronous**

You can bind these annotations at the class level or at the method level. An annotation bound to a class applies to all of the business methods of that class.

The following rules apply to binding interceptors:

- If a component class declares or inherits a class-level interceptor binding, the following restrictions apply:
  - The class must not be declared final.
  - The class must not contain any static, private, or final methods.
- If a non-static, non-private method of a component class declares a method level interceptor binding, neither the method nor the component class may be declared final.

Fault tolerance operations have the following restrictions:

- Fault tolerance interceptor bindings must be applied to a bean class or bean class method.
- When invoked, the invocation must be the business method invocation as defined in the Jakarta Contexts and Dependency Injection specification.
- An operation is not considered fault tolerant if both of the following conditions are true:
  - The method itself is not bound to any fault tolerance interceptor.
  - The class containing the method is not bound to any fault tolerance interceptor.

The **microprofile-fault-tolerance-smallrye** subsystem provides the following configuration options, in addition to the configuration options provided by MicroProfile Fault Tolerance:

- **io.smallrye.faulttolerance.mainThreadPoolSize**
- **io.smallrye.faulttolerance.mainThreadPoolQueueSize**

Additional Resources

- MicroProfile Fault Tolerance Specification
- SmallRye Fault Tolerance project

2.3. MICROPROFILE HEALTH

2.3.1. MicroProfile Health in JBoss EAP

JBoss EAP includes the SmallRye Health component, which you can use to determine whether the JBoss EAP instance is responding as expected. This capability is enabled by default.

MicroProfile Health is only available when running JBoss EAP as a standalone server.

The MicroProfile Health specification defines the following health checks:

Readiness

Determines whether an application is ready to process requests. The annotation **@Readiness** provides this health check.

Liveness

Determines whether an application is running. The annotation **@Liveness** provides this health check.

Startup

Determines whether an application has already started. The annotation **@Startup** provides this health check.

The **@Health** annotation was removed in MicroProfile Health 3.0.

MicroProfile Health 3.1 includes a new **Startup** health check probe.

For more information about the changes in MicroProfile Health 3.1, see Release Notes for MicroProfile Health 3.1.

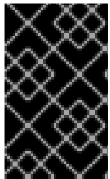

IMPORTANT

The **:empty-readiness-checks-status**, **:empty-liveness-checks-status**, and **:empty-startup-checks-status** management attributes specify the global status when no **readiness**, **liveness**, or **startup** probes are defined.

Additional Resources

- Global status when probes are not defined
- SmallRye Health
- MicroProfile Health
- Custom health check example

2.4. MICROPROFILE JWT

2.4.1. MicroProfile JWT integration in JBoss EAP

The subsystem **microprofile-jwt-smallrye** provides MicroProfile JWT integration in JBoss EAP.

The following functionalities are provided by the **microprofile-jwt-smallrye** subsystem:

- Detecting deployments that use MicroProfile JWT security.
- Activating support for MicroProfile JWT.

The subsystem contains no configurable attributes or resources.

In addition to the **microprofile-jwt-smallrye** subsystem, the **org.eclipse.microprofile.jwt.auth.api** module provides MicroProfile JWT integration in JBoss EAP.

Additional Resources

- SmallRye JWT

2.4.2. Differences between a traditional deployment and an MicroProfile JWT deployment

MicroProfile JWT deployments do not depend on managed SecurityDomain resources like traditional JBoss EAP deployments. Instead, a virtual SecurityDomain is created and used across the MicroProfile JWT deployment.

As the MicroProfile JWT deployment is configured entirely within the MicroProfile Config properties and the **microprofile-jwt-smallrye** subsystem, the virtual SecurityDomain does not need any other managed configuration for the deployment.

2.4.3. MicroProfile JWT activation in JBoss EAP

MicroProfile JWT is activated for applications based on the presence of an **auth-method** in the application.

The MicroProfile JWT integration is activated for an application in the following way:

- As part of the deployment process, JBoss EAP scans the application archive for the presence of an **auth-method**.
- If an **auth-method** is present and defined as **MP-JWT**, the MicroProfile JWT integration is activated.

The **auth-method** can be specified in either or both of the following files:

- the file containing the class that extends **javax.ws.rs.core.Application**, annotated with the **@LoginConfig**
- the **web.xml** configuration file

If **auth-method** is defined both in a class, using annotation, and in the web.xml configuration file, the definition in **web.xml** configuration file is used.

2.4.4. Limitations of MicroProfile JWT in JBoss EAP

The MicroProfile JWT implementation in JBoss EAP has certain limitations.

The following limitations of MicroProfile JWT implementation exist in JBoss EAP:

- The MicroProfile JWT implementation parses only the first key from the JSON Web Key Set (JWKS) supplied in the **mp.jwt.verify.publickey** property. Therefore, if a token claims to be signed by the second key or any key after the second key, the token fails verification and the request containing the token is not authorized.
- Base64 encoding of JWKS is not supported.

In both cases, a clear text JWKS can be referenced instead of using the **mp.jwt.verify.publickey.location** config property.

2.5. MICROPROFILE METRICS

2.5.1. MicroProfile Metrics in JBoss EAP

JBoss EAP includes the SmallRye Metrics component. The SmallRye Metrics component provides the MicroProfile Metrics functionality using the **microprofile-metrics-smallrye** subsystem.

The **microprofile-metrics-smallrye** subsystem provides monitoring data for the JBoss EAP instance. The subsystem is enabled by default.

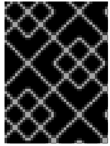

IMPORTANT

The **microprofile-metrics-smallrye** subsystem is only enabled in standalone configurations.

Additional Resources

- SmallRye Metrics
- MicroProfile Metrics

2.6. MICROPROFILE OPENAPI

2.6.1. MicroProfile OpenAPI in JBoss EAP

MicroProfile OpenAPI is integrated in JBoss EAP using the **microprofile-openapi-smallrye** subsystem.

The MicroProfile OpenAPI specification defines an HTTP endpoint that serves an OpenAPI 3.0 document. The OpenAPI 3.0 document describes the REST services for the host. The OpenAPI endpoint is registered using the configured path, for example http://localhost:8080/openapi, local to the root of the host associated with a deployment.

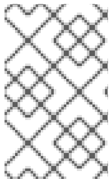

NOTE

Currently, the OpenAPI endpoint for a virtual host can only document a single deployment. To use OpenAPI with multiple deployments registered with different context paths on the same virtual host, each deployment must use a distinct endpoint path.

The OpenAPI endpoint returns a YAML document by default. You can also request a JSON document using an Accept HTTP header, or a format query parameter.

If the Undertow server or host of a given application defines an HTTPS listener then the OpenAPI document is also available using HTTPS. For example, an endpoint for HTTPS is https://localhost:8443/openapi.

2.7. MICROPROFILE OPENTRACING

2.7.1. MicroProfile OpenTracing

The ability to trace requests across service boundaries is important, especially in a microservices environment where a request can flow through multiple services during its life cycle.

The MicroProfile OpenTracing specification defines behaviors and an API for accessing an OpenTracing compliant **Tracer** interface within a CDI-bean application. The **Tracer** interface automatically traces JAX-RS applications.

The behaviors specify how OpenTracing Spans are created automatically for incoming and outgoing requests. The API defines how to explicitly disable or enable tracing for given endpoints.

Additional Resources

- For more information about MicroProfile OpenTracing specification, see MicroProfile OpenTracing documentation.
- For more information about the **Tracer** interface, see **Tracer** javadoc.

2.7.2. MicroProfile OpenTracing in JBoss EAP

You can use the **microprofile-opentracing-smallrye** subsystem to configure the distributed tracing of Jakarta EE applications. This subsystem uses the SmallRye OpenTracing component to provide the MicroProfile OpenTracing functionality for JBoss EAP.

MicroProfile OpenTracing 2.0 supports tracing requests for applications. You can configure the default Jaeger Java Client tracer, plus a set of instrumentation libraries for components commonly used in Jakarta EE, using JBoss EAP management API with the management CLI or the management console.

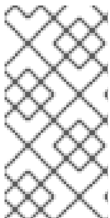

NOTE

Each individual WAR deployed to the JBoss EAP server automatically has its own **Tracer** instance. Each WAR within an EAR is treated as an individual WAR, and each has its own **Tracer** instance. By default, the service name used with the Jaeger Client is derived from the deployment's name, which is usually the WAR file name.

Within the **microprofile-opentracing-smallrye** subsystem, you can configure the Jaeger Java Client by setting system properties or environment variables.

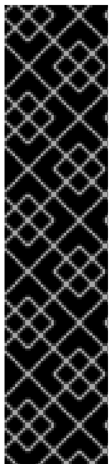

IMPORTANT

Configuring the Jaeger Client tracer using system properties and environment variables is provided as a Technology Preview. The system properties and environment variables affiliated with the Jaeger Client tracer might change and become incompatible with each other in future releases.

Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see Technology Preview Features Support Scope.

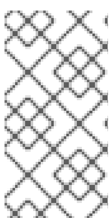

NOTE

By default, the probabilistic sampling strategy of the Jaeger Client for Java is set to **0.001**, meaning that only approximately one in one thousand traces are sampled. To sample every request, set the system properties **JAEGER_SAMPLER_TYPE** to **const** and **JAEGER_SAMPLER_PARAM** to **1**.

Additional Resources

- For more information about SmallRye OpenTracing functionality, see the SmallRye OpenTracing component.
- For more information about the default tracer, see the Jaeger Java Client.
- For more information about the **Tracer** interface, see **Tracer** javadoc.
- For more information about overriding the default tracer and tracing Jakarta Contexts and Dependency Injection beans, see Using Eclipse MicroProfile OpenTracing to Trace Requests in the *Development Guide*.
- For more information about configuring the Jaeger Client, see the Jaeger documentation.
- For more information about valid system properties, see Configuration via Environment in the Jaeger documentation.

2.8. MICROPROFILE REST CLIENT

2.8.1. MicroProfile REST client

JBoss EAP XP 4.0.0 supports the MicroProfile REST client 2.0 that builds on Jakarta RESTful Web Services 2.1.6 client APIs to provide a type-safe approach to invoke RESTful services over HTTP. The MicroProfile Type Safe REST clients are defined as Java interfaces. With the MicroProfile REST clients, you can write client applications with executable code.

Use the MicroProfile REST client to avail the following capabilities:

- An intuitive syntax
- Programmatic registration of providers
- Declarative registration of providers
- Declarative specification of headers
- **ResponseExceptionMapper**
- Jakarta Contexts and Dependency Injection integration
- Access to server-sent events (SSE)

Additional resources

- A comparison between MicroProfile REST client and Jakarta RESTful Web Services syntaxes
- Programmatic registration of providers in MicroProfile REST client
- Declarative registration of providers in MicroProfile REST client
- Declarative specification of headers in MicroProfile REST client
- ResponseExceptionMapper in MicroProfile REST client
- Context dependency injection with MicroProfile REST client

2.8.2. The **resteasy.original.webapplicationexception.behavior** MicroProfile Config

property

MicroProfile Config is the name of a specification that developers can use to configure applications and microservices to run in multiple environments without having to modify or repackage those apps. Previously, MicroProfile Config was available for JBoss EAP as a technology preview, but it has since been removed. MicroProfile Config is now available only on JBoss EAP XP.

Defining the resteasy.original.webapplicationexception.behavior MicroProfile Config property
You can set the **resteasy.original.webapplicationexception.behavior** parameter as either a **web.xml** servlet property or a system property. Here's an example of one such servlet property in **web.xml**:

```
<context-param>
    <param-name>resteasy.original.webapplicationexception.behavior</param-name>
    <param-value>true</param-value>
</context-param>
```

You can also use MicroProfile Config to configure any other RESTEasy property.

Additional resources

- For more information about MicroProfile Config on JBoss EAP XP, see *Understand MicroProfile*.
- For more information about the MicroProfile REST Client, see MicroProfile REST Client.
- For more information about RESTEasy, see Jakarta RESTful Web Services Request Processing.

2.9. MICROPROFILE REACTIVE MESSAGING

2.9.1. MicroProfile reactive messaging

When you upgrade to JBoss EAP XP 4.0.0, you can enable the newest version of MicroProfile Reactive Messaging, which includes reactive messaging extensions and subsystems.

A "reactive stream" is a succession of event data, along with processing protocols and standards, that is pushed across an asynchronous boundary (like a scheduler) without any buffering. An "event" might be a scheduled and repeating temperature check in a weather app, for example. The primary benefit of reactive streams is the seamless interoperability of your various applications and implementations.

Reactive messaging provides a framework for building event-driven, data-streaming, and event-sourcing applications. Reactive messaging results in the constant and smooth exchange of event data, the reactive stream, from one app to another. You can use MicroProfile Reactive Messaging for asynchronous messaging through reactive streams so that your application can interact with others, like Apache Kafka, for example.

After you upgrade your instance of MicroProfile Reactive Messaging to the latest version, you can do the following:

- Provision a server with MicroProfile Reactive Messaging for the Apache Kafka data-streaming platform.
- Interact with reactive messaging in-memory and backed by Apache Kafka topics through the latest reactive messaging APIs.
- Use MicroProfile Metrics to find out how many messages are streamed on a given channel.

Additional resources

- For more information about Apache Kafka, see What is Apache Kafka?

2.9.2. MicroProfile reactive messaging connectors

You can use connectors to integrate MicroProfile Reactive Messaging with a number of external messaging systems. MicroProfile for JBoss EAP comes with the Apache Kafka connector. Use the Eclipse MicroProfile Config specification to configure your connectors.

The Apache Kafka connector and incorporated layers
MicroProfile Reactive Messaging includes the Kafka connector, which you can configure with MicroProfile Config. The Kafka connector incorporates **microprofile-reactive-messaging-kafka** and **microprofile-reactive-messaging** Galleon layers. The **microprofile-reactive-messaging** layer provides the core MicroProfile Reactive Messaging functionality.

Table 2.1. Reactive messaging and Apache Kafka connector Galleon layers

<table>
<tr><th>Layer</th><th>Definition</th></tr>
<tr><td>microprofile-reactive-streams-operators</td><td>Provides MicroProfile Reactive Streams Operators APIs and supporting implementing modules.
Contains MicroProfile Reactive Streams Operators with SmallRye extension and subsystem.
Depends on cdi layer.
cdi stands for Jakarta Contexts and Dependency Injection; provides subsystems that add @Inject functionality.</td></tr>
<tr><td>microprofile-reactive-messaging</td><td>Provides MicroProfile Reactive Messaging APIs and supporting implementing modules.
Contains MicroProfile with SmallRye extension and subsystem.
Depends on microprofile-config and microprofile-reactive-streams-operators layers.</td></tr>
<tr><td>microprofile-reactive-messaging-kafka</td><td>Provides Kafka connector modules that enable MicroProfile Reactive Messaging to interact with Kafka.
Depends on microprofile-reactive-messaging layer.</td></tr>
</table>

2.9.3. The Apache Kafka event streaming platform

Apache Kafka is an open source distributed event (data) streaming platform that can publish, subscribe to, store, and process streams of records in real time. It handles event streams from multiple sources and delivers them to multiple consumers, moving large amounts of data from points A to Z and everywhere else, all at the same time. MicroProfile Reactive Messaging uses Apache Kafka to deliver these event records in as few as two microseconds, store them safely in distributed, fault-tolerant clusters, all while making them available across any team-defined zones or geographic regions.

Additional resources

- What is Apache Kafka?
- Red Hat OpenShift Streams for Apache Kafka
- Red Hat AMQ

CHAPTER 3. ADMINISTER MICROPROFILE IN JBOSS EAP

3.1. MICROPROFILE OPENTRACING ADMINISTRATION

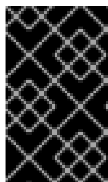

IMPORTANT

If you see duplicate traces exported for REST calls, disable the **microprofile-opentracing-smallrye** subsystem. For information about disabling the **microprofile-opentracing-smallrye**, see Removing the microprofile-opentracing-smallrye subsystem.

3.1.1. Enabling MicroProfile Open Tracing

Use the following management CLI commands to enable the MicroProfile Open Tracing feature globally for the server instance by adding the subsystem to the server configuration.

Procedure

1. Enable the **microprofile-opentracing-smallrye** subsystem using the following management command:

```
/subsystem=microprofile-opentracing-smallrye:add()
```

2. Reload the server for the changes to take effect.

```
reload
```

3.1.2. Removing the microprofile-opentracing-smallrye subsystem

The **microprofile-opentracing-smallrye** subsystem is included in the default JBoss EAP 7.4 configuration. This subsystem provides MicroProfile OpenTracing functionality for JBoss EAP 7.4. If you experience system memory or performance degradation with MicroProfile OpenTracing enabled, you might want to disable the **microprofile-opentracing-smallrye** subsystem.

You can use the **remove** operation in the management CLI to disable the MicroProfile OpenTracing feature globally for a given server.

Procedure

1. Remove the **microprofile-opentracing-smallrye** subsystem.

```
/subsystem=microprofile-opentracing-smallrye:remove()
```

2. Reload the server for the changes to take effect.

```
reload
```

3.1.3. Installing Jaeger

Install Jaeger using **docker**.

Prerequisites

- **docker** is installed.

Procedure

1. Install Jaeger using **docker** by issuing the following command in CLI:

```
$ docker run -d --name jaeger   -p 6831:6831/udp   -p 5778:5778   -p 14268:14268   -p
16686:16686   jaegertracing/all-in-one:1.16
```

3.2. MICROPROFILE CONFIG CONFIGURATION

3.2.1. Adding properties in a ConfigSource management resource

You can store properties directly in a **config-source** subsystem as a management resource.

Procedure

- Create a ConfigSource and add a property:

```
/subsystem=microprofile-config-smallrye/config-source=props:add(properties={"name" =
"jim"})
```

3.2.2. Configuring directories as ConfigSources

When a property is stored in a directory as a file, the file-name is the name of a property and the file content is the value of the property.

Procedure

1. Create a directory where you want to store the files:

```
$ mkdir -p ~/config/prop-files/
```

2. Navigate to the directory:

```
$ cd ~/config/prop-files/
```

3. Create a file **name** to store the value for the property **name**:

```
$ touch name
```

4. Add the value of the property to the file:

```
$ echo "jim" > name
```

5. Create a ConfigSource in which the file name is the property and the file contents the value of the property:

```
/subsystem=microprofile-config-smallrye/config-source=file-props:add(dir=
{path=~/config/prop-files})
```

This results in the following XML configuration:

```
<subsystem xmlns="urn:wildfly:microprofile-config-smallrye:1.0">
    <config-source name="file-props">
        <dir path="/etc/config/prop-files"/>
    </config-source>
</subsystem>
```

3.2.3. Obtaining ConfigSource from a ConfigSource class

You can create and configure a custom **org.eclipse.microprofile.config.spi.ConfigSource** implementation class to provide a source for the configuration values.

Procedure

- The following management CLI command creates a **ConfigSource** for the implementation class named **org.example.MyConfigSource** that is provided by a JBoss module named **org.example**.
  If you want to use a **ConfigSource** from the **org.example** module, add the **<module name="org.eclipse.microprofile.config.api"/>** dependency to the **path/to/org/example/main/module.xml** file.

```
/subsystem=microprofile-config-smallrye/config-source=my-config-source:add(class=
{name=org.example.MyConfigSource, module=org.example})
```

This command results in the following XML configuration for the **microprofile-config-smallrye** subsystem.

```
<subsystem xmlns="urn:wildfly:microprofile-config-smallrye:1.0">
    <config-source name="my-config-source">
        <class name="org.example.MyConfigSource" module="org.example"/>
    </config-source>
</subsystem>
```

Properties provided by the custom **org.eclipse.microprofile.config.spi.ConfigSource** implementation class are available to any JBoss EAP deployment.

3.2.4. Obtaining ConfigSource configuration from a ConfigSourceProvider class

You can create and configure a custom **org.eclipse.microprofile.config.spi.ConfigSourceProvider** implementation class that registers implementations for multiple **ConfigSource** instances.

Procedure

- Create a **config-source-provider**:

```
/subsystem=microprofile-config-smallrye/config-source-provider=my-config-source-
provider:add(class={name=org.example.MyConfigSourceProvider, module=org.example})
```

The command creates a **config-source-provider** for the implementation class named **org.example.MyConfigSourceProvider** that is provided by a JBoss Module named **org.example**.

If you want to use a **config-source-provider** from the **org.example** module, add the **<module name="org.eclipse.microprofile.config.api"/>** dependency to the **path/to/org/example/main/module.xml** file.

This command results in the following XML configuration for the **microprofile-config-smallrye** subsystem:

```
<subsystem xmlns="urn:wildfly:microprofile-config-smallrye:1.0">
   <config-source-provider name="my-config-source-provider">
      <class name="org.example.MyConfigSourceProvider" module="org.example"/>
   </config-source-provider>
</subsystem>
```

Properties provided by the **ConfigSourceProvider** implementation are available to any JBoss EAP deployment.

Additional resources

- For information about how to add a global module to the JBoss EAP server, see Define Global Modules in the *Configuration Guide* for JBoss EAP.

3.3. MICROPROFILE FAULT TOLERANCE CONFIGURATION

3.3.1. Adding the MicroProfile Fault Tolerance extension

The MicroProfile Fault Tolerance extension is included in **standalone-microprofile.xml** and **standalone-microprofile-ha.xml** configurations that are provided as part of JBoss EAP XP.

The extension is not included in the standard **standalone.xml** configuration. To use the extension, you must manually enable it.

Prerequisites

- EAP XP pack is installed.

Procedure

1. Add the MicroProfile Fault Tolerance extension using the following management CLI command:

```
/extension=org.wildfly.extension.microprofile.fault-tolerance-smallrye:add
```

2. Enable the **microprofile-fault-tolerance-smallrye** subsystem using the following managenent command:

```
/subsystem=microprofile-fault-tolerance-smallrye:add
```

3. Reload the server with the following management command:

```
reload
```

3.4. MICROPROFILE HEALTH CONFIGURATION

3.4.1. Examining health using the management CLI

You can check system health using the management CLI.

Procedure

- Examine health:

```
/subsystem=microprofile-health-smallrye:check
{
    "outcome" => "success",
    "result" => {
        "status" => "UP",
        "checks" => []
    }
}
```

3.4.2. Examining health using the management console

You can check system health using the management console.

A check runtime operation shows the health checks and the global outcome as boolean value.

Procedure

1. Navigate to the **Runtime** tab and select the server.
2. In the **Monitor** column, click **MicroProfile Health** → **View**.

3.4.3. Examining health using the HTTP endpoint

Health check is automatically deployed to the health context on JBoss EAP, so you can obtain the current health using the HTTP endpoint.

The default address for the **/health** endpoint, accessible from the management interface, is **http://127.0.0.1:9990/health**.

Procedure

- To obtain the current health of the server using the HTTP endpoint, use the following URL:.

```
http://<host>:<port>/health
```

  Accessing this context displays the health check in JSON format, indicating if the server is healthy.

3.4.4. Enabling authentication for MicroProfile Health

You can configure the **health** context to require authentication for access.

Procedure

1. Set the **security-enabled** attribute to **true** on the **microprofile-health-smallrye** subsystem.

```
/subsystem=microprofile-health-smallrye:write-attribute(name=security-enabled,value=true)
```

2. Reload the server for the changes to take effect.

```
reload
```

Any subsequent attempt to access the **/health** endpoint triggers an authentication prompt.

3.4.5. Readiness probes that determine server health and readiness

JBoss EAP XP 4.0.0 supports three readiness probes to determine server health and readiness.

- **server-status** - returns **UP** when the server-state is **running**.
- **boot-errors** - returns **UP** when the probe detects no boot errors.
- **deployment-status** - returns **UP** when the status for all deployments is **OK**.

These readiness probes are enabled by default. You can disable the probes using the MicroProfile Config property **mp.health.disable-default-procedures**.

The following example illustrates the use of the three probes with the **check** operation:

```
[standalone@localhost:9990 /] /subsystem=microprofile-health-smallrye:check
{
    "outcome" => "success",
    "result" => {
        "status" => "UP",
        "checks" => [
            {
                "name" => "boot-errors",
                "status" => "UP"
            },
            {
                "name" => "server-state",
                "status" => "UP",
                "data" => {"value" => "running"}
            },
            {
                "name" => "empty-readiness-checks",
                "status" => "UP"
            },
            {
                "name" => "deployments-status",
                "status" => "UP"
            },
            {
                "name" => "empty-liveness-checks",
                "status" => "UP"
            },
            {
                "name" => "empty-startup-checks",
                "status" => "UP"
            }
```

```
    ]
  }
}
```

Additional resources

- MicroProfile Health in JBoss EAP
- Global status when probes are not defined

3.4.6. Global status when probes are not defined

The **:empty-readiness-checks-status**, **:empty-liveness-checks-status**, and **:empty-startup-checks-status** management attributes specify the global status when no **readiness**, **liveness**, or **startup** probes are defined.

These attributes allow applications to report 'DOWN' until their probes verify that the application is ready, live, or started up. By default, applications report 'UP'.

- The **:empty-readiness-checks-status** attribute specifies the global status for **readiness** probes if no **readiness** probes have been defined:

```
/subsystem=microprofile-health-smallrye:read-attribute(name=empty-readiness-checks-status)
{
    "outcome" => "success",
    "result" => expression "${env.MP_HEALTH_EMPTY_READINESS_CHECKS_STATUS:UP}"
}
```

- The **:empty-liveness-checks-status** attribute specifies the global status for **liveness** probes if no **liveness** probes have been defined:

```
/subsystem=microprofile-health-smallrye:read-attribute(name=empty-liveness-checks-status)
{
    "outcome" => "success",
    "result" => expression "${env.MP_HEALTH_EMPTY_LIVENESS_CHECKS_STATUS:UP}"
}
```

- The **:empty-startup-checks-status** attribute specifies the global status for **startup** probes if no **startup** probes have been defined:

```
/subsystem=microprofile-health-smallrye:read-attribute(name=empty-startup-checks-status)
{
    "outcome" => "success",
    "result" => expression "${env.MP_HEALTH_EMPTY_STARTUP_CHECKS_STATUS:UP}"
}
```

  The **/health** HTTP endpoint and the **:check** operation that check **readiness**, **liveness**, and **startup** probes also take into account these attributes.

You can also modify these attributes as shown in the following example:

```
/subsystem=microprofile-health-smallrye:write-attribute(name=empty-readiness-checks-
```

```
status,value=DOWN)
{
    "outcome" => "success",
    "response-headers" => {
        "operation-requires-reload" => true,
        "process-state" => "reload-required"
    }
}
```

3.5. MICROPROFILE JWT CONFIGURATION

3.5.1. Enabling microprofile-jwt-smallrye subsystem

The MicroProfile JWT integration is provided by the **microprofile-jwt-smallrye** subsystem and is included in the default configuration. If the subsystem is not present in the default configuration, you can add it as follows.

Prerequisites

- EAP XP is installed.

Procedure

1. Enable the MicroProfile JWT smallrye extension in JBoss EAP:

```
/extension=org.wildfly.extension.microprofile.jwt-smallrye:add
```

2. Enable the **microprofile-jwt-smallrye** subsystem:

```
/subsystem=microprofile-jwt-smallrye:add
```

3. Reload the server:

```
reload
```

The **microprofile-jwt-smallrye** subsystem is enabled.

3.6. MICROPROFILE METRICS ADMINISTRATION

3.6.1. Metrics available on the management interface

The JBoss EAP subsystem metrics are exposed in Prometheus format.

Metrics are automatically available on the JBoss EAP management interface, with the following contexts:

- **/metrics/** - Contains metrics specified in the MicroProfile 3.0 specification.
- **/metrics/vendor** - Contains vendor-specific metrics, such as memory pools.
- **/metrics/application** - Contains metrics from deployed applications and subsystems that use the MicroProfile Metrics API.

The metric names are based on subsystem and attribute names. For example, the subsystem **undertow** exposes a metric attribute **request-count** for every servlet in an application deployment. The name of this metric is **jboss_undertow_request_count**. The prefix **jboss** identifies JBoss EAP as the source of the metrics.

3.6.2. Examining metrics using the HTTP endpoint

Examine the metrics that are available on the JBoss EAP management interface using the HTTP endpoint.

Procedure

- Use the curl command:

```
$ curl -v http://localhost:9990/metrics | grep -i type
```

3.6.3. Enabling Authentication for the MicroProfile Metrics HTTP Endpoint

Configure the **metrics** context to require users to be authorized to access the context. This configuration extends to all the subcontexts of the **metrics** context.

Procedure

1. Set the **security-enabled** attribute to **true** on the **microprofile-metrics-smallrye** subsystem.

```
/subsystem=microprofile-metrics-smallrye:write-attribute(name=security-enabled,value=true)
```

2. Reload the server for the changes to take effect.

```
reload
```

Any subsequent attempt to access the **metrics** endpoint results in an authentication prompt.

3.6.4. Obtaining the request count for a web service

Obtain the request count for a web service that exposes its request count metric.

The following procedure uses **helloworld-rs** quickstart as the web service for obtaining request count. The quickstart is available at Download the quickstart from: jboss-eap-quickstarts.

Prerequsites

- The web service exposes request count.

Procedure

1. Enable statistics for the **undertow** subsystem:
    - Start the standalone server with statistics enabled:

    ```
    $ ./standalone.sh -Dwildfly.statistics-enabled=true
    ```

    - For an already running server, enable the statistics for the **undertow** subsystem:

```
/subsystem=undertow:write-attribute(name=statistics-enabled,value=true)
```

2. Deploy the **helloworld-rs** quickstart:
   - In the root directory of the quickstart, deploy the web application using Maven:

     ```
     $ mvn clean install wildfly:deploy
     ```

3. Query the HTTP endpoint in the CLI using the **curl** command and filter for **request_count**:

   ```
   $ curl -v http://localhost:9990/metrics |  grep request_count
   ```

   Expected output:

   ```
   jboss_undertow_request_count_total{server="default-server",http_listener="default",} 0.0
   ```

   The attribute value returned is **0.0**.

4. Access the quickstart, located at http://localhost:8080/helloworld-rs/, in a web browser and click any of the links.

5. Query the HTTP endpoint from the CLI again:

   ```
   $ curl -v http://localhost:9990/metrics |  grep request_count
   ```

   Expected output:

   ```
   jboss_undertow_request_count_total{server="default-server",http_listener="default",} 1.0
   ```

   The value is updated to **1.0**.

   Repeat the last two steps to verify that the request count is updated.

3.7. MICROPROFILE OPENAPI ADMINISTRATION

3.7.1. Enabling MicroProfile OpenAPI

The **microprofile-openapi-smallrye** subsystem is provided in the **standalone-microprofile.xml** configuration. However, JBoss EAP XP uses the **standalone.xml** by default. You must include the subsystem in **standalone.xml** to use it.

Alternatively, you can follow the procedure Updating standalone configurations with MicroProfile subsystems and extensions to update the **standalone.xml** configuration file.

Procedure

1. Enable the MicroProfile OpenAPI smallrye extension in JBoss EAP:

   ```
   /extension=org.wildfly.extension.microprofile.openapi-smallrye:add()
   ```

2. Enable the **microprofile-openapi-smallrye** subsystem using the following management command:

```
/subsystem=microprofile-openapi-smallrye:add()
```

3. Reload the server.

```
reload
```

The **microprofile-openapi-smallrye** subsystem is enabled.

3.7.2. Requesting an MicroProfile OpenAPI document using Accept HTTP header

Request an MicroProfile OpenAPI document, in the JSON format, from a deployment using an Accept HTTP header.

By default, the OpenAPI endpoint returns a YAML document.

Prerequisites

- The deployment being queried is configured to return an MicroProfile OpenAPI document.

Procedure

- Issue the following **curl** command to query the **/openapi** endpoint of the deployment:

```
$ curl -v -H'Accept: application/json' http://localhost:8080/openapi
< HTTP/1.1 200 OK
...
{"openapi": "3.0.1" ... }
```

  Replace http://localhost:8080 with the URL and port of the deployment.

  The Accept header indicates that the JSON document is to be returned using the **application/json** string.

3.7.3. Requesting an MicroProfile OpenAPI document using an HTTP parameter

Request an MicroProfile OpenAPI document, in the JSON format, from a deployment using a query parameter in an HTTP request.

By default, the OpenAPI endpoint returns a YAML document.

Prerequisites

- The deployment being queried is configured to return an MicroProfile OpenAPI document.

Procedure

- Issue the following **curl** command to query the **/openapi** endpoint of the deployment:

```
$ curl -v http://localhost:8080/openapi?format=JSON
< HTTP/1.1 200 OK
...
```

  Replace http://localhost:8080 with the URL and port of the deployment.

The HTTP parameter **format=JSON** indicates that JSON document is to be returned.

3.7.4. Configuring JBoss EAP to serve a static OpenAPI document

Configure JBoss EAP to serve a static OpenAPI document that describes the REST services for the host.

When JBoss EAP is configured to serve a static OpenAPI document, the static OpenAPI document is processed before any Jakarta RESTful Web Services and MicroProfile OpenAPI annotations.

In a production environment, disable annotation processing when serving a static document. Disabling annotation processing ensures that an immutable and versioned API contract is available for clients.

Procedure

1. Create a directory in the application source tree:

   ```
   $ mkdir APPLICATION_ROOT/src/main/webapp/META-INF
   ```

   APPLICATION_ROOT is the directory containing the **pom.xml** configuration file for the application.

2. Query the OpenAPI endpoint, redirecting the output to a file:

   ```
   $ curl http://localhost:8080/openapi?format=JSON > src/main/webapp/META-INF/openapi.json
   ```

   By default, the endpoint serves a YAML document, **format=JSON** specifies that a JSON document is returned.

3. Configure the application to skip annotation scanning when processing the OpenAPI document model:

   ```
   $ echo "mp.openapi.scan.disable=true" > APPLICATION_ROOT/src/main/webapp/META-INF/microprofile-config.properties
   ```

4. Rebuild the application:

   ```
   $ mvn clean install
   ```

5. Deploy the application again using the following management CLI commands:

   a. Undeploy the application:

      ```
      undeploy microprofile-openapi.war
      ```

   b. Deploy the application:

      ```
      deploy APPLICATION_ROOT/target/microprofile-openapi.war
      ```

JBoss EAP now serves a static OpenAPI document at the OpenAPI endpoint.

3.7.5. Disabling microprofile-openapi-smallrye

You can disable the **microprofile-openapi-smallrye** subsystem in JBoss EAP XP using the management CLI.

Procedure

- Disable the **microprofile-openapi-smallrye** subsystem:

```
/subsystem=microprofile-openapi-smallrye:remove()
```

3.8. MICROPROFILE REACTIVE MESSAGING ADMINISTRATION

3.8.1. Configuring the required MicroProfile reactive messaging extension and subsystem for JBoss EAP

If you want to enable asynchronous reactive messaging to your instance of JBoss EAP, you must add its extension through the JBoss EAP management CLI.

Prerequisites

- You added the Reactive Streams Operators with SmallRye extension and subsystem. For more information, see MicroProfile Reactive Streams Operators Subsystem Configuration: Required Extension.
- You added the Reactive Messaging with SmallRye extension and subsystem.

Procedure

1. Open the JBoss EAP management CLI.
2. Enter the following code:

```
[standalone@localhost:9990 /] /extension=org.wildfly.extension.microprofile.reactive-messaging-smallrye:add
{"outcome" => "success"}

[standalone@localhost:9990 /] /subsystem=microprofile-reactive-messaging-smallrye:add
{
    "outcome" => "success",
    "response-headers" => {
        "operation-requires-reload" => true,
        "process-state" => "reload-required"
    }
}
```

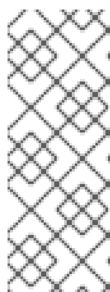

NOTE

If you provision a server using Galleon, either on OpenShift or not, make sure you include the **microprofile-reactive-messaging** Galleon layer to get the core MicroProfile 2.0.1 and reactive messaging functionality, and to enable the required subsystems and extensions. Note that this configuration does not contain the JBoss EAP modules you need to enable Kafka connector functionality. To do this, use the **microprofile-reactive-messaging-kafka** layer.

Verification

You have successfully added the required MicroProfile Reactive Messaging extension and subsystem for JBoss EAP if you see **success** in two places in the resulting code in the management CLI.

TIP

If the resulting code says **reload-required**, you have to reload your server configuration to completely apply all of your changes. To reload, in a standalone server CLI, enter **reload**.

3.9. STANDALONE SERVER CONFIGURATION

3.9.1. Standalone server configuration files

The JBoss EAP XP includes additional standalone server configuration files, **standalone-microprofile.xml** and **standalone-microprofile-ha.xml**.

Standard configuration files that are included with JBoss EAP remain unchanged. Note that JBoss EAP XP 4.0.0 does not support the use of **domain.xml** files or domain mode.

Table 3.1. Standalone configuration files available in JBoss EAP XP

Configuration File	Purpose	Included capabilities	Excluded capabilities
standalone.xml	This is the default configuration that is used when you start your standalone server.	Includes information about the server, including subsystems, networking, deployments, socket bindings, and other configurable details.	Excludes subsystems necessary for messaging or high availability.
standalone-microprofile.xml	This configuration file supports applications that use MicroProfile.	Includes information about the server, including subsystems, networking, deployments, socket bindings, and other configurable details.	Excludes the following capabilities: • Jakarta Enterprise Beans • Messaging • Jakarta EE Batch • Jakarta Server Faces • Jakarta Enterprise Beans timers

Configuration File	Purpose	Included capabilities	Excluded capabilities
standalone-ha.xml		Includes default subsystems and adds the **modcluster** and **jgroups** subsystems for high availability.	Excludes subsystems necessary for messaging.
standalone-microprofile-ha.xml	This standalone file supports applications that use MicroProfile.	Includes the **modcluster** and **jgroups** subsystems for high availability in addition to default subsystems.	Excludes subsystems necessary for messaging.
standalone-full.xml		Includes the **messaging-activemq** and **iiop-openjdk** subsystems in addition to default subsystems.	
standalone-full-ha.xml	Support for every possible subsystem.	Includes subsystems for messaging and high availability in addition to default subsystems.	
standalone-load-balancer.xml	Support for the minimum subsystems necessary to use the built-in mod_cluster front-end load balancer to load balance other JBoss EAP instances.		

By default, starting JBoss EAP as a standalone server uses the **standalone.xml** file. To start JBoss EAP with a standalone MicroProfile configuration, use the **-c** argument. For example,

```
$ EAP_HOME/bin/standalone.sh -c=standalone-microprofile.xml
```

Additional Resources

- Starting and Stopping JBoss EAP
- Configuration Data

3.9.2. Updating standalone configurations with MicroProfile subsystems and extensions

You can update standard standalone server configuration files with MicroProfile subsystems and extensions using the **docs/examples/enable-microprofile.cli** script. The **enable-microprofile.cli** script is intended as an example script for updating standard standalone server configuration files, not custom

configurations.

The **enable-microprofile.cli** script modifies the existing standalone server configuration and adds the following MicroProfile subsystems and extensions if they do not exist in the standalone configuration file:

- **microprofile-config-smallrye**
- **microprofile-fault-tolerance-smallrye**
- **microprofile-health-smallrye**
- **microprofile-jwt-smallrye**
- **microprofile-metrics-smallrye**
- **microprofile-openapi-smallrye**
- **microprofile-opentracing-smallrye**

The **enable-microprofile.cli** script outputs a high-level description of the modifications. The configuration is secured using the **elytron** subsystem. The **security** subsystem, if present, is removed from the configuration.

Prerequisites

- JBoss EAP XP is installed.

Procedure

1. Run the following CLI script to update the default **standalone.xml** server configuration file:

```
$ EAP_HOME/bin/jboss-cli.sh --file=docs/examples/enable-microprofile.cli
```

2. Select a standalone server configuration other than the default **standalone.xml** server configuration file using the following command:

```
$ EAP_HOME/bin/jboss-cli.sh --file=docs/examples/enable-microprofile.cli -Dconfig=
<standalone-full.xml|standalone-ha.xml|standalone-full-ha.xml>
```

3. The specified configuration file now includes MicroProfile subsystems and extensions.

CHAPTER 4. DEVELOP MICROPROFILE APPLICATIONS FOR JBOSS EAP

4.1. MAVEN AND THE JBOSS EAP MICROPROFILE MAVEN REPOSITORY

4.1.1. Downloading the JBoss EAP MicroProfile Maven repository patch as an archive file

Whenever an MicroProfile Expansion Pack is released for JBoss EAP, a corresponding patch is provided for the JBoss EAP MicroProfile Maven repository. This patch is provided as an incremental archive file that is extracted into the existing Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository. The incremental archive file does not overwrite or remove any existing files, so there is no rollback requirement.

Prerequisites

- You have set up an account on the Red Hat Customer Portal.

Procedure

1. Open a browser and log in to the Red Hat Customer Portal.
2. Select **Downloads** from the menu at the top of the page.
3. Find the **Red Hat JBoss Enterprise Application Platform** entry in the list and select it.
4. From the **Product** drop-down list, select **JBoss EAP XP**.
5. From the **Version** drop-down list, select **4.0.0**.
6. Click the **Releases** tab.
7. Find **JBoss EAP XP 4.0.0 Incremental Maven Repository** in the list, and then click **Download**.
8. Save the archive file to your local directory.

Additional Resources

- To learn more about the JBoss EAP Maven repository, see About the Maven Repository in the JBoss EAP *Development Guide*.

4.1.2. Applying the JBoss EAP MicroProfile Maven repository patch on your local system

You can install the JBoss EAP MicroProfile Maven repository patch on your local file system.

When you apply a patch in the form of an incremental archive file to the repository, new files are added to this repository. The incremental archive file does not overwrite or remove any existing files on the repository, so there is no rollback requirement.

Prerequisites

- You have downloaded and installed the Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository on your local system.
    - Check that you have this minor version of the Red Hat JBoss Enterprise Application Platform 7.4 Maven repository installed on your local system.
- You have downloaded the JBoss EAP XP 4.0.0 Incremental Maven Repository on your local system.

Procedure

1. Locate the path to your Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository. For example, **/path/to/repo/jboss-eap-7.4.0.GA-maven-repository/maven-repository/**.
2. Extract the downloaded JBoss EAP XP 4.0.0 Incremental Maven Repository directly into the directory of the Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository. For example, open a terminal and issue the following command, replacing the value for your Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository path:

```
$ unzip -o jboss-eap-xp-4.0.0-incremental-maven-repository.zip -d
EAP_MAVEN_REPOSITORY_PATH
```

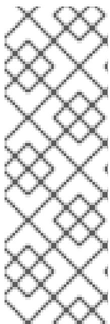

NOTE

The *EAP_MAVEN_REPOSITORY_PATH* points to the **jboss-eap-7.4.0.GA-maven-repository**. For example, this procedure demonstrated the use of the path **/path/to/repo/jboss-eap-7.4.0.GA-maven-repository/**.

After you extract the JBoss EAP XP Incremental Maven repository into the Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository, the repository name becomes JBoss EAP MicroProfile Maven repository.

Additional Resources

- To determine the URL of the JBoss EAP Maven repository, see Determining the URL for the JBoss EAP Maven repository in the JBoss EAP *Development Guide*.

4.1.3. Supported JBoss EAP MicroProfile BOM

JBoss EAP XP 4.0.0 includes the JBoss EAP MicroProfile BOM. This BOM is named **jboss-eap-xp-microprofile**, and its use case supports JBoss EAP MicroProfile APIs.

Table 4.1. JBoss EAP MicroProfile BOM

BOM Artifact ID	Use Case
jboss-eap-xp-microprofile	This BOM, whose **groupId** is **org.jboss.bom**, packages many JBoss EAP MicroProfile supported API dependencies, such as **microprofile-openapi-api** and **microprofile-config-api**. If you use this BOM, you need not specify a version for a supported API dependency, because the **jboss-eap-xp-microprofile** BOM specifies this value for the dependency.

4.1.4. Using the JBoss EAP MicroProfile Maven repository

You can access the **jboss-eap-xp-microprofile** BOM after you install the Red Hat JBoss Enterprise Application Platform 7.4.0.GA Maven repository and apply the JBoss EAP XP Incremental Maven repository to it. The repository name then becomes JBoss EAP MicroProfile Maven repository. The BOM is shipped inside the JBoss EAP XP Incremental Maven repository.

You must configure one of the following to use the JBoss EAP MicroProfile Maven repository:

- The Maven global or user settings
- The project's POM files

Maven settings used with a repository manager or repository on a shared server provide better control and manageability of projects.

You can use an alternative mirror to redirect all lookup requests for a specific repository to your repository manager without changing the project files.

WARNING

Configuring the JBoss EAP MicroProfile Maven repository by modifying the POM file overrides the global and user Maven settings for the configured project.

Prerequisites

- You have installed the Red Hat JBoss Enterprise Application Platform 7.4 Maven repository on your local system, and you have applied the JBoss EAP XP Incremental Maven repository to it.

Procedure

1. Choose a configuration method and configure the JBoss EAP MicroProfile Maven repository.
2. After you have configured the JBoss EAP MicroProfile Maven repository, add the **jboss-eap-xp-microprofile** BOM to the project POM file. The following example shows how to configure the BOM in the **<dependencyManagement>** section of the **pom.xml** file:

```
<dependencyManagement>
  <dependencies>
    ...
    <dependency>
      <groupId>org.jboss.bom</groupId>
      <artifactId>jboss-eap-xp-microprofile</artifactId>
      <version>4.0.0.GA</version>
      <type>pom</type>
      <scope>import</scope>
  </dependency>
    ...
  </dependencies>
</dependencyManagement>
```

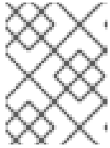

NOTE

If you do not specify a value for the **type** element in the **pom.xml** file, Maven specifies a **jar** value for the element.

Additional Resources

- For more information about selecting methods to configure the JBoss EAP Maven repository, see Use the Maven Repository in the JBoss EAP *Development Guide*.
- For more information about managing dependencies, see Dependency Management.

4.2. MICROPROFILE CONFIG DEVELOPMENT

4.2.1. Creating a Maven project for MicroProfile Config

Create a Maven project with the required dependencies and the directory structure for creating an MicroProfile Config application.

Prerequisites

- Maven is installed.

Procedure

1. Set up the Maven project.

```
$ mvn archetype:generate \
    -DgroupId=com.example \
    -DartifactId=microprofile-config \
    -DinteractiveMode=false \
    -DarchetypeGroupId=org.apache.maven.archetypes \
    -DarchetypeArtifactId=maven-archetype-webapp
cd microprofile-config
```

This creates the directory structure for the project and **pom.xml** configuration file.

2. To let the POM file automatically manage the versions for the MicroProfile Config artifact and the MicroProfile REST Client artifact in the **jboss-eap-xp-microprofile** BOM, import the BOM to the **<dependencyManagement>** section of the project POM file.

```
<dependencyManagement>
  <dependencies>
    <!-- importing the microprofile BOM adds MicroProfile specs -->
    <dependency>
        <groupId>org.jboss.bom</groupId>
        <artifactId>jboss-eap-xp-microprofile</artifactId>
        <version>4.0.0.GA</version>
        <type>pom</type>
        <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

3. Add the MicroProfile Config artifact and the MicroProfile REST Client artifact and other dependencies, managed by the BOM, to the **<dependency>** section of the project POM file. The following example demonstrates adding the MicroProfile Config and the MicroProfile REST Client dependencies to the file:

```
<!-- Add the MicroProfile REST Client API. Set provided for the <scope> tag, as the API is
included in the server. -->
<dependency>
  <groupId>org.eclipse.microprofile.rest.client</groupId>
  <artifactId>microprofile-rest-client-api</artifactId>
  <scope>provided</scope>
</dependency>
<!-- Add the MicroProfile Config API. Set provided for the <scope> tag, as the API is
included in the server. -->
<dependency>
  <groupId>org.eclipse.microprofile.config</groupId>
  <artifactId>microprofile-config-api</artifactId>
  <scope>provided</scope>
</dependency>
<!-- Add the {JAX-RS} API. Set provided for the <scope> tag, as the API is included in the
server. -->
<dependency>
  <groupId>org.jboss.spec.javax.ws.rs</groupId>
  <artifactId>jboss-jaxrs-api_2.1_spec</artifactId>
  <scope>provided</scope>
</dependency>
<!-- Add the CDI API. Set provided for the <scope> tag, as the API is included in the server.
-->
<dependency>
  <groupId>jakarta.enterprise</groupId>
  <artifactId>jakarta.enterprise.cdi-api</artifactId>
  <scope>provided</scope>
</dependency>
```

4.2.2. Using MicroProfile Config property in an application

Create an application that uses a configured **ConfigSource**.

Prerequisites

- MicroProfile Config is enabled in JBoss EAP.
- The latest POM is installed.
- The Maven project is configured for creating an MicroProfile Config application.

Procedure

1. Create the directory to store class files:

```
$ mkdir -p APPLICATION_ROOT/src/main/java/com/example/microprofile/config/
```

Where ***APPLICATION_ROOT*** is the directory containing the **pom.xml** configuration file for the application.

2. Navigate to the new directory:

```
$ cd APPLICATION_ROOT/src/main/java/com/example/microprofile/config/
```

Create all class files described in this procedure in this directory.

3. Create a class file named **HelloApplication.java** with the following content:

```
package com.example.microprofile.config;

import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

@ApplicationPath("/")
public class HelloApplication extends Application {

}
```

This class defines the application as a Jakarta RESTful Web Services application.

4. Create a class file named **HelloService.java** with the following content:

```
package com.example.microprofile.config;

public class HelloService {
 String createHelloMessage(String name){
        return "Hello " + name;
    }
}
```

5. Create a class file named **HelloWorld.java** with the following content:

```
package com.example.microprofile.config;

import javax.inject.Inject;
import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import org.eclipse.microprofile.config.inject.ConfigProperty;

@Path("/config")
public class HelloWorld {

    @Inject
    @ConfigProperty(name="name", defaultValue="jim") (1)
    String name;

    @Inject
    HelloService helloService;

    @GET
    @Path("/json")
    @Produces({ "application/json" })
    public String getHelloWorldJSON() {
        String message = helloService.createHelloMessage(name);
```

```
        return "{\"result\":\"" + message + "\"}";
    }
}
```

1. A MicroProfile Config property is injected into the class with the annotation **@ConfigProperty(name="name", defaultValue="jim")**. If no **ConfigSource** is configured, the value **jim** is returned.

6. Create an empty file named **beans.xml** in the **src/main/webapp/WEB-INF/** directory:

```
$ touch APPLICATION_ROOT/src/main/webapp/WEB-INF/beans.xml
```

   Where ***APPLICATION_ROOT*** is the directory containing the **pom.xml** configuration file for the application.

7. Navigate to the root directory of the application:

```
$ cd APPLICATION_ROOT
```

   Where ***APPLICATION_ROOT*** is the directory containing the **pom.xml** configuration file for the application.

8. Build the project:

```
$ mvn clean install wildfly:deploy
```

9. Test the output:

```
$ curl http://localhost:8080/microprofile-config/config/json
```

   The following is the expected output:

```
{"result":"Hello jim"}
```

4.3. MICROPROFILE FAULT TOLERANCE APPLICATION DEVELOPMENT

4.3.1. Adding the MicroProfile Fault Tolerance extension

The MicroProfile Fault Tolerance extension is included in **standalone-microprofile.xml** and **standalone-microprofile-ha.xml** configurations that are provided as part of JBoss EAP XP.

The extension is not included in the standard **standalone.xml** configuration. To use the extension, you must manually enable it.

Prerequisites

- EAP XP pack is installed.

Procedure

1. Add the MicroProfile Fault Tolerance extension using the following management CLI command:

```
/extension=org.wildfly.extension.microprofile.fault-tolerance-smallrye:add
```

2. Enable the **microprofile-fault-tolerance-smallrye** subsystem using the following managenent command:

```
/subsystem=microprofile-fault-tolerance-smallrye:add
```

3. Reload the server with the following management command:

```
reload
```

4.3.2. Configuring Maven project for MicroProfile Fault Tolerance

Create a Maven project with the required dependencies and the directory structure for creating an MicroProfile Fault Tolerance application.

Prerequisites

- Maven is installed.

Procedure

1. Set up the Maven project:

```
mvn archetype:generate \
    -DgroupId=com.example.microprofile.faulttolerance \
    -DartifactId=microprofile-fault-tolerance \
    -DarchetypeGroupId=org.apache.maven.archetypes \
    -DarchetypeArtifactId=maven-archetype-webapp \
    -DinteractiveMode=false
cd microprofile-fault-tolerance
```

The command creates the directory structure for the project and the **pom.xml** configuration file.

2. To let the POM file automatically manage the versions for the MicroProfile Fault Tolerance artifact in the **jboss-eap-xp-microprofile** BOM, import the BOM to the **<dependencyManagement>** section of the project POM file.

```
<dependencyManagement>
  <dependencies>
    <!-- importing the microprofile BOM adds MicroProfile specs -->
    <dependency>
        <groupId>org.jboss.bom</groupId>
        <artifactId>jboss-eap-xp-microprofile</artifactId>
        <version>${version.microprofile.bom}</version>
        <type>pom</type>
        <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

Replace ${version.microprofile.bom} with the installed version of BOM.

3. Add the MicroProfile Fault Tolerance artifact, managed by the BOM, to the **<dependency>** section of the project POM file. The following example demonstrates adding the MicroProfile Fault Tolerance dependency to the file:

```
<!-- Add the MicroProfile Fault Tolerance API. Set provided for the <scope> tag, as the API
is included in the server. -->
<dependency>
  <groupId>org.eclipse.microprofile.fault.tolerance</groupId>
  <artifactId>microprofile-fault-tolerance-api</artifactId>
  <scope>provided</scope>
</dependency>
```

4.3.3. Creating a fault tolerant application

Create a fault-tolerant application that implements retry, timeout, and fallback patterns for fault tolerance.

Prerequisites

- Maven dependencies have been configured.

Procedure

1. Create the directory to store class files:

```
$ mkdir -p APPLICATION_ROOT/src/main/java/com/example/microprofile/faulttolerance
```

APPLICATION_ROOT is the directory containing the **pom.xml** configuration file for the application.

2. Navigate to the new directory:

```
$ cd APPLICATION_ROOT/src/main/java/com/example/microprofile/faulttolerance
```

For the following steps, create all class files in the new directory.

3. Create a simple entity representing a coffee sample as **Coffee.java** with the following content:

```
package com.example.microprofile.faulttolerance;

public class Coffee {

    public Integer id;
    public String name;
    public String countryOfOrigin;
    public Integer price;

    public Coffee() {
    }

    public Coffee(Integer id, String name, String countryOfOrigin, Integer price) {
        this.id = id;
        this.name = name;
        this.countryOfOrigin = countryOfOrigin;
```

```
        this.price = price;
    }
}
```

4. Create a class file **CoffeeApplication.java** with the following content:

```
package com.example.microprofile.faulttolerance;

import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

@ApplicationPath("/")
public class CoffeeApplication extends Application {
}
```

5. Create a Jakarta Contexts and Dependency Injection Bean as **CoffeeRepositoryService.java** with the following content:

```
package com.example.microprofile.faulttolerance;

import java.util.ArrayList;
import java.util.Collections;
import java.util.HashMap;
import java.util.List;
import java.util.Map;
import java.util.stream.Collectors;
import javax.enterprise.context.ApplicationScoped;

@ApplicationScoped
public class CoffeeRepositoryService {

    private Map<Integer, Coffee> coffeeList = new HashMap<>();

    public CoffeeRepositoryService() {
        coffeeList.put(1, new Coffee(1, "Fernandez Espresso", "Colombia", 23));
        coffeeList.put(2, new Coffee(2, "La Scala Whole Beans", "Bolivia", 18));
        coffeeList.put(3, new Coffee(3, "Dak Lak Filter", "Vietnam", 25));
    }

    public List<Coffee> getAllCoffees() {
        return new ArrayList<>(coffeeList.values());
    }

    public Coffee getCoffeeById(Integer id) {
        return coffeeList.get(id);
    }

    public List<Coffee> getRecommendations(Integer id) {
        if (id == null) {
            return Collections.emptyList();
        }
        return coffeeList.values().stream()
                .filter(coffee -> !id.equals(coffee.id))
                .limit(2)
```

```
                .collect(Collectors.toList());
    }
}
```

6. Create a class file **CoffeeResource.java** with the following content:

```
package com.example.microprofile.faulttolerance;

import java.util.List;
import java.util.Random;
import java.util.concurrent.atomic.AtomicLong;
import javax.inject.Inject;
import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import javax.ws.rs.core.MediaType;
import java.util.Collections;
import javax.ws.rs.PathParam;
import org.eclipse.microprofile.faulttolerance.Fallback;
import org.eclipse.microprofile.faulttolerance.Timeout;
import org.eclipse.microprofile.faulttolerance.Retry;

@Path("/coffee")
@Produces(MediaType.APPLICATION_JSON)
public class CoffeeResource {

    @Inject
    private CoffeeRepositoryService coffeeRepository;

    private AtomicLong counter = new AtomicLong(0);

    @GET
    @Retry(maxRetries = 4) 1
    public List<Coffee> coffees() {
        final Long invocationNumber = counter.getAndIncrement();
        return coffeeRepository.getAllCoffees();
    }


    @GET
    @Path("/{id}/recommendations")
    @Timeout(250) 2
    public List<Coffee> recommendations(@PathParam("id") int id) {
            return coffeeRepository.getRecommendations(id);
        }

    @GET
    @Path("fallback/{id}/recommendations")
    @Fallback(fallbackMethod = "fallbackRecommendations") 3
    public List<Coffee> recommendations2(@PathParam("id") int id) {
        return coffeeRepository.getRecommendations(id);
        }

    public List<Coffee> fallbackRecommendations(int id) {
        //always return a default coffee
```

```
        return Collections.singletonList(coffeeRepository.getCoffeeById(1));
    }
}
```

1. Define number of re-tries to **4**.
2. Define the timeout interval in milliseconds.
3. Define a fallback method to call when invocation fails.

7. Navigate to the root directory of the application:

```
$ cd APPLICATION_ROOT
```

8. Build the application using the following Maven command:

```
$ mvn clean install wildfly:deploy
```

Access the application at **http://localhost:8080/microprofile-fault-tolerance/coffee**.

Additional Resources

- For a detailed example of fault tolerant application, which includes artificial failures to test the fault tolerance of the application, see the **microprofile-fault-tolerance** quickstart.

4.4. MICROPROFILE HEALTH DEVELOPMENT

4.4.1. The custom health check example

The default implementation provided by the **microprofile-health-smallrye** subsystem performs a basic health check. For more detailed information, on either the server or application status, custom health checks might be included. Any Jakarta Contexts and Dependency Injection beans that include the **org.eclipse.microprofile.health.Liveness**, **org.eclipse.microprofile.health.Readiness**, or **org.eclipse.microprofile.health.Startup** annotations at the class level are automatically discovered and invoked at runtime.

The following example demonstrates how to create a new implementation of a health check that returns an **UP** state.

```
import org.eclipse.microprofile.health.HealthCheck;
import org.eclipse.microprofile.health.HealthCheckResponse;
import org.eclipse.microprofile.health.Liveness;

import javax.enterprise.context.ApplicationScoped;

@Liveness
@ApplicationScoped
public class HealthTest implements HealthCheck {

    @Override
    public HealthCheckResponse call() {
```

```
        return HealthCheckResponse.named("health-test").up().build();
    }
}
```

After you deploy a health check, any subsequent health check queries include the custom checks, as demonstrated in the following example.

```
[standalone@localhost:9990 /] /subsystem=microprofile-health-smallrye:check
{
    "outcome" => "success",
    "result" => {
        "status" => "UP",
        "checks" => [
            {
                "name" => "deployments-status",
                "status" => "UP",
                "data" => {"<deployment_name>.war" => "OK"}
            },
            {
                "name" => "server-state",
                "status" => "UP",
                "data" => {"value" => "running"}
            },
            {
                "name" => "boot-errors",
                "status" => "UP"
            },
            {
                "name" => "health-test",
                "status" => "UP"
            },
            {
                "name" => "ready-deployment.<deployment_name>.war,
                "status" => "UP"
            },
            {
                "name" => "started-deployment.<deployment_name>.war",
                "status" => "UP"
            }
        ]
    }
}
```

NOTE

You can use the following commands for liveness, readiness, and startup checks:

- **/subsystem=microprofile-health-smallrye:check-live**
- **/subsystem=microprofile-health-smallrye:check-ready**
- **/subsystem=microprofile-health-smallrye:check-started**

4.4.2. The @Liveness annotation example

The following example demonstrates how to use the **@Liveness** annotation in an application.

```
@Liveness
@ApplicationScoped
public class DataHealthCheck implements HealthCheck {

   @Override
   public HealthCheckResponse call() {
      return HealthCheckResponse.named("Health check with data")
         .up()
         .withData("foo", "fooValue")
         .withData("bar", "barValue")
         .build();
   }
}
```

4.4.3. The @Readiness annotation example

The following example demonstrates how to check the connection to a database. If the database is down, the readiness check reports an error.

```
@Readiness
@ApplicationScoped
public class DatabaseConnectionHealthCheck implements HealthCheck {

   @Inject
   @ConfigProperty(name = "database.up", defaultValue = "false")
   private boolean databaseUp;

   @Override
   public HealthCheckResponse call() {

      HealthCheckResponseBuilder responseBuilder = HealthCheckResponse.named("Database
connection health check");

      try {
         simulateDatabaseConnectionVerification();
         responseBuilder.up();
      } catch (IllegalStateException e) {
         // cannot access the database
         responseBuilder.down()
            .withData("error", e.getMessage()); // pass the exception message
      }

      return responseBuilder.build();
   }

   private void simulateDatabaseConnectionVerification() {
      if (!databaseUp) {
         throw new IllegalStateException("Cannot contact database");
      }
   }
}
```

4.4.4. The @Startup annotation example

The following is an example of using the **@Startup** annotation in an application.

```
@Startup
@ApplicationScoped
public class StartupHealthCheck implements HealthCheck {

    @Override
    public HealthCheckResponse call() {
        return HealthCheckResponse.up("Application started");
    }
}
```

4.5. MICROPROFILE JWT APPLICATION DEVELOPMENT

4.5.1. Enabling microprofile-jwt-smallrye subsystem

The MicroProfile JWT integration is provided by the **microprofile-jwt-smallrye** subsystem and is included in the default configuration. If the subsystem is not present in the default configuration, you can add it as follows.

Prerequisites

- EAP XP is installed.

Procedure

1. Enable the MicroProfile JWT smallrye extension in JBoss EAP:

   ```
   /extension=org.wildfly.extension.microprofile.jwt-smallrye:add
   ```

2. Enable the **microprofile-jwt-smallrye** subsystem:

   ```
   /subsystem=microprofile-jwt-smallrye:add
   ```

3. Reload the server:

   ```
   reload
   ```

The **microprofile-jwt-smallrye** subsystem is enabled.

4.5.2. Configuring Maven project for developing JWT applications

Create a Maven project with the required dependencies and the directory structure for developing a JWT application.

Prerequisites

- Maven is installed.
- **microprofile-jwt-smallrye** subsystem is enabled.

Procedure

1. Set up the maven project:

```
$ mvn archetype:generate -DinteractiveMode=false \
    -DarchetypeGroupId=org.apache.maven.archetypes \
    -DarchetypeArtifactId=maven-archetype-webapp \
    -DgroupId=com.example -DartifactId=microprofile-jwt \
    -Dversion=1.0.0.Alpha1-SNAPSHOT
  cd microprofile-jwt
```

   The command creates the directory structure for the project and the **pom.xml** configuration file.

2. To let the POM file automatically manage the versions for the MicroProfile JWT artifact in the **jboss-eap-xp-microprofile** BOM, import the BOM to the **<dependencyManagement>** section of the project POM file.

```
<dependencyManagement>
  <dependencies>
    <!-- importing the microprofile BOM adds MicroProfile specs -->
    <dependency>
        <groupId>org.jboss.bom</groupId>
        <artifactId>jboss-eap-xp-microprofile</artifactId>
        <version>${version.microprofile.bom}</version>
        <type>pom</type>
        <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

   Replace ${version.microprofile.bom} with the installed version of BOM.

3. Add the MicroProfile JWT artifact, managed by the BOM, to the **<dependency>** section of the project POM file. The following example demonstrates adding the MicroProfile JWT dependency to the file:

```
<!-- Add the MicroProfile JWT API. Set provided for the <scope> tag, as the API is included in the server. -->
<dependency>
  <groupId>org.eclipse.microprofile.jwt</groupId>
  <artifactId>microprofile-jwt-auth-api</artifactId>
  <scope>provided</scope>
</dependency>
```

4.5.3. Creating an application with MicroProfile JWT

Create an application that authenticates requests based on JWT tokens and implements authorization based on the identity of the token bearer.

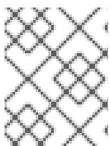

NOTE

The following procedure provides example code for generating tokens. For a production environment, use an identity provider such as Red Hat single sign-on (SSO).

Prerequisites

- Maven project is configured with the correct dependencies.

Procedure

1. Create a token generator.
   This step serves as a reference. For a production environment, use an identity provider such as Red Hat SSO.

   a. Create a directory **src/test/java** for token the generator utility and navigate to it:

```
$ mkdir -p src/test/java
$ cd src/test/java
```

   b. Create a class file **TokenUtil.java** with the following content:

```
package  com.example.mpjwt;

import java.io.FileInputStream;
import java.io.InputStream;
import java.nio.charset.StandardCharsets;
import java.security.KeyFactory;
import java.security.PrivateKey;
import java.security.spec.PKCS8EncodedKeySpec;
import java.util.Base64;
import java.util.UUID;

import javax.json.Json;
import javax.json.JsonArrayBuilder;
import javax.json.JsonObjectBuilder;

import com.nimbusds.jose.JOSEObjectType;
import com.nimbusds.jose.JWSAlgorithm;
import com.nimbusds.jose.JWSHeader;
import com.nimbusds.jose.JWSObject;
import com.nimbusds.jose.JWSSigner;
import com.nimbusds.jose.Payload;
import com.nimbusds.jose.crypto.RSASSASigner;

public class TokenUtil {

    private static PrivateKey loadPrivateKey(final String fileName) throws Exception {
        try (InputStream is = new FileInputStream(fileName)) {
            byte[] contents = new byte[4096];
            int length = is.read(contents);
            String rawKey = new String(contents, 0, length, StandardCharsets.UTF_8)
                    .replaceAll("-----BEGIN (.*)-----", "")
                    .replaceAll("-----END (.*)----", "")
                    .replaceAll("\r\n", "").replaceAll("\n", "").trim();

            PKCS8EncodedKeySpec keySpec = new
PKCS8EncodedKeySpec(Base64.getDecoder().decode(rawKey));
            KeyFactory keyFactory = KeyFactory.getInstance("RSA");

            return keyFactory.generatePrivate(keySpec);
```

```
        }
    }

    public static String generateJWT(final String principal, final String birthdate, final
String...groups) throws Exception {
     PrivateKey privateKey = loadPrivateKey("private.pem");

        JWSSigner signer = new RSASSASigner(privateKey);
        JsonArrayBuilder groupsBuilder = Json.createArrayBuilder();
        for (String group : groups) { groupsBuilder.add(group); }

        long currentTime = System.currentTimeMillis() / 1000;
        JsonObjectBuilder claimsBuilder = Json.createObjectBuilder()
                .add("sub", principal)
                .add("upn", principal)
                .add("iss", "quickstart-jwt-issuer")
                .add("aud", "jwt-audience")
                .add("groups", groupsBuilder.build())
                .add("birthdate", birthdate)
                .add("jti", UUID.randomUUID().toString())
                .add("iat", currentTime)
                .add("exp", currentTime + 14400);

        JWSObject jwsObject = new JWSObject(new
JWSHeader.Builder(JWSAlgorithm.RS256)
                .type(new JOSEObjectType("jwt"))
                .keyID("Test Key").build(),
                new Payload(claimsBuilder.build().toString()));

        jwsObject.sign(signer);

        return jwsObject.serialize();
    }

    public static void main(String[] args) throws Exception {
        if (args.length < 2) throw new IllegalArgumentException("Usage TokenUtil {principal}
{birthdate} {groups}");
        String principal = args[0];
        String birthdate = args[1];
        String[] groups = new String[args.length - 2];
        System.arraycopy(args, 2, groups, 0, groups.length);

        String token = generateJWT(principal, birthdate, groups);
        String[] parts = token.split("\\.");
        System.out.println(String.format("\nJWT Header - %s", new
String(Base64.getDecoder().decode(parts[0]), StandardCharsets.UTF_8)));
        System.out.println(String.format("\nJWT Claims - %s", new
String(Base64.getDecoder().decode(parts[1]), StandardCharsets.UTF_8)));
        System.out.println(String.format("\nGenerated JWT Token \n%s\n", token));
    }
}
```

2. Create the **web.xml** file in the **src/main/webapp/WEB-INF** directory with the following content:

```
<context-param>
    <param-name>resteasy.role.based.security</param-name>
```

```
    <param-value>true</param-value>
</context-param>

<security-role>
    <role-name>Subscriber</role-name>
</security-role>
```

3. Create a class file **SampleEndPoint.java** with the following content:

```
package com.example.mpjwt;

import javax.ws.rs.GET;
import javax.ws.rs.Path;

import java.security.Principal;
import javax.ws.rs.core.Context;
import javax.ws.rs.core.SecurityContext;

import javax.annotation.security.RolesAllowed;
import javax.inject.Inject;

import java.time.LocalDate;
import java.time.Period;
import java.util.Optional;

import org.eclipse.microprofile.jwt.Claims;
import org.eclipse.microprofile.jwt.Claim;

import org.eclipse.microprofile.jwt.JsonWebToken;

@Path("/Sample")
public class SampleEndPoint {

    @GET
    @Path("/helloworld")
    public String helloworld(@Context SecurityContext securityContext) {
        Principal principal = securityContext.getUserPrincipal();
        String caller = principal == null ? "anonymous" : principal.getName();

        return "Hello " + caller;
    }

    @Inject
  JsonWebToken jwt;

  @GET()
  @Path("/subscription")
  @RolesAllowed({"Subscriber"})
  public String helloRolesAllowed(@Context SecurityContext ctx) {
      Principal caller =  ctx.getUserPrincipal();
      String name = caller == null ? "anonymous" : caller.getName();
      boolean hasJWT = jwt.getClaimNames() != null;
      String helloReply = String.format("hello + %s, hasJWT: %s", name, hasJWT);

      return helloReply;
  }
```

```
@Inject
@Claim(standard = Claims.birthdate)
Optional<String> birthdate;

@GET()
@Path("/birthday")
@RolesAllowed({ "Subscriber" })
public String birthday() {
    if (birthdate.isPresent()) {
        LocalDate birthdate = LocalDate.parse(this.birthdate.get().toString());
        LocalDate today = LocalDate.now();
        LocalDate next = birthdate.withYear(today.getYear());
        if (today.equals(next)) {
            return "Happy Birthday";
        }
        if (next.isBefore(today)) {
            next = next.withYear(next.getYear() + 1);
        }

        Period wait = today.until(next);

        return String.format("%d months and %d days until your next birthday.",
wait.getMonths(), wait.getDays());
    }

    return "Sorry, we don't know your birthdate.";

}

}
```

The methods annotated with **@Path** are the Jakarta RESTful Web Services endpoints.

The annotation **@Claim** defines a JWT claim.

4. Create a class file **App.java** to enable Jakarta RESTful Web Services:

```
package com.example.mpjwt;

import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

import org.eclipse.microprofile.auth.LoginConfig;

@ApplicationPath("/rest")
@LoginConfig(authMethod="MP-JWT", realmName="MP JWT Realm")
public class App extends Application {}
```

The annotation **@LoginConfig(authMethod="MP-JWT", realmName="MP JWT Realm")** enables JWT RBAC during deployment.

5. Compile the application with the following Maven command:

```
$ mvn package
```

6. Generate JWT token using the token generator utility:

```
$ mvn exec:java -Dexec.mainClass=org.wildfly.quickstarts.mpjwt.TokenUtil -
Dexec.classpathScope=test -Dexec.args="testUser 2017-09-15 Echoer Subscriber"
```

7. Build and deploy the application using the following Maven command:

```
$ mvn package wildfly:deploy
```

8. Test the application.
   - Call the **Sample/subscription** endpoint using the bearer token:

   ```
   $ curl -H "Authorization: Bearer ey..rg" http://localhost:8080/microprofile-
   jwt/rest/Sample/subscription
   ```

   - Call the **Sample/birthday** endpoint:

   ```
   $ curl -H "Authorization: Bearer ey..rg" http://localhost:8080/microprofile-
   jwt/rest/Sample/birthday
   ```

4.6. MICROPROFILE METRICS DEVELOPMENT

4.6.1. Creating an MicroProfile Metrics application

Create an application that returns the number of requests made to the application.

Procedure

1. Create a class file **HelloService.java** with the following content:

```
package com.example.microprofile.metrics;

public class HelloService {
    String createHelloMessage(String name){
        return "Hello" + name;
    }
}
```

2. Create a class file **HelloWorld.java** with the following content:

```
package com.example.microprofile.metrics;

import javax.inject.Inject;
import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import org.eclipse.microprofile.metrics.annotation.Counted;

@Path("/")
public class HelloWorld {
@Inject
    HelloService helloService;
```

```
@GET
@Path("/json")
   @Produces({ "application/json" })
   @Counted(name = "requestCount",
    absolute = true,
description = "Number of times the getHelloWorldJSON was requested")
   public String getHelloWorldJSON() {
      return "{\"result\":\"" + helloService.createHelloMessage("World") + "\"}";
   }
}
```

3. Update the **pom.xml** file to include the following dependency:

```
<dependency>
   <groupId>org.eclipse.microprofile.metrics</groupId>
   <artifactId>microprofile-metrics-api</artifactId>
   <scope>provided</scope>
</dependency>
```

4. Build the application using the following Maven command:

```
$ mvn clean install wildfly:deploy
```

5. Test the metrics:

   a. Issue the following command in the CLI:

   ```
   $ curl -v http://localhost:9990/metrics |  grep request_count | grep helloworld-rs-metrics
   ```

   Expected output:

   ```
   jboss_undertow_request_count_total{deployment="helloworld-rs-
   metrics.war",servlet="org.jboss.as.quickstarts.rshelloworld.JAXActivator",subdeployment="h
   elloworld-rs-metrics.war",microprofile_scope="vendor"} 0.0
   ```

   b. In a browser, navigate to the URL http://localhost:8080/helloworld-rs/rest/json.

   c. Re-Issue the following command in the CLI:

   ```
   $ curl -v http://localhost:9990/metrics |  grep request_count | grep helloworld-rs-metrics
   ```

   Expected output:

   ```
   jboss_undertow_request_count_total{deployment="helloworld-rs-
   metrics.war",servlet="org.jboss.as.quickstarts.rshelloworld.JAXActivator",subdeployment="h
   elloworld-rs-metrics.war",microprofile_scope="vendor"} 1.0
   ```

4.7. DEVELOPING AN MICROPROFILE OPENAPI APPLICATION

4.7.1. Enabling MicroProfile OpenAPI

The **microprofile-openapi-smallrye** subsystem is provided in the **standalone-microprofile.xml** configuration. However, JBoss EAP XP uses the **standalone.xml** by default. You must include the subsystem in **standalone.xml** to use it.

Alternatively, you can follow the procedure Updating standalone configurations with MicroProfile subsystems and extensions to update the **standalone.xml** configuration file.

Procedure

1. Enable the MicroProfile OpenAPI smallrye extension in JBoss EAP:

```
/extension=org.wildfly.extension.microprofile.openapi-smallrye:add()
```

2. Enable the **microprofile-openapi-smallrye** subsystem using the following management command:

```
/subsystem=microprofile-openapi-smallrye:add()
```

3. Reload the server.

```
reload
```

The **microprofile-openapi-smallrye** subsystem is enabled.

4.7.2. Configuring Maven project for MicroProfile OpenAPI

Create a Maven project to set up the dependencies for creating an MicroProfile OpenAPI application.

Prerequisites

- Maven is installed.
- JBoss EAP Maven repository is configured.

Procedure

1. Initialize the project:

```
mvn archetype:generate \
    -DgroupId=com.example.microprofile.openapi \
    -DartifactId=microprofile-openapi\
    -DarchetypeGroupId=org.apache.maven.archetypes \
    -DarchetypeArtifactId=maven-archetype-webapp \
    -DinteractiveMode=false
cd microprofile-openapi
```

The command creates the directory structure for the project and the **pom.xml** configuration file.

2. Edit the **pom.xml** configuration file to contain:

```
<?xml version="1.0" encoding="UTF-8"?>

<project xmlns="http://maven.apache.org/POM/4.0.0"
```

```
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:schemaLocation="http://maven.apache.org/POM/4.0.0
http://maven.apache.org/xsd/maven-4.0.0.xsd">
    <modelVersion>4.0.0</modelVersion>

    <groupId>com.example.microprofile.openapi</groupId>
    <artifactId>microprofile-openapi</artifactId>
    <version>1.0-SNAPSHOT</version>
    <packaging>war</packaging>

    <name>microprofile-openapi Maven Webapp</name>
    <!-- Update the value with the URL of the project -->
    <url>http://www.example.com</url>

    <properties>
        <project.build.sourceEncoding>UTF-8</project.build.sourceEncoding>
        <maven.compiler.source>1.8</maven.compiler.source>
        <maven.compiler.target>1.8</maven.compiler.target>
        <version.server.bom>4.0.0.GA</version.server.bom>
    </properties>

    <dependencyManagement>
        <dependencies>
            <dependency>
                <groupId>org.jboss.bom</groupId>
                <artifactId>jboss-eap-xp-microprofile</artifactId>
                <version>${version.server.bom}</version>
                <type>pom</type>
                <scope>import</scope>
            </dependency>
        </dependencies>
    </dependencyManagement>

    <dependencies>
        <dependency>
            <groupId>org.jboss.spec.javax.ws.rs</groupId>
            <artifactId>jboss-jaxrs-api_2.1_spec</artifactId>
            <scope>provided</scope>
        </dependency>
    </dependencies>

    <build>
        <!-- Set the name of the archive -->
        <finalName>${project.artifactId}</finalName>
        <plugins>
            <plugin>
                <artifactId>maven-clean-plugin</artifactId>
                <version>3.1.0</version>
            </plugin>
            <!-- see http://maven.apache.org/ref/current/maven-core/default-
bindings.html#Plugin_bindings_for_war_packaging -->
            <plugin>
                <artifactId>maven-resources-plugin</artifactId>
                <version>3.0.2</version>
            </plugin>
            <plugin>
```

```
                <artifactId>maven-compiler-plugin</artifactId>
                <version>3.8.0</version>
            </plugin>
            <plugin>
                <artifactId>maven-surefire-plugin</artifactId>
                <version>2.22.1</version>
            </plugin>
            <plugin>
                <artifactId>maven-war-plugin</artifactId>
                <version>3.2.2</version>
            </plugin>
            <plugin>
                <artifactId>maven-install-plugin</artifactId>
                <version>2.5.2</version>
            </plugin>
            <plugin>
                <artifactId>maven-deploy-plugin</artifactId>
                <version>2.8.2</version>
            </plugin>
            <!-- Allows to use mvn wildfly:deploy -->
            <plugin>
                <groupId>org.wildfly.plugins</groupId>
                <artifactId>wildfly-maven-plugin</artifactId>
            </plugin>
        </plugins>
    </build>
</project>
```

Use the **pom.xml** configuration file and directory structure to create an application.

Additional resources

- For information about configuring the JBoss EAP Maven repository, see Configuring the JBoss EAP Maven repository with the POM file.

4.7.3. Creating an MicroProfile OpenAPI application

Create an application that returns an OpenAPI v3 document.

Prerequisites

- Maven project is configured for creating an MicroProfile OpenAPI application.

Procedure

1. Create the directory to store class files:

   ```
   $ mkdir -p APPLICATION_ROOT/src/main/java/com/example/microprofile/openapi/
   ```

   APPLICATION_ROOT is the directory containing the **pom.xml** configuration file for the application.

2. Navigate to the new directory:

   ```
   $ cd APPLICATION_ROOT/src/main/java/com/example/microprofile/openapi/
   ```

All the class files in the following steps must be created in this directory.

3. Create the class file **InventoryApplication.java** with the following content:

```
package com.example.microprofile.openapi;

import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

@ApplicationPath("/inventory")
public class InventoryApplication extends Application {
}
```

This class serves as the REST endpoint for the application.

4. Create a class file **Fruit.java** with the following content:

```
package com.example.microprofile.openapi;

public class Fruit {

    private final String name;
    private final String description;

    public Fruit(String name, String description) {
        this.name = name;
        this.description = description;
    }

    public String getName() {
        return this.name;
    }

    public String getDescription() {
        return this.description;
    }
}
```

5. Create a class file **FruitResource.java** with the following content:

```
package com.example.microprofile.openapi;

import java.util.Collections;
import java.util.LinkedHashMap;
import java.util.Set;

import javax.ws.rs.Consumes;
import javax.ws.rs.DELETE;
import javax.ws.rs.GET;
import javax.ws.rs.POST;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import javax.ws.rs.core.MediaType;

@Path("/fruit")
```

```
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
public class FruitResource {

    private final Set<Fruit> fruits =
Collections.newSetFromMap(Collections.synchronizedMap(new LinkedHashMap<>()));

    public FruitResource() {
        this.fruits.add(new Fruit("Apple", "Winter fruit"));
        this.fruits.add(new Fruit("Pineapple", "Tropical fruit"));
    }

    @GET
    public Set<Fruit> all() {
        return this.fruits;
    }

    @POST
    public Set<Fruit> add(Fruit fruit) {
        this.fruits.add(fruit);
        return this.fruits;
    }

    @DELETE
    public Set<Fruit> remove(Fruit fruit) {
        this.fruits.removeIf(existingFruit ->
existingFruit.getName().contentEquals(fruit.getName()));
        return this.fruits;
    }
}
```

6. Navigate to the root directory of the application:

```
$ cd APPLICATION_ROOT
```

7. Build and deploy the application using the following Maven command:

```
$ mvn wildfly:deploy
```

8. Test the application.

   - Access the OpenAPI documentation of the sample application using **curl**:

   ```
   $ curl http://localhost:8080/openapi
   ```

   - The following output is returned:

   ```
   openapi: 3.0.1
   info:
     title: Archetype Created Web Application
     version: "1.0"
   servers:
   - url: /microprofile-openapi
   paths:
     /inventory/fruit:
   ```

```
    get:
      responses:
        "200":
          description: OK
          content:
            application/json:
              schema:
                type: array
                items:
                  $ref: '#/components/schemas/Fruit'
    post:
      requestBody:
        content:
          application/json:
            schema:
              $ref: '#/components/schemas/Fruit'
      responses:
        "200":
          description: OK
          content:
            application/json:
              schema:
                type: array
                items:
                  $ref: '#/components/schemas/Fruit'
    delete:
      requestBody:
        content:
          application/json:
            schema:
              $ref: '#/components/schemas/Fruit'
      responses:
        "200":
          description: OK
          content:
            application/json:
              schema:
                type: array
                items:
                  $ref: '#/components/schemas/Fruit'
components:
  schemas:
    Fruit:
      type: object
      properties:
        description:
          type: string
        name:
          type: string
```


Additional Resources

- For a list of annotations defined in MicroProfile SmallRye OpenAPI, see MicroProfile OpenAPI annotations.

4.7.4. Configuring JBoss EAP to serve a static OpenAPI document

Configure JBoss EAP to serve a static OpenAPI document that describes the REST services for the host.

When JBoss EAP is configured to serve a static OpenAPI document, the static OpenAPI document is processed before any Jakarta RESTful Web Services and MicroProfile OpenAPI annotations.

In a production environment, disable annotation processing when serving a static document. Disabling annotation processing ensures that an immutable and versioned API contract is available for clients.

Procedure

1. Create a directory in the application source tree:

```
$ mkdir APPLICATION_ROOT/src/main/webapp/META-INF
```

   APPLICATION_ROOT is the directory containing the **pom.xml** configuration file for the application.

2. Query the OpenAPI endpoint, redirecting the output to a file:

```
$ curl http://localhost:8080/openapi?format=JSON > src/main/webapp/META-INF/openapi.json
```

   By default, the endpoint serves a YAML document, **format=JSON** specifies that a JSON document is returned.

3. Configure the application to skip annotation scanning when processing the OpenAPI document model:

```
$ echo "mp.openapi.scan.disable=true" > APPLICATION_ROOT/src/main/webapp/META-INF/microprofile-config.properties
```

4. Rebuild the application:

```
$ mvn clean install
```

5. Deploy the application again using the following management CLI commands:

   a. Undeploy the application:

```
undeploy microprofile-openapi.war
```

   b. Deploy the application:

```
deploy APPLICATION_ROOT/target/microprofile-openapi.war
```

JBoss EAP now serves a static OpenAPI document at the OpenAPI endpoint.

4.8. MICROPROFILE REST CLIENT DEVELOPMENT

4.8.1. A comparison of MicroProfile REST client and Jakarta RESTful Web Services syntaxes

The MicroProfile REST client enables a version of distributed object communication, which is also implemented in CORBA, Java Remote Method Invocation (RMI), the JBoss Remoting Project, and RESTEasy. For example, consider the resource:

```
@Path("resource")
public class TestResource {
  @Path("test")
  @GET
  String test() {
    return "test";
  }
 }
```

The following example demonstrates the use of the Jakarta RESTful Web Services-native way to access the **TestResource** class:

```
Client client = ClientBuilder.newClient();
String response = client.target("http://localhost:8081/test").request().get(String.class);
```

However, Microprofile REST client supports a more intuitive syntax by directly calling the **test()** method, as the following example demonstrates:

```
@Path("resource")
public interface TestResourceIntf {
   @Path("test")
   @GET
   public String test();
}

TestResourceIntf service = RestClientBuilder.newBuilder()
                              .baseUrl(http://localhost:8081/))
                              .build(TestResourceIntf.class);
String s = service.test();
```

In the preceding example, making calls on the **TestResource** class becomes much easier with the **TestResourceIntf** class, as illustrated by the call **service.test()**.

The following example is a more elaborate version of the **TestResourceIntf** class:

```
@Path("resource")
public interface TestResourceIntf2 {
  @Path("test/{path}")
  @Consumes("text/plain")
  @Produces("text/html")
  @POST
  public String test(@PathParam("path") String path, @QueryParam("query") String query, String
entity);
}
```

Calling the **service.test("p", "q", "e")** method results in an HTTP message as shown in the following example:

```
POST /resource/test/p/?query=q HTTP/1.1
Accept: text/html
Content-Type: text/plain
Content-Length: 1

e
```

4.8.2. Programmatic registration of providers in MicroProfile REST client

With the MicroProfile REST client, you can configure the client environment by registering providers. For example:

```
TestResourceIntf service = RestClientBuilder.newBuilder()
                    .baseUrl(http://localhost:8081/))
                    .register(MyClientResponseFilter.class)
                    .register(MyMessageBodyReader.class)
                    .build(TestResourceIntf.class);
```

4.8.3. Declarative registration of providers in MicroProfile REST client

Use the MicroProfile REST client to register providers declaratively by adding the **org.eclipse.microprofile.rest.client.annotation.RegisterProvider** annotation to the target interface, as shown in the following example:

```
@Path("resource")
@RegisterProvider(MyClientResponseFilter.class)
@RegisterProvider(MyMessageBodyReader.class)
public interface TestResourceIntf2 {
  @Path("test/{path}")
  @Consumes("text/plain")
  @Produces("text/html")
  @POST
  public String test(@PathParam("path") String path, @QueryParam("query") String query, String
entity);
}
```

Declaring the **MyClientResponseFilter** class and the **MyMessageBodyReader** class with annotations eliminates the need to call the **RestClientBuilder.register()** method.

4.8.4. Declarative specification of headers in MicroProfile REST client

You can specify a header for an HTTP request in the following ways:

- By annotating one of the resource method parameters.
- By declaratively using the **org.eclipse.microprofile.rest.client.annotation.ClientHeaderParam** annotation.

The following example illustrates setting a header by annotating one of the resource method parameters with the annotation **@HeaderParam**:

```
@POST
@Produces(MediaType.TEXT_PLAIN)
```

```
@Consumes(MediaType.TEXT_PLAIN)
String contentLang(@HeaderParam(HttpHeaders.CONTENT_LANGUAGE) String contentLanguage,
String subject);
```

The following example illustrates setting a header using the **org.eclipse.microprofile.rest.client.annotation.ClientHeaderParam** annotation:

```
@POST
@Produces(MediaType.TEXT_PLAIN)
@Consumes(MediaType.TEXT_PLAIN)
@ClientHeaderParam(name=HttpHeaders.CONTENT_LANGUAGE, value="{getLanguage}")
String contentLang(String subject);

default String getLanguage() {
  return ...;
}
```

4.8.5. ResponseExceptionMapper in MicroProfile REST client

The **org.eclipse.microprofile.rest.client.ext.ResponseExceptionMapper** class is the client-side inverse of the **javax.ws.rs.ext.ExceptionMapper** class, which is defined in Jakarta RESTful Web Services. The **ExceptionMapper.toResponse()** method turns an **Exception** class thrown during the server-side processing into a **Response** class. The **ResponseExceptionMapper.toThrowable()** method turns a **Response** class received on the client-side with an HTTP error status into an **Exception** class.

You can register the **ResponseExceptionMapper** class either programmatically or declaratively. In the absence of a registered **ResponseExceptionMapper** class, a default **ResponseExceptionMapper** class maps any response with status **>= 400** to a **WebApplicationException** class.

4.8.6. Context dependency injection with MicroProfile REST client

With the MicroProfile REST client, you must annotate any interface that is managed as a Jakarta contexts and dependency injection (Jakarta Contexts and Dependency Injection) bean with the **@RegisterRestClient** class. For example:

```
@Path("resource")
@RegisterProvider(MyClientResponseFilter.class)
public static class TestResourceImpl {
    @Inject TestDataBase db;

    @Path("test/{path}")
    @Consumes("text/plain")
    @Produces("text/html")
    @POST
    public String test(@PathParam("path") String path, @QueryParam("query")
    String query, String entity) {
      return db.getByName(query);
    }
  }
  @Path("database")
  @RegisterRestClient
  public interface TestDataBase {
```

```
    @Path("")
    @POST
    public String getByName(String name);
}
```

Here, the MicroProfile REST client implementation creates a client for a **TestDataBase** class service, allowing easy access by the **TestResourceImpl** class. However, it does not include the information about the path to the **TestDataBase** class implementation. This information can be supplied by the optional **@RegisterProvider** parameter **baseUri**:

```
@Path("database")
@RegisterRestClient(baseUri="https://localhost:8080/webapp")
public interface TestDataBase {
  @Path("")
  @POST
  public String getByName(String name);
}
```

This indicates that you can access the implementation of **TestDataBase** at https://localhost:8080/webapp. You can also use MicroProfile configuration to supply the information externally:

```
<fully qualified name of TestDataBase>/mp-rest/url=<URL>
```

For example, the following property indicates that you can access an implementation of the **com.bluemonkeydiamond.TestDatabase** class at https://localhost:8080/webapp:

```
com.bluemonkeydiamond.TestDatabase/mp-rest/url=https://localhost:8080/webapp
```

You can supply a number of other properties to Jakarta Contexts and Dependency Injection clients. For example, **com.mycompany.remoteServices.MyServiceClient/mp-rest/providers**, comma-separated list of fully-qualified provider class names to include in the client.

Additional resources

- For more information about the MicroProfile REST Client specification, see Rest Client for MicroProfile.
- For more information about MicroProfile REST Client 2.0 features, see MicroProfile REST Client 2.0.

CHAPTER 5. BUILD AND RUN MICROSERVICES APPLICATIONS ON THE OPENSHIFT IMAGE FOR JBOSS EAP XP

You can build and run your microservices applications on the OpenShift image for JBoss EAP XP.

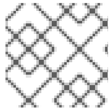

NOTE

JBoss EAP XP is supported only on OpenShift 4 and later versions.

Use the following workflow to build and run a microservices application on the OpenShift image for JBoss EAP XP by using the source-to-image (S2I) process.

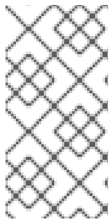

NOTE

The OpenShift images for JBoss EAP XP 4.0.0 provide a default standalone configuration file, which is based on the **standalone-microprofile-ha.xml** file. For more information about the server configuration files included in JBoss EAP XP, see the *Standalone server configuration files* section.

This workflow uses the **microprofile-config** quickstart as an example. The quickstart provides a small, specific working example that can be used as a reference for your own project. See the **microprofile-config** quickstart that ships with JBoss EAP XP 4.0.0 for more information.

Additional resources

- For more information about the server configuration files included in JBoss EAP XP, see Standalone server configuration files.

5.1. PREPARING OPENSHIFT FOR APPLICATION DEPLOYMENT

Prepare OpenShift for application deployment.

Prerequisites

You have installed an operational OpenShift instance. For more information, see the *Installing and Configuring OpenShift Container Platform Clusters* book on Red Hat Customer Portal.

Procedure

1. Log in to your OpenShift instance using the **oc login** command.
2. Create a new project in OpenShift.
   A project allows a group of users to organize and manage content separately from other groups. You can create a project in OpenShift using the following command.

   ```
   $ oc new-project PROJECT_NAME
   ```

   For example, for the **microprofile-config** quickstart, create a new project named **eap-demo** using the following command.

   ```
   $ oc new-project eap-demo
   ```

5.2. CONFIGURING AUTHENTICATION TO THE RED HAT CONTAINER REGISTRY

Before you can import and use the OpenShift image for JBoss EAP XP, you must configure authentication to the Red Hat Container Registry.

Create an authentication token using a registry service account to configure access to the Red Hat Container Registry. You need not use or store your Red Hat account's username and password in your OpenShift configuration when you use an authentication token.

Procedure

1. Follow the instructions on Red Hat Customer Portal to create an authentication token using a Registry Service Account management application.
2. Download the YAML file containing the OpenShift secret for the token.
   You can download the YAML file from the **OpenShift Secret** tab on your token's **Token Information** page.
3. Create the authentication token secret for your OpenShift project using the YAML file that you downloaded:

   ```
   oc create -f 1234567_myserviceaccount-secret.yaml
   ```

4. Configure the secret for your OpenShift project using the following commands, replacing the secret name below with the name of your secret created in the previous step.

   ```
   oc secrets link default 1234567-myserviceaccount-pull-secret --for=pull
   oc secrets link builder 1234567-myserviceaccount-pull-secret --for=pull
   ```

Additional resources

- Configuring authentication to the Red Hat Container Registry
- Registry Service Account management application
- Configuring access to secured registries

5.3. IMPORTING THE LATEST OPENSHIFT IMAGESTREAMS AND TEMPLATES FOR JBOSS EAP XP

Import the latest OpenShift imagestreams and templates for JBoss EAP XP.

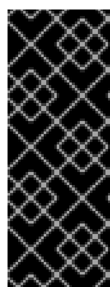

IMPORTANT

OpenJDK 8 images and imagestreams on OpenShift are deprecated.

The images and imagestreams are still supported on OpenShift. However, no enhancements are made to these images and imagestreams and they might be removed in the future. Red Hat continues to provide full support and bug fixes OpenJDK 8 images and imagestreams under its standard support terms and conditions.

Procedure

1. To import the latest imagestreams and templates for the OpenShift image for JBoss EAP XP into your OpenShift project's namespace, use the following commands:

   a. Import JDK 11 imagestream:

   ```
   oc replace --force -f https://raw.githubusercontent.com/jboss-container-images/jboss-eap-openshift-templates/eap-xp4/eap-xp4-openjdk11-image-stream.json
   ```

   This command imports the following imagestreams and templates:

   - The JDK 11 builder imagestream: jboss-eap-xp4-openjdk11-openshift
   - The JDK 11 runtime imagestream: jboss-eap-xp4-openjdk11-runtime-openshift

   b. Import the OpenShift templates:

   ```
   oc replace --force -f https://raw.githubusercontent.com/jboss-container-images/jboss-eap-openshift-templates/eap-xp4/templates/eap-xp4-basic-s2i.json
   ```

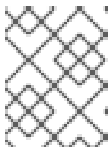

NOTE

The JBoss EAP XP imagestreams and templates imported using the above command are only available within that OpenShift project.

2. If you have administrative access to the general **openshift** namespace and want the imagestreams and templates to be accessible by all projects, add **-n openshift** to the **oc replace** line of the command. For example:

   ```
   ...
   oc replace -n openshift --force -f \
   ...
   ```

3. If you want to import the imagestreams and templates into a different project, add the **-n *PROJECT_NAME*** to the **oc replace** line of the command. For example:

   ```
   ...
   oc replace -n PROJECT_NAME --force -f
   ...
   ```

   If you use the cluster-samples-operator, see the OpenShift documentation on configuring the cluster samples operator. See Configuring the Cluster Samples Operator for details about configuring the cluster samples operator.

5.4. DEPLOYING A JBOSS EAP XP SOURCE-TO-IMAGE (S2I) APPLICATION ON OPENSHIFT

Deploy a JBoss EAP XP source-to-image (S2I) application on OpenShift.

Prerequisites

- Optional: A template can specify default values for many template parameters, and you might

have to override some, or all, of the defaults. To see template information, including a list of parameters and any default values, use the command **oc describe template TEMPLATE_NAME**.

Procedure

1. Create a new OpenShift application using the JBoss EAP XP image and your Java application's source code. Use one of the provided JBoss EAP XP templates for S2I builds.

```
$ oc new-app --template=eap-xp4-basic-s2i \ (1)
 -p EAP_IMAGE_NAME=jboss-eap-xp4-openjdk11-openshift:latest \
 -p EAP_RUNTIME_IMAGE_NAME=jboss-eap-xp4-openjdk11-runtime-openshift:latest \
 -p IMAGE_STREAM_NAMESPACE=eap-demo \ (2)
 -p SOURCE_REPOSITORY_URL=https://github.com/jboss-developer/jboss-eap-quickstarts
\ (3)
 -p SOURCE_REPOSITORY_REF=xp-4.0.x \ (4)
 -p CONTEXT_DIR=microprofile-config (5)
```

(1) The template to use. The application image is tagged with the **latest** tag.

(2) The latest images streams and templates were imported into the project's namespace, so you must specify the namespace of where to find the imagestream. This is usually the project's name.

(3) URL to the repository containing the application source code.

(4) The Git repository reference to use for the source code. This can be a Git branch or tag reference.

(5) The directory within the source repository to build.

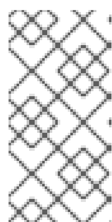

NOTE

A template can specify default values for many template parameters, and you might have to override some, or all, of the defaults. To see template information, including a list of parameters and any default values, use the command **oc describe template *TEMPLATE_NAME***.

You might also want to configure environment variables when creating your new OpenShift application.

2. Retrieve the name of the build configurations.

```
$ oc get bc -o name
```

3. Use the name of the build configurations from the previous step to view the Maven progress of the builds.

```
$ oc logs -f buildconfig/${APPLICATION_NAME}-build-artifacts

…
Push successful
```

```
$ oc logs -f buildconfig/${APPLICATION_NAME}
...
Push successful
```

For example, for the **microprofile-config**, the following command shows the progress of the Maven builds.

```
$ oc logs -f buildconfig/eap-xp4-basic-app-build-artifacts

...
Push successful
$ oc logs -f buildconfig/eap-xp4-basic-app
...
Push successful
```

Additional resources

- Importing the latest OpenShift imagestreams and templates for JBoss EAP XP .
- Preparing OpenShift for application deployment.

5.5. COMPLETING POST-DEPLOYMENT TASKS FOR JBOSS EAP XP SOURCE-TO-IMAGE (S2I) APPLICATION

Depending on your application, you might need to complete some tasks after your OpenShift application has been built and deployed.

Examples of post-deployment tasks include the following:

- Exposing a service so that the application is viewable from outside of OpenShift.
- Scaling your application to a specific number of replicas.

Procedure

1. Get the service name of your application using the following command.

```
$ oc get service
```

2. **Optional**: Expose the main service as a route so you can access your application from outside of OpenShift. For example, for the **microprofile-config** quickstart, use the following command to expose the required service and port.

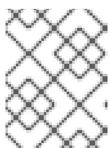

NOTE

If you used a template to create the application, the route might already exist. If it does, continue on to the next step.

```
$ oc expose service/eap-xp4-basic-app --port=8080
```

3. Get the URL of the route.

```
$ oc get route
```

4. Access the application in your web browser using the URL. The URL is the value of the **HOST/PORT** field from previous command's output.

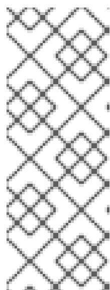

NOTE

For JBoss EAP XP 4.0.0 GA distribution, the Microprofile Config quickstart does not reply to HTTPS GET requests to the application's root context. This enhancement is only available in the {JBossXPShortName101} GA distribution.

For example, to interact with the Microprofile Config application, the URL might be **http://*HOST_PORT_Value*/config/value** in your browser.

If your application does not use the JBoss EAP root context, append the context of the application to the URL. For example, for the **microprofile-config** quickstart, the URL might be **http://*HOST_PORT_VALUE*/microprofile-config/**.

5. Optionally, you can scale up the application instance by running the following command. This command increases the number of replicas to 3.

```
$ oc scale deploymentconfig DEPLOYMENTCONFIG_NAME --replicas=3
```

For example, for the **microprofile-config** quickstart, use the following command to scale up the application.

```
$ oc scale deploymentconfig/eap-xp4-basic-app --replicas=3
```

Additional Resources

For more information about JBoss EAP XP Quickstarts, see the JBoss EAP XP quickstart.

CHAPTER 6. CAPABILITY TRIMMING

When building a bootable JAR, you can decide which JBoss EAP features and subsystems to include.

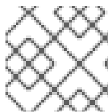

NOTE

Capability trimming is supported only on OpenShift or when building a bootable JAR.

Additional resources

- About the bootable JAR

6.1. AVAILABLE JBOSS EAP LAYERS

Red Hat makes available a number of layers to customize provisioning of the JBoss EAP server in OpenShift or a bootable JAR.

Three layers are base layers that provide core functionality. The other layers are decorator layers that enhance the base layers with additional capabilities.

Most decorator layers can be used to build S2I images in JBoss EAP for OpenShift or to build a bootable JAR. A few layers do not support S2I images; the description of the layer notes this limitation.

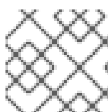

NOTE

Only the listed layers are supported. Layers not listed here are not supported.

6.1.1. Base layers

Each base layer includes core functionality for a typical server user case.

datasources-web-server
This layer includes a servlet container and the ability to configure a datasource.

This layer does not include MicroProfile capabilities.

The following Jakarta EE specifications are supported in this layer:

- Jakarta JSON Processing 1.1
- Jakarta JSON Binding 1.0
- Jakarta Servlet 4.0
- Jakarta Expression Language 3.0
- Jakarta Server Pages 2.3
- Jakarta Standard Tag Library 1.2
- Jakarta Concurrency 1.1
- Jakarta Annotations 1.3
- Jakarta XML Binding 2.3

- Jakarta Debugging Support for Other Languages 1.0
- Jakarta Transaction 1.3
- Jakarta Connector API 1.7

jaxrs-server
This layer enhances the **datasources-web-server** layer with the following JBoss EAP subsystems:

- **jaxrs**
- **weld**
- **jpa**

This layer also adds Infinispan-based second-level entity caching locally in the container.

The following MicroProfile capability is included in this layer:

- MicroProfile REST Client

The following Jakarta EE specifications are supported in this layer in addition to those supported in the **datasources-web-server** layer:

- Jakarta Contexts and Dependency Injection 2.0
- Jakarta Bean Validation 2.0
- Jakarta Interceptors 1.2
- Jakarta RESTful Web Services 2.1
- Jakarta Persistence 2.2

cloud-server
This layer enhances the **jaxrs-server** layer with the following JBoss EAP subsystems:

- **resource-adapters**
- **messaging-activemq** (remote broker messaging, not embedded messaging)

This layer also adds the following observability features to the **jaxrs-server** layer:

- MicroProfile Health
- MicroProfile Metrics
- MicroProfile Config
- MicroProfile OpenTracing

The following Jakarta EE specification is supported in this layer in addition to those supported in the **jaxrs-server** layer:

- Jakarta Security 1.0

6.1.2. Decorator layers

Decorator layers are not used alone. You can configure one or more decorator layers with a base layer to deliver additional functionality.

ejb-lite
This decorator layer adds a minimal Jakarta Enterprise Beans implementation to the provisioned server. The following support is not included in this layer:

- IIOP integration
- MDB instance pool
- Remote connector resource

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

Jakarta Enterprise Beans
This decorator layer extends the **ejb-lite** layer. This layer adds the following support to the provisioned server, in addition to the base functionality included in the **ejb-lite** layer:

- MDB instance pool
- Remote connector resource

Use this layer if you want to use message-driven beans (MDBs) or Jakarta Enterprise Beans remoting capabilities, or both. If you do not need these capabilities, use the **ejb-lite** layer.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

ejb-local-cache
This decorator layer adds local caching support for Jakarta Enterprise Beans to the provisioned server.

Dependencies: You can only include this layer if you have included the **ejb-lite** layer or the **ejb** layer.

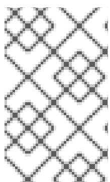

NOTE

This layer is not compatible with the **ejb-dist-cache** layer. If you include the **ejb-dist-cache** layer, you cannot include the **ejb-local-cache** layer. If you include both layers, the resulting build might include an unexpected Jakarta Enterprise Beans configuration.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

ejb-dist-cache
This decorator layer adds distributed caching support for Jakarta Enterprise Beans to the provisioned server.

Dependencies: You can only include this layer if you have included the **ejb-lite** layer or the **ejb** layer.

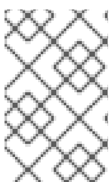

NOTE

This layer is not compatible with the **ejb-local-cache** layer. If you include the **ejb-dist-cache** layer, you cannot include the **ejb-local-cache** layer. If you include both layers, the resulting build might result in an unexpected configuration.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

jdr

This decorator layer adds the JBoss Diagnostic Reporting (**jdr**) subsystem to gather diagnostic data when requesting support from Red Hat.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

Jakarta Persistence
This decorator layer adds persistence capabilities for a single-node server. Note that distributed caching only works if the servers are able to form a cluster.

The layer adds Hibernate libraries to the provisioned server, with the following support:

- Configurations of the **jpa** subsystem
- Configurations of the **infinispan** subsystem
- A local Hibernate cache container

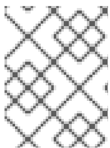

NOTE

This layer is not compatible with the **jpa-distributed** layer. If you include the **jpa** layer, you cannot include the **jpa-distributed** layer.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

jpa-distributed
This decorator layer adds persistence capabilities for servers operating in a cluster. The layer adds Hibernate libraries to the provisioned server, with the following support:

- Configurations of the **jpa** subsystem
- Configurations of the **infinispan** subsystem
- A local Hibernate cache container
- Invalidation and replication Hibernate cache containers
- Configuration of the **jgroups** subsystem

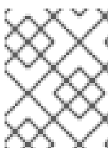

NOTE

This layer is not compatible with the **jpa** layer. If you include the **jpa** layer, you cannot include the **jpa-distributed** layer.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

Jakarta Server Faces
This decorator layer adds the **jsf** subsystem to the provisioned server.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

microprofile-platform
This decorator layer adds the following MicroProfile capabilities to the provisioned server:

- MicroProfile Config
- MicroProfile Fault Tolerance

- MicroProfile Health
- MicroProfile JWT
- MicroProfile Metrics
- MicroProfile OpenAPI
- MicroProfile OpenTracing

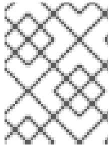

NOTE

This layer includes MicroProfile capabilities that are also included in the **observability** layer. If you include this layer, you do not need to include the **observability** layer.

observability
This decorator layer adds the following observability features to the provisioned server:

- MicroProfile Health
- MicroProfile Metrics
- MicroProfile Config
- MicroProfile OpenTracing

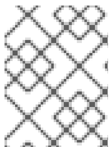

NOTE

This layer is built in to the **cloud-server** layer. You do not need to add this layer to the **cloud-server** layer.

remote-activemq
This decorator layer adds the ability to communicate with a remote ActiveMQ broker to the provisioned server, integrating messaging support.

The pooled connection factory configuration specifies **guest** as the value for the **user** and **password** attributes. You can use a CLI script to change these values at runtime.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

sso
This decorator layer adds Red Hat Single Sign-On integration to the provisioned server.

This layer should only be used when provisioning a server using S2I.

web-console
This decorator layer adds the management console to the provisioned server.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

web-clustering
This decorator layer adds support for distributable web applications by configuring a non-local Infinispan-based container web cache for data session handling suitable to clustering environments.

web-passivation

This decorator layer adds support for distributable web applications by configuring a local Infinispan-based container web cache for data session handling suitable to single node environments.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

webservices
This layer adds web services functionality to the provisioned server, supporting Jakarta web services deployments.

This layer is only supported when building a bootable JAR. This layer is not supported when using S2I.

Additional resources

- Pooled Connection Factory Attributes

CHAPTER 7. ENABLE MICROPROFILE APPLICATION DEVELOPMENT FOR JBOSS EAP ON RED HAT CODEREADY STUDIO

If you want to incorporate MicroProfile capabilities in applications that you develop on CodeReady Studio, you must enable MicroProfile support for JBoss EAP in CodeReady Studio.

JBoss EAP expansion packs provide support for MicroProfile.

JBoss EAP expansion packs are not supported on JBoss EAP 7.2 and earlier.

Each version of the JBoss EAP expansion pack supports specific patches of JBoss EAP. For details, see the JBoss EAP expansion pack Support and Life Cycle Policies page.

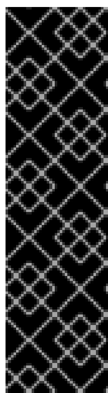

IMPORTANT

The JBoss EAP XP Quickstarts for Openshift are provided as Technology Preview only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

See Technology Preview Features Support Scope on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

7.1. CONFIGURING CODEREADY STUDIO TO USE MICROPROFILE CAPABILITIES

To enable MicroProfile support on JBoss EAP, register a new runtime server for JBoss EAP XP, and then create the new JBoss EAP 7.4 server.

Give the server an appropriate name that helps you recognize that it supports MicroProfile capabilities.

This server uses a newly created JBoss EAP XP runtime that points to the runtime installed previously and uses the **standalone-microprofile.xml** configuration file.

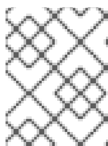

NOTE

If you set the **Target runtime** to **7.4** or a later runtime version in Red Hat CodeReady Studio, your project is compatible with the Jakarta EE 8 specification.

Prerequisites

- JBoss EAP XP 4.0.0 has been installed .

Procedure

1. Set up the new server on the **New Server** dialog box.

    a. In the **Select server type** list, select *Red Hat JBoss Enterprise Application Platform 7.4* .

    b. In the **Server's host name** field, enter *localhost*.

c. In the **Server name** field, enter *JBoss EAP 7.4 XP*.

d. Click **Next**.

2. Configure the new server.

   a. In the **Home directory** field, if you do not want to use the default setting, specify a new directory; for example: *home/myname/dev/microprofile/runtimes/jboss-eap-7.4*.

   b. Make sure the **Execution Environment** is set to *JavaSE-1.8*.

   c. Optional: Change the values in the **Server base directory** and **Configuration file** fields.

   d. Click **Finish**.

Result

You are now ready to begin developing applications using MicroProfile capabilities, or to begin using the MicroProfile quickstarts for JBoss EAP.

7.2. USING MICROPROFILE QUICKSTARTS FOR CODEREADY STUDIO

Enabling the MicroProfile quickstarts makes the simple examples available to run and test on your installed server.

These examples illustrate the following MicroProfile capabilities.

- MicroProfile Config
- MicroProfile Fault Tolerance
- MicroProfile Health
- MicroProfile JWT
- MicroProfile Metrics
- MicroProfile OpenAPI
- MicroProfile OpenTracing
- MicroProfile REST Client

Procedure

1. Import the **pom.xml** file from the Quickstart Parent Artifact.

2. If the quickstart you are using requires environment variables, configure the environment variables.
   Define environment variables on the launch configuration on the server **Overview** dialog box.

   For example, the **microprofile-opentracing** quickstart uses the following environment variables:

   - **JAEGER_REPORTER_LOG_SPANS** set to **true**
   - **JAEGER_SAMPLER_PARAM** set to **1**

- **JAEGER_SAMPLER_TYPE** set to **const**

Additional resources

About Microprofile

About JBoss Enterprise Application Platform expansion pack

Red Hat JBoss Enterprise Application Platform expansion pack Support and Life Cycle Policies

CHAPTER 8. THE BOOTABLE JAR

You can build and package a microservices application as a bootable JAR with the JBoss EAP JAR Maven plug-in. You can then run the application on a JBoss EAP bare-metal platform or a JBoss EAP OpenShift platform.

8.1. ABOUT THE BOOTABLE JAR

You can build and package a microservices application as a bootable JAR with the JBoss EAP JAR Maven plug-in.

A bootable JAR contains a server, a packaged application, and the runtime required to launch the server.

The JBoss EAP JAR Maven plug-in uses Galleon trimming capability to reduce the size and memory footprint of the server. Thus, you can configure the server according to your requirements, including only the Galleon layers that provide the capabilities that you need.

The JBoss EAP JAR Maven plug-in supports the execution of JBoss EAP CLI script files to customize your server configuration. A CLI script includes a list of CLI commands for configuring the server.

A bootable JAR is like a standard JBoss EAP server in the following ways:

- It supports JBoss EAP common management CLI commands.
- It can be managed using the JBoss EAP management console.

The following limitations exist when packaging a server in a bootable JAR:

- CLI management operations that require a server restart are not supported.
- The server cannot be restarted in admin-only mode, which is a mode that starts services related to server administration.
- If you shut down the server, updates that you applied to the server are lost.

Additionally, you can provision a hollow bootable JAR. This JAR contains only the server, so you can reuse the server to run a different application.

Additional resources

For information about capability trimming, see Capability Trimming.

8.2. JBOSS EAP MAVEN PLUG-IN

You can use the JBoss EAP JAR Maven plug-in to build an application as a bootable JAR.

You can retrieve the latest Maven plug-in version from the Maven repository, which is available at Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.

In a Maven project, the **src** directory contains all the source files required to build your application. After the JBoss EAP JAR Maven plug-in builds the bootable JAR, the generated JAR is located in **target/<application>-bootable.jar**.

The JBoss EAP JAR Maven plug-in also provides the following functionality:

- Applies CLI script commands to the server.
- Uses the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack and some of its layers for customizing the server configuration file.
- Supports the addition of extra files into the packaged bootable JAR, such as a keystore file.
- Includes the capability to create a hollow bootable JAR; that is, a bootable JAR that does not contain an application.

After you use the JBoss EAP JAR Maven plug-in to create the bootable JAR, you can start the application by issuing the following command. Replace **target/myapp-bootable.jar** with the path to your bootable JAR. For example:

```
$ java -jar target/myapp-bootable.jar
```

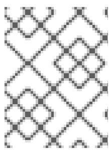

NOTE

To get a list of supported bootable JAR startup commands, append **--help** to the end of the startup command. For example, **java -jar target/myapp-bootable.jar --help**.

Additional resources

- For information about supported JBoss EAP Galleon layers, see Available JBoss EAP layers.
- For information about supported Galleon plug-ins to build feature packs for your project, see the WildFly Galleon Maven Plugin Documentation .
- For information about selecting methods to configure the JBoss EAP Maven repository, see Use the Maven Repository.
- For information about Maven project directories, see Introduction to the Standard Directory Layout in the *Apache Maven* documentation.

8.3. BOOTABLE JAR ARGUMENTS

View the arguments in the following table to learn about supported arguments for use with the bootable JAR.

Table 8.1. Supported bootable JAR executable arguments

Argument	Description
--help	Displays the help message for the specified command and exit.

Argument	Description
--cli-script=<path>	Specifies the path to a JBoss CLI script that executes when starting the bootable JAR. If the path specified is relative, the path is resolved against the working directory of the Java VM instance used to launch the bootable JAR.
--deployment=<path>	Argument specific to the hollow bootable JAR. Specifies the path to the WAR, JAR, EAR file or exploded directory that contains the application you want to deploy on a server.
--display-galleon-config	Print the content of the generated Galleon configuration file.
--install-dir=<path>	By default, the JVM settings are used to create a *TEMP* directory after the bootable JAR is started. You can use the **--install-dir** argument to specify a directory to install the server.
-secmgr	Runs the server with a security manager installed.
-b<interface>=<value>	Set system property **jboss.bind.address.<interface>** to the given value. For example, **bmanagement=IP_ADDRESS**.
-b=<value>	Set system property **jboss.bind.address**, which is used in configuring the bind address for the public interface. This defaults to 127.0.0.1 if no value is specified.
-D<name>[=<value>]	Specifies system properties that are set by the server at server runtime. The bootable JAR JVM does not set these system properties.

Argument	Description
--properties=<url>	Loads system properties from a specified URL.
-S<name>[=value]	Set a security property.
-u=<value>	Set system property **jboss.default.multicast.address**, which is used in configuring the multicast address in the socket-binding elements in the configuration files. This defaults to 230.0.0.4 if no value is specified.
--version	Display the application server version and exit.

8.4. SPECIFYING GALLEON LAYERS FOR YOUR BOOTABLE JAR SERVER

You can specify Galleon layers to build a custom configuration for your server. Additionally, you can specify Galleon layers that you want excluded from the server.

To reference a single feature pack, use the **<feature-pack-location>** element to specify its location. The following example specifies **org.jboss.eap:wildfly-galleon-pack:4.0.0.GA-redhat-00002** in the **<feature-pack-location>** element of the Maven plug-in configuration file.

```
<configuration>
  <feature-pack-location>org.jboss.eap:wildfly-galleon-pack:4.0.0.GA-redhat-00002</feature-pack-location>
</configuration>
```

If you need to reference more than one feature pack, list them in the **<feature-packs>** element. The following example shows the addition of the Red Hat Single Sign-On feature pack to the **<feature-packs>** element:

```
<configuration>
    <feature-packs>
        <feature-pack>
            <location>org.jboss.eap:wildfly-galleon-pack:4.0.0.GA-redhat-00002</location>
        </feature-pack>
        <feature-pack>
            <location>org.keycloak:keycloak-adapter-galleon-pack:15.0.4.redhat-00001</location>
        </feature-pack>
    </feature-packs>
</configuration>
```

You can combine Galleon layers from multiple feature packs to configure the bootable JAR server to include only the supported Galleon layers that provide the capabilities that you need.

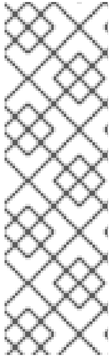

NOTE

On a bare-metal platform, if you do not specify Galleon layers in your configuration file then the provisioned server contains a configuration identical to that of a default **standalone-microprofile.xml** configuration.

On an OpenShift platform, after you have added the **<cloud/>** configuration element in the plug-in configuration and you choose not to specify Galleon layers in your configuration file, the provisioned server contains a configuration that is adjusted for the cloud environment and is similar to a default **standalone-microprofile-ha.xml**.

Prerequisites

- Maven is installed.
- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.

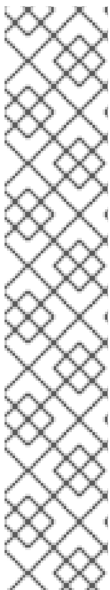

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Identify the supported JBoss EAP Galleon layers that provide the capabilities that you need to run your application.
2. Reference a JBoss EAP feature pack location in the **<plugin>** element of the Maven project **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack, as demonstrated in the following

example. The following example also displays the inclusion of a single feature-pack, which includes the **jaxrs-server** base layer and the **jpa-distributed** layer . The **jaxrs-server** base layer provides additional support for the server.

```
<plugins>
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
         <artifactId>wildfly-jar-maven-plugin</artifactId>
         <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
            <feature-pack-location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
            <layers>
                <layer>jaxrs-server</layer>
                <layer>jpa-distributed</layer>
            </layers>
            <excluded-layers>
                <layer>jpa</layer>
            </excluded-layers>
            ...
</plugins>
```

This example also shows the exclusion of the **jpa** layer from the project.

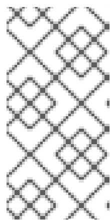

NOTE

If you include the **jpa-distributed** layer in your project, you must exclude the **jpa** layer from the **jaxrs-server** layer. The **jpa** layer configures a local infinispan hibernate cache, while the **jpa-distributed** layer configures a remote infinispan hibernate cache.

Additional resources

- For information about available base layers, see Base layers.
- For information about supported Galleon plug-ins to build feature packs for your project, see the WildFly Galleon Maven Plugin Documentation .
- For information about selecting methods to configure the JBoss EAP Maven repository, see Maven and the JBoss EAP MicroProfile Maven repository .
- For information about managing your Maven dependencies, see Dependency Management in the *Apache Maven Project* documentation.

8.5. USING A BOOTABLE JAR ON A JBOSS EAP BARE-METAL PLATFORM

You can package an application as a bootable JAR on a JBoss EAP bare-metal platform.

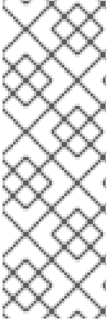

NOTE

- To use the custom Galleon feature-pack and layers when building a bootable JAR on a JBoss EAP bare-metal platform, see Building and using custom Galleon layers for JBoss EAP.
- When building the application image by using the **oc new-build** command, ensure to use this S2I builder image **jboss-eap-xp4-openjdk11-openshift:latest**, instead of **jboss-eap74-openjdk11-openshift:latest**.

A bootable JAR contains a server, a packaged application, and the runtime required to launch the server.

This procedure demonstrates packaging the MicroProfile Config microservices application as a bootable JAR with the JBoss EAP JAR Maven plug-in. See MicroProfile Config development.

You can use CLI scripts to configure the server during the packaging of the bootable JAR.

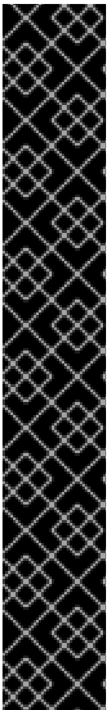

IMPORTANT

On building a web application that must be packaged inside a bootable JAR, you must specify **war** in the **<packaging>** element of your **pom.xml** file. For example:

```
<packaging>war</packaging>
```

This value is required to package the build application as a WAR file and not as the default JAR file.

In a Maven project that is used solely to build a hollow bootable JAR, set the packaging value to **pom**. For example:

```
<packaging>pom</packaging>
```

You are not limited to using **pom** packaging when you build a hollow bootable JAR for a Maven project. You can create one by specifying **true** in the **<hollow-jar>** element for any type of packaging, such as **war**. See Creating a hollow bootable JAR on a JBoss EAP bare-metal platform.

Prerequisites

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.
- You have created a Maven project, set up a parent dependency, and added dependencies for creating an MicroProfile application. See MicroProfile Config development.

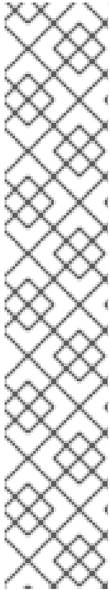

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
        <artifactId>wildfly-jar-maven-plugin</artifactId>
        <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
             <feature-pack-location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
            <layers>
                <layer>jaxrs-server</layer>
                <layer>microprofile-platform</layer>
            </layers>
         </configuration>
        <executions>
            <execution>
                <goals>
                    <goal>package</goal>
                </goals>
            </execution>
        </executions>
    </plugin>
</plugins>
```

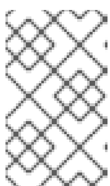

NOTE

If you do not specify Galleon layers in your **pom.xml** file then the bootable JAR server contains a configuration that is identical to a **standalone-microprofile.xml** configuration.

2. Package the application as a bootable JAR:

```
$ mvn package
```

3. Start the application:

```
$ NAME="foo" java -jar target/microprofile-config-bootable.jar
```

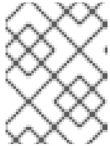

NOTE

The example uses **NAME** as the environment variable, but you can choose to use **jim**, which is the default value.

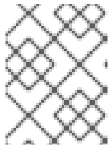

NOTE

To view a list of supported bootable JAR arguments, append **--help** to the end of the **java -jar target/microprofile-config-bootable.jar** command.

4. Specify the following URL in your web browser to access the MicroProfile Config application:

```
http://localhost:8080/config/json
```

5. *Verification:* Test the application behaves properly by issuing the following command in your terminal:

```
curl http://localhost:8080/config/json
```

The following is the expected output:

```
{"result":"Hello foo"}
```

Additional resources

- For information about available MicroProfile Config functionality, see MicroProfile Config.
- For information about **ConfigSources**, see MicroProfile Config reference.

8.6. CREATING A HOLLOW BOOTABLE JAR ON A JBOSS EAP BARE-METAL PLATFORM

You can package an application as a hollow bootable JAR on a JBoss EAP bare-metal platform.

A hollow bootable JAR contains only the JBoss EAP server. The hollow bootable JAR is packaged by the JBoss EAP JAR Maven plug-in. The application is provided at server runtime. The hollow bootable JAR is useful if you need to re-use the server configuration for a different application.

Prerequisites

- You have created a Maven project, set up a parent dependency, and added dependencies for creating an application. See MicroProfile Config development.
- You have completed the **pom.xml** file configuration steps outlined in Using a bootable JAR on a JBoss EAP bare-metal platform.

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.

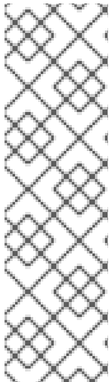

NOTE

The example shown in the procedure specifies **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version, but you must set the property in your project. For example:

```
<properties>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. To build a hollow bootable JAR, you must set the **<hollow-jar>** plug-in configuration element to true in the project **pom.xml** file. For example:

```
<plugins>
       <plugin>
           ...
           <configuration>
               <!-- This example configuration does not show a complete plug-in configuration -->
               ...
               <feature-pack-location>org.jboss.eap:wildfly-galleon-pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
               <hollow-jar>true</hollow-jar>
           </configuration>
       </plugin>
</plugins>
```

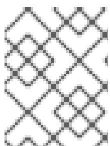

NOTE

By specifying **true** in the **<hollow-jar>** element, the JBoss EAP JAR Maven plug-in does not include an application in the JAR.

1. Build the hollow bootable JAR:

```
$ mvn clean package
```

2. Run the hollow bootable JAR:

```
$ java -jar target/microprofile-config-bootable.jar --deployment=target/microprofile-config.war
```

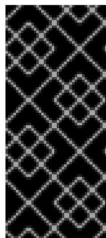

IMPORTANT

To specify the path to the WAR file that you want to deploy on the server, use the following argument, where **<PATH_NAME>** is the path to your deployment.

```
--deployment=<PATH_NAME>
```

3. Access the application:

```
$ curl http://localhost:8080/microprofile-config/config/json
```

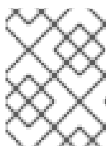

NOTE

To register your web application in the root directory, name the application **ROOT.war**.

Additional resources

- For information about available MicroProfile functionality, see MicroProfile Config.
- For more information about the JBoss EAP JAR Maven plug-in supported in JBoss EAP XP 4.0.0, see JBoss EAP Maven plug-in.

8.7. CLI SCRIPTS EXECUTED AT BUILD TIME

You can create CLI scripts to configure the server during the packaging of the bootable JAR.

A CLI script is a text file that contains a sequence of CLI commands that you can use to apply additional server configurations. For example, you can create a script to add a new logger to the **logging** subsystem.

You can also specify more complex operations in a CLI script. For example, you can group security management operations into a single command to enable HTTP authentication for the management HTTP endpoint.

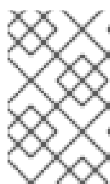

NOTE

You must define CLI scripts in the **<cli-session>** element of the plug-in configuration before you package an application as a bootable JAR. This ensures the server configuration settings persist after packaging the bootable JAR.

Although you can combine predefined Galleon layers to configure a server that deploys your application, limitations do exist. For example, you cannot enable the HTTPS **undertow** listener using Galleon layers when packaging the bootable JAR. Instead, you must use a CLI script.

You must define the CLI scripts in the **<cli-session>** element of the **pom.xml** file. The following table shows types of CLI session attributes:

Table 8.2. CLI script attributes

Argument	Description
script-files	List of paths to script files.
properties-file	Optional attribute that specifies a path to a properties file. This file lists Java properties that scripts can reference by using the **${my.prop}** syntax. The following example sets **public inet-address** to the value of **all.addresses**: **/interface=public:write-attribute(name=inet-address,value=${all.addresses})**
resolve-expressions	Optional attribute that contains a boolean value. Indicates if system properties or expressions are resolved before sending the operation requests to the server. Value is **true** by default.

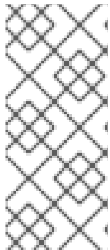

NOTE

- CLI scripts are started in the order that they are defined in the **<cli-session>** element of the **pom.xml** file.
- The JBoss EAP JAR Maven plug-in starts the embedded server for each CLI session. Thus, your CLI script does not have to start or stop the embedded server.

8.8. EXECUTING CLI SCRIPT AT RUNTIME

You can apply changes to the server configuration during runtime; this gives you the flexibility to adjust the server with respect to the execution context. However, the preferred way to apply changes to the server is during build time.

Procedure

- Launch the bootable JAR, and the **--cli-script** argument.
  For Example:

```
java -jar myapp-bootable.jar --cli-scipt=my-scli-scipt.cli
```

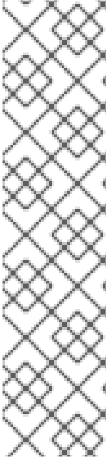

NOTE

- The CLI script must be a text file (UTF-8), the file extension if present is meaningless although **.cli** extension is advised.
- Operations that require your server to restart will terminate your bootable JAR instance.
- CLI commands such as **connect**, **reload**, **shutdown**, and any command related to embedded server and **patch** are not operational.
- CLI commands such as **jdbc-driver-info** that cannot be executed in admin-mode are not supported.

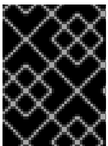

IMPORTANT

If you restart the server without executing the CLI script, your new server instance will not contain the changes from your previous server instance.

8.9. USING A BOOTABLE JAR ON A JBOSS EAP OPENSHIFT PLATFORM

After you packaged an application as a bootable JAR, you can run the application on a JBoss EAP OpenShift platform.

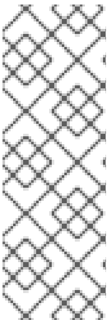

NOTE

- To use the custom Galleon feature-pack and layers when building a bootable JAR on a JBoss EAP OpenShift platform, see Building and using custom Galleon layers for JBoss EAP.
- When building the application image by using the **oc new-build** command, ensure to use this S2I builder image **jboss-eap-xp4-openjdk11-openshift:latest**, instead of **jboss-eap74-openjdk11-openshift:latest**.

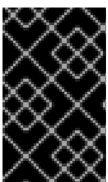

IMPORTANT

On OpenShift, you cannot use the EAP Operator automated transaction recovery feature with your bootable JAR. A fix for this technical limitation is planned for a future JBoss EAP XP 4.0.0 patch release.

Prerequisites

- You have created a Maven project for MicroProfile Config development.
- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP 4 and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during

the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.

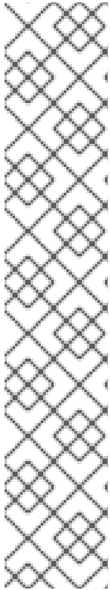

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
   <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-00001</bootable.jar.maven.plugin.version>
   <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
   <plugin>
      <groupId>org.wildfly.plugins</groupId>
      <artifactId>wildfly-jar-maven-plugin</artifactId>
      <version>${bootable.jar.maven.plugin.version}</version>
      <configuration>
         <feature-pack-location>org.jboss.eap:wildfly-galleon-pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
         <layers>
            <layer>jaxrs-server</layer>
            <layer>microprofile-platform</layer>
         </layers>
         <cloud/>
         </configuration>
         <executions>
            <execution>
               <goals>
                  <goal>package</goal>
               </goals>
            </execution>
         </executions>
   </plugin>
</plugins>
```

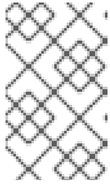

NOTE

You must include the **<cloud/>** element in the **<configuration>** element of the plug-in configuration, so the JBoss EAP Maven JAR plug-in can identify that you choose the OpenShift platform.

2. Package the application:

```
$ mvn package
```

3. Log in to your OpenShift instance using the **oc login** command.

4. Create a new project in OpenShift. For example:

```
$ oc new-project bootable-jar-project
```

5. Enter the following **oc** commands to create an application image:

```
$ mkdir target/openshift && cp target/microprofile-config-bootable.jar target/openshift (1)

$ oc import-image ubi8/openjdk-11 --from=registry.redhat.io/ubi8/openjdk-11 --confirm (2)

$ oc new-build --strategy source --binary --image-stream openjdk-11 --name microprofile-config-app (3)

$ oc start-build microprofile-config-app --from-dir target/openshift (4)
```

(1) Creates an openshift sub-directory in the target directory. The packaged application is copied into the created sub-directory.

(2) Imports the latest OpenJDK 11 imagestream tag and image information into the OpenShift project.

(3) Creates a build configuration based on the microprofile-config-app directory and the OpenJDK 11 imagestream.

(4) Uses the **target/openshift** sub-directory as the binary input to build the application.

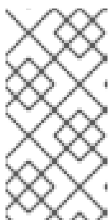

NOTE

OpenShift applies a set of CLI script commands to the bootable JAR configuration file to adjust it to the cloud environment. You can access this script by opening the **bootable-jar-build-artifacts/generated-cli-script.txt** file in the Maven project /**target directory**.

6. *Verification:*
   View a list of OpenShift pods available and check the pods build statuses by issuing the following command:

```
$ oc get pods
```

Verify the built application image:

```
$ oc get is microprofile-config-app
```

The output shows the built application image details, such as name and image repository, tag, and so on. For the example in this procedure, the imagestream name and tag output displays **microprofile-config-app:latest**.

7. Deploy the application:

```
$ oc new-app microprofile-config-app

$ oc expose svc/microprofile-config-app
```

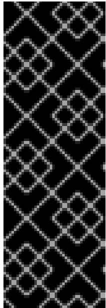

IMPORTANT

To provide system properties to the bootable JAR, you must use the **JAVA_OPTS_APPEND** environment variable. The following example demonstrates usage of the **JAVA_OPTS_APPEND** environment variable:

```
$ oc new-app <_IMAGESTREAM_> -e JAVA_OPTS_APPEND="-
Xlog:gc*:file=/tmp/gc.log:time -Dwildfly.statistics-enabled=true"
```

A new application is created and started. The application configuration is exposed as a new service.

8. *Verification*: Test the application behaves properly by issuing the following command in your terminal:

```
$ curl http://$(oc get route microprofile-config-app --template='{{ .spec.host }}')/config/json
```

Expected output:

```
{"result":"Hello jim"}
```

Additional resources

- For information about MicroProfile, see MicroProfile Config.
- For information about **ConfigSources**, see Default MicroProfile Config attributes.

8.10. CONFIGURE THE BOOTABLE JAR FOR OPENSHIFT

Before using your bootable JAR, you can configure JVM settings to ensure that your standalone server operates correctly on JBoss EAP for OpenShift.

Use the **JAVA_OPTS_APPEND** environment variable to configure JVM settings. Use the **JAVA_ARGS** command to provide arguments to the bootable JAR.

You can use environment variables to set values for properties. For example, you can use the **JAVA_OPTS_APPEND** environment variable to set the **-Dwildfly.statistics-enabled** property to **true**:

```
JAVA_OPTS_APPEND="-Xlog:gc*:file=/tmp/gc.log:time -Dwildfly.statistics-enabled=true"
```

Statistics are now enabled for your server.

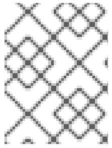

NOTE

Use the **JAVA_ARGS** environment variable, if you need to provide arguments to the bootable JAR.

JBoss EAP for OpenShift provides a JDK 11 image. To run the application associated with your bootable JAR, you must first import the latest OpenJDK 11 imagestream tag and image information into your OpenShift project. You can then use environment variables to configure the JVM in the imported image.

You can apply the same configuration options for configuring the JVM used for JBoss EAP for OpenShift S2I image, but with the following differences:

- Optional: The **-Xlog** capability is not available, but you can set garbage collection logging by enabling **-Xlog:gc**. For example: **JAVA_OPTS_APPEND="-Xlog:gc*:file=/tmp/gc.log:time"**.
- To increase initial metaspace size, you can set the **GC_METASPACE_SIZE** environment variable. For best metadata capacity performance, set the value to **96**.
- For better random file generation, use the **JAVA_OPTS_APPEND** environment variable to set **java.security.egd** property as **-Djava.security.egd=file:/dev/urandom**.

These configurations improve the memory settings and garbage collection capability of JVM when running on your imported OpenJDK 11 image.

8.11. USING A CONFIGMAP IN YOUR APPLICATION ON OPENSHIFT

For OpenShift, you can use a deployment controller (dc) to mount the configmap into the pods used to run the application.

A **ConfigMap** is an OpenShift resource that is used to store non-confidential data in key-value pairs.

After you specify the **microprofile-platform** Galleon layer to add **microprofile-config-smallrye** subsystem and any extensions to the server configuration file, you can use a CLI script to add a new **ConfigSource** to the server configuration. You can save CLI scripts in an accessible directory, such as the **/scripts** directory, in the root directory of your Maven project.

MicroProfile Config functionality is implemented in JBoss EAP using the SmallRye Config component and is provided by the **microprofile-config-smallrye** subsystem. This subsystem is included in the **microprofile-platform** Galleon layer.

Prerequisites

- You have installed Maven.
- You have configured the JBoss EAP Maven repository.
- You have packaged an application as a bootable JAR and you can run the application on a JBoss EAP OpenShift platform. For information about building an application as a bootable JAR on an OpenShift platform, see Using a bootable JAR on a JBoss EAP OpenShift platform .

Procedure

1. Create a directory named **scripts** at the root directory of your project. For example:

```
$ mkdir scripts
```

2. Create a **cli.properties** file and save the file in the **/scripts** directory. Define the **config.path** and the **config.ordinal** system properties in this file. For example:

```
config.path=/etc/config
config.ordinal=200
```

3. Create a CLI script, such as **mp-config.cli**, and save it in an accessible directory in the bootable JAR, such as the **/scripts** directory. The following example shows the contents of the **mp-config.cli** script:

```
# config map

/subsystem=microprofile-config-smallrye/config-source=os-map:add(dir=
{path=${config.path}}, ordinal=${config.ordinal})
```

   The **mp-config.cli** CLI script creates a new **ConfigSource**, to which ordinal and path values are retrieved from a properties file.

4. Save the script in the **/scripts** directory, which is located at the root directory of the project.

5. Add the following configuration extract to the existing plug-in **<configuration>** element:

```
<cli-sessions>
    <cli-session>
        <properties-file>
            scripts/cli.properties
        </properties-file>
        <script-files>
            <script>scripts/mp-config.cli</script>
        </script-files>
    </cli-session>
</cli-sessions>
```

6. Package the application:

```
$ mvn package
```

7. Log in to your OpenShift instance using the **oc login** command.

8. *Optional:* If you have not previously created a **target/openshift** subdirectory, you must create the suddirectory by issuing the following command:

```
$ mkdir target/openshift
```

9. Copy the packaged application into the created subdirectory.

```
$ cp target/microprofile-config-bootable.jar target/openshift
```

10. Use the **target/openshift** subdirectory as the binary input to build the application:

```
$ oc start-build microprofile-config-app --from-dir target/openshift
```

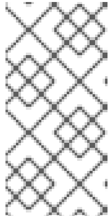

NOTE

OpenShift applies a set of CLI script commands to the bootable JAR configuration file to enable it for the cloud environment. You can access this script by opening the **bootable-jar-build-artifacts/generated-cli-script.txt** file in the Maven project /**target** directory.

11. Create a **ConfigMap**. For example:

```
$ oc create configmap microprofile-config-map --from-literal=name="Name comes from Openshift ConfigMap"
```

12. Mount the **ConfigMap** into the application with the dc. For example:

```
$ oc set volume deployments/microprofile-config-app --add --name=config-volume \
--mount-path=/etc/config \
--type=configmap \
--configmap-name=microprofile-config-map
```

After executing the **oc set volume** command, the application is re-deployed with the new configuration settings.

13. Test the output:

```
$ curl http://$(oc get route microprofile-config-app --template='{{ .spec.host }}')/config/json
```

The following is the expected output:

```
{"result":"Hello Name comes from Openshift ConfigMap"}
```

Additional resources

- For information about MicroProfile Config **ConfigSources** attributes, see Default MicroProfile Config attributes.
- For information about bootable JAR arguments, see Supported bootable JAR arguments.

8.12. CREATING A BOOTABLE JAR MAVEN PROJECT

Follow the steps in the procedure to create an example Maven project. You must create a Maven project before you can perform the following procedures:

- Enabling JSON logging for your bootable JAR
- Enabling web session data storage for multiple bootable JAR instances
- Enabling HTTP authentication for bootable JAR with a CLI script
- Securing your JBoss EAP bootable JAR application with Red Hat Single Sign-On

In the project **pom.xml** file, you can configure Maven to retrieve the project artifacts required to build your bootable JAR.

Procedure

1. Set up the Maven project:

```
$ mvn archetype:generate \
-DgroupId=GROUP_ID \
-DartifactId=ARTIFACT_ID \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
```

   Where *GROUP_ID* is the **groupId** of your project and *ARTIFACT_ID* is the **artifactId** of your project.

2. In the **pom.xml** file, configure Maven to retrieve the JBoss EAP BOM file from a remote repository.

```
<repositories>
    <repository>
        <id>jboss</id>
        <url>https://maven.repository.redhat.com/ga</url>
        <snapshots>
            <enabled>false</enabled>
        </snapshots>
    </repository>
</repositories>
<pluginRepositories>
  <pluginRepository>
      <id>jboss</id>
        <url>https://maven.repository.redhat.com/ga</url>
        <snapshots>
            <enabled>false</enabled>
        </snapshots>
  </pluginRepository>
</pluginRepositories>
```

3. To configure Maven to automatically manage versions for the Jakarta EE artifacts in the **jboss-eap-jakartaee8** BOM, add the BOM to the **<dependencyManagement>** section of the project **pom.xml** file. For example:

```
<dependencyManagement>
  <dependencies>
    <dependency>
        <groupId>org.jboss.bom</groupId>
        <artifactId>jboss-eap-jakartaee8</artifactId>
        <version>7.3.4.GA</version>
        <type>pom</type>
        <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

4. Add the servlet API artifact, which is managed by the BOM, to the **<dependency>** section of the project **pom.xml** file, as shown in the following example:

```
<dependency>
    <groupId>org.jboss.spec.javax.servlet</groupId>
    <artifactId>jboss-servlet-api_4.0_spec</artifactId>
    <scope>provided</scope>
</dependency>
```

Additional resources

- For information about the JBoss EAP Maven plug-in, see JBoss EAP Maven plug-in.
- For information about the Galleon layers, see Specifying Galleon layers for your bootable JAR server.
- For information about including the Red Hat Single Sign-On Galleon feature pack in your project, see Securing your JBoss EAP bootable JAR application with Red Hat Single Sign-On .

8.13. ENABLING JSON LOGGING FOR YOUR BOOTABLE JAR

You can enable JSON logging for your bootable JAR by configuring the server logging configuration with a CLI script. When you enable JSON logging, you can use the JSON formatter to view log messages in JSON format.

The example in this procedure shows you how to enable JSON logging for your bootable JAR on a bare-metal platform and an OpenShift platform.

Prerequisites

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the minor version of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.
- You have created a Maven project, set up a parent dependency, and added dependencies for creating an application. See Creating a bootable JAR Maven project .

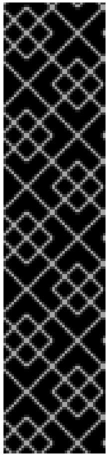

IMPORTANT

In the Maven archetype of your Maven project, you must specify the groupID and artifactID that are specific to your project. For example:

```
$ mvn archetype:generate \
-DgroupId=com.example.logging \
-DartifactId=logging \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
cd logging
```

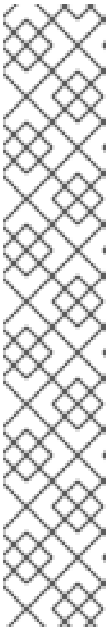

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Add the JBoss Logging and Jakarta RESTful Web Services dependencies, which are managed by the BOM, to the **<dependencies>** section of the project **pom.xml** file. For example:

```
<dependencies>
    <dependency>
        <groupId>org.jboss.logging</groupId>
        <artifactId>jboss-logging</artifactId>
        <scope>provided</scope>
    </dependency>
    <dependency>
        <groupId>org.jboss.spec.javax.ws.rs</groupId>
        <artifactId>jboss-jaxrs-api_2.1_spec</artifactId>
        <scope>provided</scope>
    </dependency>
</dependencies>
```

2. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
```

```
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
        <artifactId>wildfly-jar-maven-plugin</artifactId>
        <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
            <feature-packs>
                <feature-pack>
                    <location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</location>
                </feature-pack>
            </feature-packs>
            <layers>
                <layer>jaxrs-server</layer>
            </layers>
        </configuration>
        <executions>
            <execution>
                <goals>
                    <goal>package</goal>
                </goals>
            </execution>
        </executions>
    </plugin>
</plugins>
```

3. Create the directory to store Java files:

```
$ mkdir -p APPLICATION_ROOT/src/main/java/com/example/logging/
```

Where **APPLICATION_ROOT** is the directory containing the **pom.xml** configuration file for the application.

4. Create a Java file **RestApplication.java** with the following content and save the file in the **APPLICATION_ROOT/src/main/java/com/example/logging/** directory:

```
package com.example.logging;
import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

@ApplicationPath("/")
public class RestApplication extends Application {
}
```

5. Create a Java file **HelloWorldEndpoint.java** with the following content and save the file in the **APPLICATION_ROOT/src/main/java/com/example/logging/** directory:

```
package com.example.logging;

import javax.ws.rs.Path;
import javax.ws.rs.core.Response;
import javax.ws.rs.GET;
import javax.ws.rs.Produces;

import org.jboss.logging.Logger;
@Path("/hello")
```

```
public class HelloWorldEndpoint {

    private static Logger log = Logger.getLogger(HelloWorldEndpoint.class.getName());
    @GET
    @Produces("text/plain")
    public Response doGet() {
        log.debug("HelloWorldEndpoint.doGet called");
        return Response.ok("Hello from XP bootable jar!").build();
    }
}
```

6. Create a CLI script, such as **logging.cli**, and save it in an accessible directory in the bootable JAR, such as the ***APPLICATION_ROOT*/scripts** directory, where **APPLICATION_ROOT** is the root directory of your Maven project. The script must contain the following commands:

```
/subsystem=logging/logger=com.example.logging:add(level=ALL)
/subsystem=logging/json-formatter=json-formatter:add(exception-output-type=formatted,
pretty-print=false, meta-data={version="1"}, key-overrides={timestamp="@timestamp"})
/subsystem=logging/console-handler=CONSOLE:write-attribute(name=level,value=ALL)
/subsystem=logging/console-handler=CONSOLE:write-attribute(name=named-formatter,
value=json-formatter)
```

7. Add the following configuration extract to the plug-in **<configuration>** element:

```
<cli-sessions>
        <cli-session>
        <script-files>
            <script>scripts/logging.cli</script>
        </script-files>
    </cli-session>
</cli-sessions>
```

This example shows the **logging.cli** CLI script, which modifies the server logging configuration file to enable JSON logging for your application.

8. Package the application as a bootable JAR.

```
$ mvn package
```

9. *Optional*: To run the application on a JBoss EAP bare-metal platform, follow the steps outlined in Using a bootable JAR on a JBoss EAP bare-metal platform , but with the following difference:

   a. Start the application:

   ```
   mvn wildfly-jar:run
   ```

   b. Verification: You can access the application by specifying the following URL in your browser: http://127.0.0.1:8080/hello.
   Expected output: You can view the JSON-formatted logs, including the **com.example.logging.HelloWorldEndpoint** debug trace, in the application console.

10. *Optional*: To run the application on a JBoss EAP OpenShift platform, complete the following steps:

a. Add the **<cloud/>** element to the plug-in configuration. For example:

```
<plugins>
  <plugin>
    ... <!-- You must evolve the existing configuration with the <cloud/> element  -->
    <configuration >
      ...
      <cloud/>
    </configuration>
  </plugin>
</plugins>
```

b. Rebuild the application:

```
$ mvn clean package
```

c. Log in to your OpenShift instance using the **oc login** command.

d. Create a new project in OpenShift. For example:

```
$ oc new-project bootable-jar-project
```

e. Enter the following **oc** commands to create an application image:

```
$ mkdir target/openshift && cp target/logging-bootable.jar target/openshift 1

$ oc import-image ubi8/openjdk-11 --from=registry.redhat.io/ubi8/openjdk-11 --confirm
2

$ oc new-build --strategy source --binary --image-stream openjdk-11 --name logging 3

$ oc start-build logging --from-dir target/openshift 4
```

1. Creates the **target/openshift** subdirectory. The packaged application is copied into the **openshift** subdirectory.
2. Imports the latest OpenJDK 11 imagestream tag and image information into the OpenShift project.
3. Creates a build configuration based on the logging directory and the OpenJDK 11 imagestream.
4. Uses the **target/openshift** subdirectory as the binary input to build the application.

f. Deploy the application:

```
$ oc new-app logging

$ oc expose svc/logging
```

g. Get the URL of the route.

```
$ oc get route logging --template='{{ .spec.host }}'
```

h. Access the application in your web browser using the URL returned from the previous command. For example:

```
http://ROUTE_NAME/hello
```

i. *Verification:* Issue the following command to view a list of OpenShift pods available, and to check the pods build statuses:

```
$ oc get pods
```

Access a running pod log of your application. Where **APP_POD_NAME** is the name of the running pod logging application.

```
$ oc logs APP_POD_NAME
```

Expected outcome: The pod log is in JSON format and includes the **com.example.logging.HelloWorldEndpoint** debug trace.

Additional resources

- For information about logging functionality for JBoss EAP, see Logging with JBoss EAP in the *Configuration Guide*.
- For information about using a bootable JAR on OpenShift, see Using a bootable JAR on a JBoss EAP OpenShift platform.
- For information about specifying the JBoss EAP JAR Maven for your project, see Specifying Galleon layers for your bootable JAR server.
- For information about creating CLI scripts, see CLI scripts.

8.14. ENABLING WEB SESSION DATA STORAGE FOR MULTIPLE BOOTABLE JAR INSTANCES

You can build and package a web-clustering application as a bootable JAR.

Prerequisites

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.
- You have created a Maven project, set up a parent dependency, and added dependencies for creating a web-clustering application. See Creating a bootable JAR Maven project .

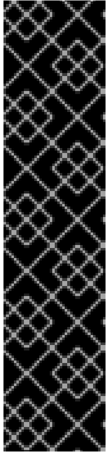

IMPORTANT

When setting up the Maven project, you must specify values in the Maven archetype configuration. For example:

```
$ mvn archetype:generate \
-DgroupId=com.example.webclustering \
-DartifactId=web-clustering \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
cd web-clustering
```

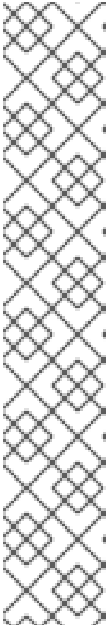

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
        <artifactId>wildfly-jar-maven-plugin</artifactId>
        <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
            <feature-pack-location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
            <layers>
                <layer>datasources-web-server</layer>
                <layer>web-clustering</layer>
            </layers>
        </configuration>
        <executions>
            <execution>
                <goals>
                    <goal>package</goal>
```

```
            </goals>
          </execution>
        </executions>
    </plugin>
</plugins>
```

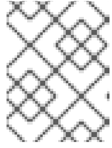

NOTE

This example makes use of the **web-clustering** Galleon layer to enable web session sharing.

2. Update the **web.xml** file in the **src/main/webapp/WEB-INF** directory with the following configuration:

```
<?xml version="1.0" encoding="UTF-8"?>

<web-app version="4.0"
         xmlns="http://xmlns.jcp.org/xml/ns/javaee"
         xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
         xsi:schemaLocation="http://xmlns.jcp.org/xml/ns/javaee
http://xmlns.jcp.org/xml/ns/javaee/web-app_4_0.xsd">
    <distributable/>
</web-app>
```

The **<distributable/>** tag indicates that this servlet can be distributed across multiple servers.

3. Create the directory to store Java files:

```
$ mkdir -p APPLICATION_ROOT
/src/main/java/com/example/webclustering/
```

Where **APPLICATION_ROOT** is the directory containing the **pom.xml** configuration file for the application.

4. Create a Java file **MyServlet.java** with the following content and save the file in the ***APPLICATION_ROOT*/src/main/java/com/example/webclustering/** directory.

```
package com.example.webclustering;

import java.io.IOException;
import java.io.PrintWriter;
import javax.servlet.ServletException;
import javax.servlet.annotation.WebServlet;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;

@WebServlet(urlPatterns = {"/clustering"})
public class MyServlet extends HttpServlet {
    @Override
    protected void doGet(HttpServletRequest request, HttpServletResponse response)
            throws IOException {
        response.setContentType("text/html;charset=UTF-8");
        long t;
```

```
        User user = (User) request.getSession().getAttribute("user");
        if (user == null) {
            t = System.currentTimeMillis();
            user = new User(t);
            request.getSession().setAttribute("user", user);
        }
        try (PrintWriter out = response.getWriter()) {
            out.println("<!DOCTYPE html>");
            out.println("<html>");
            out.println("<head>");
            out.println("<title>Web clustering demo</title>");
            out.println("</head>");
            out.println("<body>");
            out.println("<h1>Session id " + request.getSession().getId() + "</h1>");
            out.println("<h1>User Created " + user.getCreated() + "</h1>");
            out.println("<h1>Host Name " + System.getenv("HOSTNAME") + "</h1>");
            out.println("</body>");
            out.println("</html>");
        }
    }
}
```

The content in **MyServlet.java** defines the endpoint to which a client sends an HTTP request.

5. Create a Java file **User.java** with the following content and save the file in the ***APPLICATION_ROOT*/src/main/java/com/example/webclustering/** directory.

```
package com.example.webclustering;

import java.io.Serializable;

public class User implements Serializable {
    private final long created;

    User(long created) {
        this.created = created;
    }
    public long getCreated() {
        return created;
    }
}
```

6. Package the application:

```
$ mvn package
```

7. *Optional*: To run the application on a JBoss EAP bare-metal platform, follow the steps outlined in Using a bootable JAR on a JBoss EAP bare-metal platform , but with the following difference:

   a. On a JBoss EAP bare-metal platform, you can use the **java -jar** command to run multiple bootable JAR instances, as demonstrated in the following examples:

```
$ java -jar target/web-clustering-bootable.jar -Djboss.node.name=node1

$ java -jar target/web-clustering-bootable.jar -Djboss.node.name=node2 -
```

```
Djboss.socket.binding.port-offset=10
```

b. *Verification*: You can access the application on the node 1 instance: http://127.0.0.1:8080/clustering. Note the user session ID and the user-creation time. After you kill this instance, you can access the node 2 instance: http://127.0.0.1:8090/clustering. The user must match the session ID and the user-creation time of the node 1 instance.

8. *Optional*: To run the application on a JBoss EAP OpenShift platform, follow the steps outlined in Using a bootable JAR on a JBoss EAP OpenShift platform , but complete the following steps:

a. Add the **<cloud/>** element to the plug-in configuration. For example:

```
<plugins>
  <plugin>
    ... <!-- You must evolve the existing configuration with the <cloud/> element  -->
    <configuration >
      ...
      <cloud/>
     </configuration>
  </plugin>
</plugins>
```

b. Rebuild the application:

```
$ mvn clean package
```

c. Log in to your OpenShift instance using the **oc login** command.

d. Create a new project in OpenShift. For example:

```
$ oc new-project bootable-jar-project
```

e. To run a web-clustering application on a JBoss EAP OpenShift platform, authorization access must be granted for the service account that the pod is running in. The service account can then access the Kubernetes REST API. The following example shows authorization access being granted to a service account:

```
$ oc policy add-role-to-user view system:serviceaccount:$(oc project -q):default
```

f. Enter the following **oc** commands to create an application image:

```
$ mkdir target/openshift && cp target/web-clustering-bootable.jar target/openshift (1)

$ oc import-image ubi8/openjdk-11 --from=registry.redhat.io/ubi8/openjdk-11 --confirm
(2)

$ oc new-build --strategy source --binary --image-stream openjdk-11 --name web-
clustering (3)

$ oc start-build web-clustering --from-dir target/openshift (4)
```

(1) Creates the **target/openshift** sub-directory. The packaged application is copied into the **openshift** sub-directory.

2. Imports the latest OpenJDK 11 imagestream tag and image information into the OpenShift project.

3. Creates a build configuration based on the web-clustering directory and the OpenJDK 11 imagestream.

4. Uses the **target/openshift** sub-directory as the binary input to build the application.

g. Deploy the application:

```
$ oc new-app web-clustering -e KUBERNETES_NAMESPACE=$(oc project -q)

$ oc expose svc/web-clustering
```

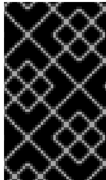

IMPORTANT

You must use the **KUBERNETES_NAMESPACE** environment variable to view other pods in the current OpenShift namespace; otherwise, the server attempts to retrieve the pods from the **default** namespace.

h. Get the URL of the route.

```
$ oc get route web-clustering --template='{{ .spec.host }}'
```

i. Access the application in your web browser using the URL returned from the previous command. For example:

```
http://ROUTE_NAME/clustering
```

Note the user session ID and user creation time.

j. Scale the application to two pods:

```
$ oc scale --replicas=2 deployments web-clustering
```

k. Issue the following command to view a list of OpenShift pods available, and to check the pods build statuses:

```
$ oc get pods
```

l. Kill the oldest pod using the **oc delete pod web-clustering-*POD_NAME*** command, where *POD_NAME* is the name of your oldest pod.

m. Access the application again:

```
http://ROUTE_NAME/clustering
```

Expected outcome: The session ID and the creation time generated by the new pod match those of the of the terminated pod. This indicates that web session data storage is enabled.

Additional resources

- For information about distributable web session management profiles, see The distributable-web subsystem for Distributable Web Session Configurations in the *Development Guide*.
- For information about configuring the JGroups protocol stack, see Configuring a JGroups Discovery Mechanism in the *Getting Started with JBoss EAP for OpenShift Container Platform* guide.

8.15. ENABLING HTTP AUTHENTICATION FOR BOOTABLE JAR WITH A CLI SCRIPT

You can enable HTTP authentication for the bootable JAR with a CLI script. This script adds a security realm and a security domain to your server.

Prerequisites

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.
- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.
- You have created a Maven project, set up a parent dependency, and added dependencies for creating an application that requires HTTP authentication. See Creating a bootable JAR Maven project.

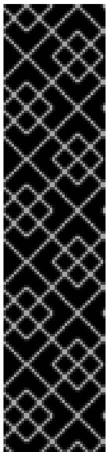

IMPORTANT

When setting up the Maven project, you must specify HTTP authentication values in the Maven archetype configuration. For example:

```
$ mvn archetype:generate \
-DgroupId=com.example.auth \
-DartifactId=authentication \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
cd authentication
```

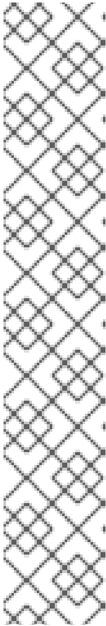

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

Procedure

1. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
        <artifactId>wildfly-jar-maven-plugin</artifactId>
        <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
            <feature-pack-location>org.jboss.eap:wildfly-galleon-pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
            <layers>
                <layer>datasources-web-server</layer>
            </layers>
        </configuration>
        <executions>
            <execution>
                <goals>
                    <goal>package</goal>
                </goals>
            </execution>
        </executions>
    </plugin>
</plugins>
```

The example shows the inclusion of the **datasources-web-server** Galleon layer that contains the **elytron** subsystem.

2. Update the **web.xml** file in the **src/main/webapp/WEB-INF** directory. For example:

```
<?xml version="1.0" encoding="UTF-8"?>

<web-app version="4.0"
        xmlns="http://xmlns.jcp.org/xml/ns/javaee"
```

```
         xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
         xsi:schemaLocation="http://xmlns.jcp.org/xml/ns/javaee
http://xmlns.jcp.org/xml/ns/javaee/web-app_4_0.xsd">

    <login-config>
        <auth-method>BASIC</auth-method>
        <realm-name>Example Realm</realm-name>
    </login-config>

</web-app>
```

3. Create the directory to store Java files:

```
$ mkdir -p APPLICATION_ROOT/src/main/java/com/example/authentication/
```

Where **APPLICATION_ROOT** is the root directory of your Maven project.

4. Create a Java file **TestServlet.java** with the following content and save the file in the **APPLICATION_ROOT/src/main/java/com/example/authentication/** directory.

```
package com.example.authentication;

import javax.servlet.annotation.HttpMethodConstraint;
import javax.servlet.annotation.ServletSecurity;
import javax.servlet.annotation.WebServlet;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;

import java.io.IOException;
import java.io.PrintWriter;

@WebServlet(urlPatterns = "/hello")
@ServletSecurity(httpMethodConstraints = { @HttpMethodConstraint(value = "GET",
rolesAllowed = { "Users" }) })
public class TestServlet extends HttpServlet {

    @Override
    protected void doGet(HttpServletRequest req, HttpServletResponse resp) throws
IOException {
        PrintWriter writer = resp.getWriter();
        writer.println("Hello " + req.getUserPrincipal().getName());
        writer.close();
    }

}
```

5. Create a CLI script, such as **authentication.cli**, and save it in an accessible directory in the bootable JAR, such as the ***APPLICATION_ROOT*/scripts** directory. The script must contain the following commands:

```
/subsystem=elytron/properties-realm=bootable-realm:add(users-properties={relative-
to=jboss.server.config.dir, path=bootable-users.properties, plain-text=true}, groups-
properties={relative-to=jboss.server.config.dir, path=bootable-groups.properties})
/subsystem=elytron/security-domain=BootableDomain:add(default-realm=bootable-realm,
```

```
permission-mapper=default-permission-mapper, realms=[{realm=bootable-realm, role-decoder=groups-to-roles}])

/subsystem=undertow/application-security-domain=other:write-attribute(name=security-domain, value=BootableDomain)
```

6. Add the following configuration extract to the plug-in **<configuration>** element:

```
<cli-sessions>
    <cli-session>
        <script-files>
            <script>scripts/authentication.cli</script>
        </script-files>
    </cli-session>
</cli-sessions>
```

This example shows the **authentication.cli** CLI script, which configures the default **undertow** security domain to the security domain defined for your server.

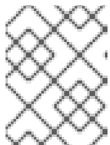

NOTE

You have the option to execute the CLI script at runtime instead of packaging time. To do so, skip this step and proceed to step 10.

7. In the root directory of your Maven project create a directory to store the properties files that the JBoss EAP JAR Maven plug-in adds to the bootable JAR:

```
$ mkdir -p APPLICATION_ROOT/extra-content/standalone/configuration/
```

Where **APPLICATION_ROOT** is the directory containing the **pom.xml** configuration file for the application.

This directory stores files such as **bootable-users.properties** and **bootable-groups.properties** files.

The **bootable-users.properties** file contains the following content:

```
testuser=bootable_password
```

The **bootable-groups.properties** file contains the following content:

```
testuser=Users
```

8. Add the following **extra-content-content-dirs** element to the existing **<configuration>** element:

```
<extra-server-content-dirs>
            <extra-content>extra-content</extra-content>
</extra-server-content-dirs>
```

The **extra-content** directory contains the properties files.

9. Package the application as a bootable JAR.

```
$ mvn package
```

10. Start the application:

```
mvn wildfly-jar:run
```

If you have chosen to skip step 6 and not execute the CLI script during build, launch the application with the following command:

```
mvn wildfly-jar:run -Dwildfly.bootable.arguments=--cli-script=scripts/authentication.cli
```

11. Call the servlet, but do not specify credentials:

```
curl -v http://localhost:8080/hello
```

Expected output:

```
HTTP/1.1 401 Unauthorized
...
WWW-Authenticate: Basic realm="Example Realm"
```

12. Call the server and specify your credentials. For example:

```
$ curl -v -u testuser:bootable_password http://localhost:8080/hello
```

A HTTP 200 status is returned that indicates HTTP authentication is enabled for your bootable JAR. For example:

```
HTTP/1.1 200 OK
....
Hello testuser
```

Additional resources

- For information about enabling HTTP authentication for the **undertow** security domain, see Enable HTTP Authentication for Applications Using the CLI Security Command in the *How to Configure Server Security*.

8.16. SECURING YOUR JBOSS EAP BOOTABLE JAR APPLICATION WITH RED HAT SINGLE SIGN-ON

You can use the Galleon **keycloak-client-oidc** layer to install a version of a server that is provisioned with Red Hat Single Sign-On 7.5 OpenID Connect client adapters.

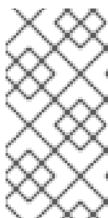

NOTE

The use of **keycloak-client-oidc** layer has been deprecated in JBoss EAP XP 4. Use the **elytron-oidc-client** layer, which provides a native OpenID Connect (OIDC) client, instead. For more information, see Developing JBoss EAP bootable jar application with OpenID Connect.

The **keycloak-client-oidc** layer provides Red Hat Single Sign-On OpenID Connect client adapters to your Maven project. This layer is included with the **keycloak-adapter-galleon-pack** Red Hat Single Sign-On feature pack.

You can add the **keycloak-adapter-galleon-pack** feature pack to your JBoss EAP Maven plug-in configuration and then add the **keycloak-client-oidc**. You can view Red Hat Single Sign-On client adapters that are compatible with JBoss EAP by visiting the Supported Configurations: Red Hat Single Sign-On 7.4 web page.

The example in this procedure shows you how to secure a JBoss EAP bootable JAR by using JBoss EAP features provided by the **keycloak-client-oidc** layer.

Prerequisites

- You have checked the latest Maven plug-in version, such as ***MAVEN_PLUGIN_VERSION*.X.GA.Final-redhat-00001**, where *MAVEN_PLUGIN_VERSION* is the major version and *X* is the microversion. See Index of /ga/org/wildfly/plugins/wildfly-jar-maven-plugin.

- You have checked the latest Galleon feature pack version, such as **4.0.X.GA-redhat-*BUILD_NUMBER***, where *X* is the microversion of JBoss EAP XP and *BUILD_NUMBER* is the build number of the Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/jboss/eap/wildfly-galleon-pack.

- You have checked the latest Red Hat Single Sign-On Galleon feature pack version, such as **org.keycloak:keycloak-adapter-galleon-pack:15.0.X.redhat-BUILD_NUMBER**, where **X** is the microversion of Red Hat Single Sign-On that depends on the Red Hat Single Sign-On server release used to secure the application, and **BUILD_NUMBER** is the build number of the Red Hat Single Sign-On Galleon feature pack. Both *X* and *BUILD_NUMBER* can evolve during the JBoss EAP XP 4.0.0 product life cycle. See Index of /ga/org/keycloak/keycloak-adapter-galleon-pack.

- You have created a Maven project, set up a parent dependency, and added dependencies for creating an application that you want secured with Red Hat Single Sign-On. See Creating a bootable JAR Maven project.

- You have a Red Hat Single Sign-On server that is running on port 8090. See Starting the Red Hat Single Sign-On server.

- You have logged in to the Red Hat Single Sign-On Admin Console and created the following metadata:

  - A realm named **demo**.

  - A role named **Users**.

  - A user and password. You must assign a **Users** role to the user.

  - A **public-client** web application with a Root URL. The example in the procedure, defines **simple-webapp** as the web application and **http://localhost:8080/simple-webapp/secured** as the Root URL.

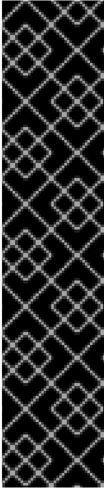

IMPORTANT

When setting up the Maven project, you must specify values for the application that you want to secure with Red Hat Single Sign-On in the Maven archetype. For example:

```
$ mvn archetype:generate \
-DgroupId=com.example.keycloak \
-DartifactId=simple-webapp \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
cd simple-webapp
```

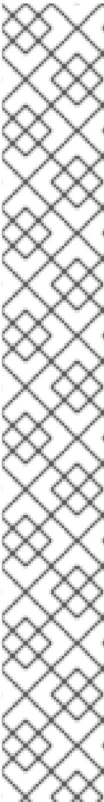

NOTE

The examples shown in the procedure specify the following properties:

- **${bootable.jar.maven.plugin.version}** for the Maven plug-in version.
- **${jboss.xp.galleon.feature.pack.version}** for the Galleon feature pack version.
- **${keycloak.feature.pack.version}** for the Red Hat Single Sign-On feature pack version.

You must set these properties in your project. For example:

```
<properties>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
    <keycloak.feature.pack.version>15.0.4.redhat-
00001</keycloak.feature.pack.version>
</properties>
```

Procedure

1. Add the following content to the **<build>** element of the **pom.xml** file. You must specify the latest version of any Maven plug-in and the latest version of the **org.jboss.eap:wildfly-galleon-pack** Galleon feature pack. For example:

```
<plugins>
    <plugin>
        <groupId>org.wildfly.plugins</groupId>
        <artifactId>wildfly-jar-maven-plugin</artifactId>
        <version>${bootable.jar.maven.plugin.version}</version>
        <configuration>
            <feature-packs>
                <feature-pack>
                    <location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</location>
                </feature-pack>
                <feature-pack>
```

```
                <location>org.keycloak:keycloak-adapter-galleon-pack:${keycloak.feature.pack.version}</location>
            </feature-pack>
        </feature-packs>
        <layers>
            <layer>datasources-web-server</layer>
            <layer>keycloak-client-oidc</layer>
        </layers>
    </configuration>
    <executions>
        <execution>
            <goals>
                <goal>package</goal>
            </goals>
        </execution>
    </executions>
  </plugin>
</plugins>
```

The Maven plug-in provisions subsystems and modules that are required for deploying the web application.

The **keycloak-client-oidc** layer provides Red Hat Single Sign-On OpenID Connect client adapters to your project by using the **keycloak** subsystem and its dependencies to activate support for Red Hat Single Sign-On authentication. Red Hat Single Sign-On client adapters are libraries that secure applications and services with Red Hat Single Sign-On.

2. In the project **pom.xml** file, set the **<context-root>** to **false** in your plug-in configuration. This registers the application in the **simple-webapp** resource path. By default, the WAR file is registered under the root-context path.

```
<configuration>
  ...
   <context-root>false</context-root>
   ...
</configuration>
```

3. Create a CLI script, such as **configure-oidc.cli** and save it in an accessible directory in the bootable JAR, such as the ***APPLICATION_ROOT*/scripts** directory, where *APPLICATION_ROOT* is the root directory of your Maven project. The script must contain commands similar to the following example:

```
/subsystem=keycloak/secure-deployment=simple-webapp.war:add( \
  realm=demo, \
  resource=simple-webapp, \
  public-client=true, \
  auth-server-url=http://localhost:8090/auth/, \
  ssl-required=EXTERNAL)
```

This script example defines the **secure-deployment=simple-webapp.war** resource in the **keycloak** subsystem. The **simple-webapp.war** resource is the name of the WAR file that is deployed in the bootable JAR.

4. In the project **pom.xml** file, add the following configuration extract to the existing plug-in **<configuration>** element:

```
<cli-sessions>
    <cli-session>
        <script-files>
            <script>scripts/configure-oidc.cli</script>
        </script-files>
    </cli-session>
</cli-sessions>
```

5. Update the **web.xml** file in the **src/main/webapp/WEB-INF** directory. For example:

```
<?xml version="1.0" encoding="UTF-8"?>

<web-app version="4.0" xmlns="http://java.sun.com/xml/ns/javaee"
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://java.sun.com/xml/ns/javaee/web-app_4_0.xsd"
   metadata-complete="false">

    <login-config>
        <auth-method>BASIC</auth-method>
        <realm-name>Simple Realm</realm-name>
    </login-config>

</web-app>
```

6. *Optional:* Alternatively to steps 7 through 9, you can embed the server configuration in the web application by adding the **keycloak.json** descriptor to the **WEB-INF** directory of the web application. For example:

```
{
    "realm" : "demo",
    "resource" : "simple-webapp",
    "public-client" : "true",
    "auth-server-url" : "http://localhost:8090/auth/",
    "ssl-required" : "EXTERNAL"
}
```

You must then set the **<auth-method>** of the web application to **KEYCLOAK**. The following example code illustrates how to set the **<auth-method>**:

```
    <login-config>
        <auth-method>KEYCLOAK</auth-method>
        <realm-name>Simple Realm</realm-name>
    </login-config>
```

7. Create a Java file named **SecuredServlet.java** with the following content and save the file in the ***APPLICATION_ROOT*/src/main/java/com/example/securedservlet/** directory.

```
package com.example.securedservlet;

import java.io.IOException;
import java.io.PrintWriter;
import java.security.Principal;
```

```
import javax.servlet.ServletException;
import javax.servlet.annotation.HttpMethodConstraint;
import javax.servlet.annotation.ServletSecurity;
import javax.servlet.annotation.WebServlet;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;

@WebServlet("/secured")
@ServletSecurity(httpMethodConstraints = { @HttpMethodConstraint(value = "GET",
   rolesAllowed = { "Users" }) })
public class SecuredServlet extends HttpServlet {

   @Override
   protected void doGet(HttpServletRequest req, HttpServletResponse resp) throws
ServletException, IOException {
       try (PrintWriter writer = resp.getWriter()) {
           writer.println("<html>");
           writer.println("<head><title>Secured Servlet</title></head>");
           writer.println("<body>");
           writer.println("<h1>Secured Servlet</h1>");
           writer.println("<p>");
           writer.print(" Current Principal '");
           Principal user = req.getUserPrincipal();
           writer.print(user != null ? user.getName() : "NO AUTHENTICATED USER");
           writer.print("'");
           writer.println("    </p>");
           writer.println("  </body>");
           writer.println("</html>");
       }
   }
}
```

8. Package the application as a bootable JAR.

```
$ mvn package
```

9. Start the application. The following example starts the **simple-webapp** web application from its specified bootable JAR path:

```
$ java -jar target/simple-webapp-bootable.jar
```

10. Specify the following URL in your web browser to access the webpage secured with Red Hat Single Sign-On. The following example shows the URL for the secured **simple-webapp** web application:

```
http://localhost:8080/simple-webapp/secured
```

11. Log in as a user from your Red Hat Single Sign-On realm.

12. *Verification:* Check that the webpage displays the following output:

```
Current Principal '<principal id>'
```

Additional resources

- For information about configuring the Red Hat Single Sign-On adapter subsystem, see JBoss EAP Adapter in the *Securing Applications and Services Guide*.
- For information about specifying the JBoss EAP JAR Maven for your project, see Specifying Galleon layers for your bootable JAR server.

8.17. PACKAGING A BOOTABLE JAR IN DEV MODE

The JBoss EAP JAR Maven plug-in **dev goal** provides **dev** mode, Development Mode, which you can use to enhance your application development process.

In **dev** mode, you do not need to rebuild the bootable JAR after you make changes to your application.

The workflow in this procedure demonstrates using **dev** mode to configure a bootable JAR.

Prerequisites

- Maven is installed.
- You have created a Maven project, set up a parent dependency, and added dependencies for creating an MicroProfile application. See MicroProfile Config development.
- You have specified the JBoss EAP JAR Maven plug-in in your Maven project **pom.xml** file.

Procedure

1. Build and start the bootable JAR in Development Mode:

   ```
   $ mvn wildfly-jar:dev
   ```

   In **dev** mode, the server deployment scanner is configured to monitor the **target/deployments** directory.

2. Prompt the JBoss EAP Maven Plug-in to build and copy your application to the **target/deployments** directory with the following command:

   ```
   $ mvn package -Ddev
   ```

   The server packaged inside the bootable JAR deploys the application stored in the **target/deployments** directory.

3. Modify the code in your application code.
4. Use the **mvn package -Ddev** to prompt the JBoss EAP Maven Plug-in to re-build your application and re-deploy it.
5. Stop the server. For example:

   ```
   $ mvn wildfly-jar:shutdown
   ```

6. After you complete your application changes, package your application as a bootable JAR:

   ```
   $ mvn package
   ```

-

8.18. UPGRADING SERVER ARTIFACTS

A server artifact is jar file located inside the JBoss Modules module, you can reference it using the Maven coordinate in your project pom.xml file.

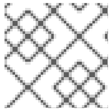

NOTE

Be aware that upgrading a server artifact can lead to an unsupported configuration.

Prerequisites

- Ensure that the Maven artifact is accessible from either your local Maven repository or remote Maven repository.

Procedure

1. Use the version of the artifact present in your dependencies during the build to successfully upgrade your server artifact. For example:

```
<dependencies>
  ...
  <dependency>
 <groupId>io.undertow</groupId>
 <artifactId>undertow-core</groupId>
  <version>2.2.5.Final-redhat-00001</version>
       <scope>provided</scope>
       <!-- In order to avoid bringing transitive dependencies to the project, exclude all
dependencies -->
       <exclusions>
         <exclusion>
           <groupId>*</groupId>
           <artifactId>*</artifactId>
         </exclusion>
       </exclusions>
  </dependency>
  ...
</dependencies>
```

2. Open the plugin **<configuration>** section and update the artifact groupId and artifactId inside the **<overridden-server-artifacts>** list, for example:

```
<configuration>
 ...
 <overridden-server-artifacts>
  <artifact>
   <groupId>io.undertow</groupId>
   <artifactId>undertow-core</groupId>
  </artifact>
 </overridden-server-artifacts>
</configuration>
```

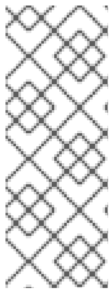

NOTE

- If the artifact added to **<overridden-server-artifacts>** is not found in the project dependencies, a failure will occur.
- If the artifact added to **<overridden-server-artifacts>**` is not among the provisioned server artifacts, a failure will occur since the target artifact for upgrade cannot be located.

8.19. UPDATING EAP 7.4.GA DEPENDENCY

When building a bootable JAR for JBoss EAP XP 4.0.0, you can update the dependency JBoss EAP XP 4.0.0 has on JBoss EAP 7.4. The JBoss EAP XP 4.0.0 galleon feature-pack **org.jboss.eap:wildfly-galleon-pack:4.0.0.GA-redhat-00001** has a dependency on the **org.jboss.eap:wildfly-ee-galleon-pack:7.4.0.GA-redhat-00001** that you can upgrade when building a Bootable JAR.

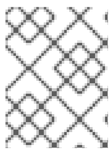

NOTE

Upgrade to the latest JBoss EAP XP version. This ensures that you get the latest updates in your JBoss EAP XP 4.0.0 bootable JAR.

Prerequisites

- You have the latest version of JBoss EAP XP.

Procedure

1. Ensure that the JBoss EAP Galleon feature-pack Maven artifact is accessible from either your local or remote Maven repositories.
2. Add the Galleon feature-pack artifact in the project dependencies:
   a. Set the scope to **provided**.
   b. Set the type to **zip**.
   c. Set the artifact version. For example:

```
<dependencies>
  ...
  <dependency>
 <groupId>org.jboss.eap</groupId>
 <artifactId>wildfly-ee-galleon-pack</groupId>
  <version>7.4.1.GA-redhat-00001</version>
<scope>provided</scope>
<type>zip</type>
  </dependency>
  ...
</dependencies>
```

3. Open the plugin **<configuration>** section and update the artifact groupId and artifactId inside the **<overridden-server-artifacts>** list, for example:

```
<configuration>
...
```

```
<overridden-server-artifacts>
 <artifact>
  <groupId>org.jboss.eap</groupId>
  <artifactId>wildfly-ee-galleon-pack</groupId>
 </artifact>
</overridden-server-artifacts>
</configuration>
```

4. Use the latest version of JBoss EAP XP Galleon feature-pack during the build to successfully the dependency.

8.20. APPLYING THE JBOSS EAP PATCH TO YOUR BOOTABLE JAR

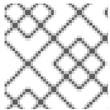

NOTE

In JBoss EAP XP 4.0.0, the legacy patching feature for bootable jar is deprecated.

On a JBoss EAP bare-metal platform, you can install the patch to your bootable JAR by using a CLI script.

The CLI script issues the **patch apply** command to apply the patch during the bootable JAR build.

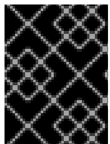

IMPORTANT

After you apply a patch to your bootable JAR, you cannot roll back from the applied patch. You must rebuild a bootable JAR without the patch.

Additionally, you can apply a legacy patch to your bootable JAR with the JBoss EAP JAR Maven plug-in. This plug-in provides a **<legacy-patch-cli-script>** configuration option to reference the CLI script that is used to patch the server.

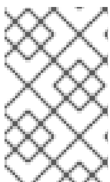

NOTE

The prefix **legacy-*** in **<legacy-patch-cli-script>** is related to applying archive patches to a bootable JAR. This method is similar to applying patches to regular JBoss EAP distributions.

You can use the **legacy-patch-cleanup** option in the JBoss EAP JAR Maven plug-in configuration to reduce the memory footprint of the bootable JAR by removing unused patch content. The option removes unused module dependencies. This option is set as **false** by default in the patch configuration file.

The **legacy-patch-cleanup** option removes the following patch content:

- The **<JBOSS_HOME>/.installation/patches** directory.
- Original locations of patch modules in the base layer.
- Unused modules that were added by the patch and are not referenced in the that existing module graph or patched modules graph.
- Overlays directories that are not listed in the **.overlays** file.

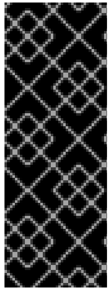

IMPORTANT

The **legacy-patch-clean-up** option variable is provided as a Technology Preview. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

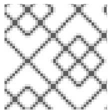

NOTE

The information outlined in this procedure also pertains to the hollow bootable JAR.

Prerequisites

- You have set up an account on the Red Hat Customer Portal.
- You have downloaded the following files from the **Product Downloads** page:
  - The JBoss EAP JBoss EAP 7.4.4 GA patch
  - The JBoss EAP XP 4.0.0 patch

Procedure

1. Create a CLI script that defines the legacy patches you want to apply to your bootable JAR. The script must contain one or more patch apply commands. The **--override-all** command is required when patching a server that was trimmed with Galleon layers, for example:

```
patch apply patch-oneoff1.zip --override-all

patch apply patch-oneoff2.zip --override-all

patch info --json-output
```

2. Reference your CLI script in the **<legacy-patch-cli-script>** element of your **pom.xml** file.
3. Rebuild the bootable JAR.

Additional resources

- For information about downloading the JBoss EAP MicroProfile Maven repository, see Downloading the JBoss EAP MicroProfile Maven repository patch as an archive file .
- For information about creating CLI scripts, see CLI Scripts.
- For information about Technology Preview features, see Technology Preview Features Support Scope on the *Red Hat Customer Portal*.

CHAPTER 9. OPENID CONNECT IN JBOSS EAP

Use the JBoss EAP native OpenID Connect (OIDC) client to secure your applications through an external OpenID provider. OIDC is an identity layer that enables clients, such as JBoss EAP, to verify a user's identity based on OpenID provider authentication. For example, you can secure your JBoss EAP applications using Red Hat Single Sign-On as the OpenID provider.

9.1. OPENID CONNECT CONFIGURATION IN JBOSS EAP

When you secure your applications using an OpenID provider, you do not need to configure any security domain resources locally. The **elytron-oidc-client** subsystem provides a native OpenID Connect (OIDC) client in JBoss EAP to connect with OpenID providers. JBoss EAP automatically creates a virtual security domain for your application, based on your OpenID provider configurations.

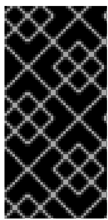

IMPORTANT

It is recommended to use the OIDC client with Red Hat Single Sign-On. You can use other OpenID providers if they can be configured to use access tokens that are JSON Web Tokens (JWTs) and can be configured to use the RS256, RS384, RS512, ES256, ES384, or ES512 signature algorithm.

To enable the use of OIDC, you can configure either the **elytron-oidc-client** subsystem or an application itself. JBoss EAP activates the OIDC authentication as follows:

- When you deploy an application to JBoss EAP, the **elytron-oidc-client** subsystem scans the deployment to detect if the OIDC authentication mechanism is required.
- If the subsystem detects OIDC configuration for the deployment in either the **elytron-oidc-client** subsystem or the application deployment descriptor, JBoss EAP enables the OIDC authentication mechanism for the application.
- If the subsystem detects OIDC configuration in both places, the configuration in the **elytron-oidc-client** subsystem **secure-deployment** attribute takes precedence over the configuration in the application deployment descriptor.

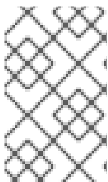

NOTE

The **keycloak-client-oidc** layer to secure your applications with Red Hat Single Sign-On is deprecated in JBoss EAP XP 4.0.0. Use the native OIDC client provided by the **elytron-oidc-client** subsystem instead.

Deployment configuration

To secure an application with OIDC by using a deployment descriptor, update the application's deployment configuration as follows:

- Create a file called **oidc.json** in the **WEB-INF** directory with the OIDC configuration information.

Example oidc.json contents

```
{
  "client-id" : "customer-portal", (1)
  "provider-url" : "http://localhost:8180/auth/realms/demo", (2)
```

```
  "ssl-required" : "external", (3)
  "credentials" : {
    "secret" : "234234-234234-234234" (4)
  }
}
```

1. The name to identify the OIDC client with the OpenID provider.
2. The OpenID provider URL.
3. Require HTTPS for external requests.
4. The client secret that was registered with the OpenID provider.

- Set the **auth-method** property to **OIDC** in the application deployment descriptor **web.xml** file.

Example deployment descriptor update

```
<login-config>
    <auth-method>OIDC</auth-method>
</login-config>
```

Subsystem configuration

You can secure applications with OIDC by configuring the **elytron-oidc-client** subsystem in the following ways:

- Create a single configuration for multiple deployments if you use the same OpenID provider for each application.
- Create a different configuration for each deployment if you use different OpenID providers for different applications.

Example XML configuration for a single deployment:

```
<subsystem xmlns="urn:wildfly:elytron-oidc-client:1.0">
    <secure-deployment name="DEPLOYMENT_RUNTIME_NAME.war"> (1)
        <client-id>customer-portal</client-id> (2)
        <provider-url>http://localhost:8180/auth/realms/demo</provider-url> (3)
        <ssl-required>external</ssl-required> (4)
        <credential name="secret" secret="0aa31d98-e0aa-404c-b6e0-e771dba1e798" /> (5)
    </secure-deployment
</subsystem>
```

1. The deployment runtime name.
2. The name to identify the OIDC client with the OpenID provider.
3. The OpenID provider URL.
4. Require HTTPS for external requests.
5. The client secret that was registered with the OpenID provider.

To secure multiple applications using the same OpenID provider, configure the **provider** separately, as shown in the example:

```
<subsystem xmlns="urn:wildfly:elytron-oidc-client:1.0">
  <provider name="${OpenID_provider_name}">
    <provider-url>http://localhost:8080/auth/realms/demo</provider-url>
    <ssl-required>external</ssl-required>
  </provider>
  <secure-deployment name="customer-portal.war"> 1
    <provider>${OpenID_provider_name}</provider>
    <client-id>customer-portal</client-id>
    <credential name="secret" secret="0aa31d98-e0aa-404c-b6e0-e771dba1e798" />
  </secure-deployment>
  <secure-deployment name="product-portal.war"> 2
    <provider>${OpenID_provider_name}</provider>
    <client-id>product-portal</client-id>
    <credential name="secret" secret="0aa31d98-e0aa-404c-b6e0-e771dba1e798" />
  </secure-deployment>
</subsystem>
```

1. A deployment: **customer-portal.war**
2. Another deployment: **product-portal.war**

Additional resources

- OpenID Connect specification
- **elytron-oidc-client** subsystem attributes
- OpenID Connect Libraries
- Securing applications using OpenID Connect with Red Hat Single Sign-On
- MicroProfile JWT

9.2. ENABLING THE ELYTRON-OIDC-CLIENT SUBSYSTEM

The **elytron-oidc-client** subsystem is provided in the **standalone-microprofile.xml** configuration file. To use it, you must start your server with the **bin/standalone.sh -c standalone-microprofile.xml** command. You can include the **elytron-oidc-client** subsystem in the **standalone.xml** configuration by enabling it using the management CLI.

Prerequisites

- You have installed JBoss EAP XP.

Procedure

1. Add the **elytron-oidc-client** extension using the management CLI.

```
/extension=org.wildfly.extension.elytron-oidc-client:add
```

2. Enable the **elytron-oidc-client** subsystem using the management CLI.

```
/subsystem=elytron-oidc-client:add
```

3. Reload JBoss EAP.

```
reload
```

You can now use the **elytron-oidc-client** subsystem by starting the server normally, with the command **bin/standalone.sh**

Additional resources

- **elytron-oidc-client** subsystem attributes

9.3. SECURING APPLICATIONS USING OPENID CONNECT WITH RED HAT SINGLE SIGN-ON

You can use OpenID Connect (OIDC) to delegate authentication to an external OpenID provider. The **elytron-oidc-client** subsystem provides a native OIDC client in JBoss EAP to connect with external OpenID providers.

To create an application secured with OpenID Connect using Red Hat Single Sign-On, follow these procedures:

- Configure Red Hat Single Sign-On as the OpenID provider
- Create a Maven project for your application
- Create an application that uses OpenID Connect
- Restrict access to your application based on user roles
- Create and assign user roles in Red Hat Single Sign-On

9.3.1. Configuring Red Hat Single Sign-On as an OpenID provider

Red Hat Single Sign-On is an identity and access management provider for securing web applications with single sign-on (SSO). It supports OpenID Connect (an extension to OAuth 2.0).

Prerequisites

- You have installed the Red Hat Single Sign-On server. For more information, see Installing the Red Hat Single Sign-On server in the Red Hat Single Sign-On *Getting Started Guide*.
- You have created a user in your Red Hat Single Sign-On server instance. For more information, see Creating a user in the Red Hat Single Sign-On *Getting Started Guide*.

Procedure

1. Start the Red Hat Single Sign-On server at a port other than 8080 because JBoss EAP default port is 8080.

   Syntax

```
$ RH_SSO_HOME/bin/standalone.sh -Djboss.socket.binding.port-offset=<offset-number>
```

Example

```
$ /home/servers/rh-sso-7.4/bin/standalone.sh -Djboss.socket.binding.port-offset=100
```

2. Log in to the Admin Console at **http://localhost:<port>/auth/**. For example, **http://localhost:8180/auth/**.

3. To create a realm, in the Admin Console, hover over **Master**, and click **Add realm**.

4. Enter a name for the realm. For example, **example_realm**. Ensure that **Enabled** is **ON** and click **Create**.

5. Click **Users**, then click **Add user** to add a user to the realm.

6. Enter a user name. For example, **jane_doe**. Ensure that User **Enabled** is **ON** and click **Save**.

7. Click **Credentials** to add a password to the user.

8. Set a password for the user. For example, **janedoep@$$**. Toggle **Temporary** to **OFF** and click **Set Password**. In the confirmation prompt, click **Set password**.

9. Click **Clients**, then click **Create** to configure a client connection.

10. Enter a client ID. For example, **my_jbeap**. Ensure that **Client Protocol** is set to **openid-connect**, and click **Save**.

11. Click **Installation**, then select **Keycloak OIDC JSON** as the **Format Option** to see the connection parameters.

```
{
  "realm": "example_realm",
  "auth-server-url": "http://localhost:8180/auth/",
  "ssl-required": "external",
  "resource": "my_jbeap",
  "public-client": true,
  "confidential-port": 0
}
```

When configuring your JBoss EAP application to use Red Hat Single Sign-On as the identity provider, you use the parameters as follows:

```
"provider-url" : "http://localhost:8180/auth/realms/example_realm",
"ssl-required": "external",
"client-id": "my_jbeap",
"public-client": true,
"confidential-port": 0
```

12. Click **Clients**, click **Edit** next to **my_jbeap** to edit the client settings.

13. In **Valid Redirect URIs**, enter the URL where the page should redirect after authentication is successful.
For this example, set this value to **http://localhost:8080/simple-oidc-example/secured/***

Additional resources

- Configuring a Maven project for creating a secure application
- Creating a realm and a user

9.3.2. Configuring a Maven project for creating a secure application

Create a Maven project with the required dependencies and the directory structure for creating a secure application.

Prerequisites

- You have installed Maven. For more information, see Downloading Apache Maven.
- You have configured your Maven repository for the latest release. For more information, see Maven and the JBoss EAP microprofile maven repository .

Procedure

1. Set up a Maven project using the **mvn** command. The command creates the directory structure for the project and the **pom.xml** configuration file.

   Syntax

   ```
   $ mvn archetype:generate \
   -DgroupId=${group-to-which-your-application-belongs} \
   -DartifactId=${name-of-your-application} \
   -DarchetypeGroupId=org.apache.maven.archetypes \
   -DarchetypeArtifactId=maven-archetype-webapp \
   -DinteractiveMode=false
   ```

   Example

   ```
   $ mvn archetype:generate \
   -DgroupId=com.example.oidc \
   -DartifactId=simple-oidc-example \
   -DarchetypeGroupId=org.apache.maven.archetypes \
   -DarchetypeArtifactId=maven-archetype-webapp \
   -DinteractiveMode=false
   ```

2. Navigate to the application root directory:

   Syntax

   ```
   $ cd <name-of-your-application>
   ```

   Example

   ```
   $ cd simple-oidc-example
   ```

3. Update the generated **pom.xml** file as follows:

a. Set the following properties:

```
<properties>
    <maven.compiler.source>1.8</maven.compiler.source>
    <maven.compiler.target>1.8</maven.compiler.target>
    <failOnMissingWebXml>false</failOnMissingWebXml>
    <version.server.bom>4.0.0.GA</version.server.bom>
    <version.server.bootable-jar>4.0.0.GA</version.server.bootable-jar>
    <version.wildfly-jar.maven.plugin>4.0.0.GA</version.wildfly-jar.maven.plugin>
</properties>
```

b. Set the following dependencies:

```
<dependencies>
    <dependency>
      <groupId>javax.servlet</groupId>
      <artifactId>javax.servlet-api</artifactId>
      <version>3.1.0.redhat-1</version>
      <scope>provided</scope>
    </dependency>
</dependencies>
```

c. Set the following build configuration to use **mvn widlfy:deploy** to deploy the application:

```
<build>
    <finalName>${project.artifactId}</finalName>
    <plugins>
        <plugin>
            <groupId>org.wildfly.plugins</groupId>
            <artifactId>wildfly-maven-plugin</artifactId>
            <version>2.1.0.Final</version>
        </plugin>
    </plugins>
</build>
```

Verification

- In the application root directory, enter the following command:

```
$ mvn install
```

You get an output similar to the following:

```
[INFO] ------------------------------------------------------------------------
[INFO] BUILD SUCCESS
[INFO] ------------------------------------------------------------------------
[INFO] Total time: 1.440 s
[INFO] Finished at: 2021-12-27T14:45:12+05:30
[INFO] ------------------------------------------------------------------------
```

You can now create your secure application.

Additional resources

- Creating a secure application that uses OpenID Connect

9.3.3. Creating a secure application that uses OpenID Connect

You can secure an application by either updating its deployment configuration or by configuring the **elytron-oidc-client** subsystem. The following example demonstrates creating a servlet that prints a logged-in user's Principal. For an existing application, only those steps that are related to updating the deployment configuration or the **elytron-oidc-client** subsystem are required.

In this example, the value of the Principal comes from the ID token from the OpenID provider. By default, the Principal is the value of the **"sub"** claim from the token. You can specify which claim value from the ID token to use as the Principal in one of the following:

- The **elytron-oidc-client** subsystem attribute **principal-attribute**.
- The **oidc.json** file.

<application_root> in the procedure denotes the **pom.xml** file directory. The **pom.xml** file contains your application's Maven configuration.

Prerequisites

- You have created a Maven project. For more information, see Configuring Maven project for creating a secure application .
- You have configured Red Hat Single Sign-On as the OpenID provider. For more information, see Configuring Red Hat Single Sign-On as an OpenID provider .
- You have enabled the **elytron-oidc-client** subsystem. For more information, see Enabling the elytron-oidc-client subsystem

Procedure

1. Create a directory to store the Java files.

   Syntax

   ```
   $ mkdir -p <application_root>/src/main/java/com/example/oidc
   ```

   Example

   ```
   $ mkdir -p simple-oidc-example/src/main/java/com/example/oidc
   ```

2. Navigate to the new directory.

   Syntax

   ```
   $ cd <application_root>/src/main/java/com/example/oidc
   ```

   Example

   ```
   $ cd simple-oidc-example/src/main/java/com/example/oidc
   ```

3. Create a servlet "SecuredServlet.java" with the following content:

```
package com.example.oidc;

import java.io.IOException;
import java.io.PrintWriter;
import java.security.Principal;

import javax.servlet.ServletException;
import javax.servlet.annotation.WebServlet;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;

/**
 * A simple secured HTTP servlet.
 *
 */
@WebServlet("/secured")
public class SecuredServlet extends HttpServlet {

    @Override
    protected void doGet(HttpServletRequest req, HttpServletResponse resp) throws
ServletException, IOException {
        try (PrintWriter writer = resp.getWriter()) {
            writer.println("<html>");
            writer.println("  <head><title>Secured Servlet</title></head>");
            writer.println("  <body>");
            writer.println("    <h1>Secured Servlet</h1>");
            writer.println("    <p>");
            writer.print(" Current Principal '");
            Principal user = req.getUserPrincipal();
            writer.print(user != null ? user.getName() : "NO AUTHENTICATED USER");
            writer.print("'");
            writer.println("    </p>");
            writer.println("  </body>");
            writer.println("</html>");
        }
    }

}
```

4. Add security rules for access to your application in the deployment descriptor **web.xml** file located in the **WEB-INF** directory of the application.

```
<?xml version="1.0" encoding="UTF-8"?>

<web-app version="2.5" xmlns="http://java.sun.com/xml/ns/javaee"
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd"
   metadata-complete="false">

    <security-constraint>
        <web-resource-collection>
            <web-resource-name>secured</web-resource-name>
```

```
            <url-pattern>/secured</url-pattern>
        </web-resource-collection>
        <auth-constraint>
            <role-name>*</role-name>
        </auth-constraint>
    </security-constraint>

    <security-role>
        <role-name>*</role-name>
    </security-role>
</web-app>
```

5. To secure the application with OpenID Connect, either update the deployment configuration or configure the **elytron-oidc-client** subsystem.

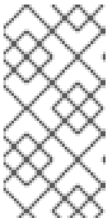

NOTE

If you configure OpenID Connect in both the deployment configuration and the **elytron-oidc-client** subsystem, the configuration in the **elytron-oidc-client** subsystem **secure-deployment** attribute takes precedence over the configuration in the application deployment descriptor.

- Updating the deployment configuration:

    i. Create a file **oidc.json** in the **WEB-INF** directory, like this:

    ```
    {
      "provider-url" : "http://localhost:8180/auth/realms/example_realm",
      "ssl-required": "external",
      "client-id": "my_jbeap",
      "public-client": true,
      "confidential-port": 0
    }
    ```

    ii. Update the deployment descriptor **web.xml** file with the following text to declare that this application uses OIDC:

    ```
    <login-config>
        <auth-method>OIDC</auth-method>
    </login-config>
    ```

- Configuring the **elytron-oidc-client** subsystem:

    - To secure your application, use the following management CLI command:

    ```
    /subsystem=elytron-oidc-client/secure-deployment=simple-oidc-example.war/:add(client-id=my_jbeap,provider-url=http://localhost:8180/auth/realms/example_realm,public-client=true,ssl-required=external)
    ```

6. In the application root directory, compile your application with the following command:

```
$ mvn package
```

7. Deploy the application.

```
$ mvn wildfly:deploy
```

Verification

1. In a browser, navigate to **http://localhost:8080/simple-oidc-example/secured**.

2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Secured Servlet
Current Principal '5cb0c4ca-0477-44c3-bdef-04db04d7e39d'
```

You can now log in to the application using the credentials you configured in the Red Hat Single Sign-On as the OpenID provider.

Additional resources

- OpenID Connect configuration in JBoss EAP
- Restricting access to applications based on user roles

9.3.4. Restricting access to applications based on user roles

You can restrict access to all, or parts, of your application based on user roles. For example, you can let users with the "public" role have access to the parts of your application that aren't sensitive, and give users with the "admin" role access to those parts that are.

Prerequisites

- You have secured your application using OpenID Connect. For more information, see Creating a secure application that uses OpenID Connect.

Procedure

1. Update the deployment descriptor **web.xml** file with the following text:

Syntax

```
<security-constraint>
    ...
    <auth-constraint>
        <role-name><allowed_role></role-name>
    </auth-constraint>
</security-constraint>
```

Example

```
<security-constraint>
    ...
    <auth-constraint>
        <role-name>example_role</role-name> 1
    </auth-constraint>
</security-constraint>
```

1. Allow only those users with the role **example_role** to access your application.

2. In the application root directory, recompile your application with the following command:

```
$ mvn package
```

3. Deploy the application.

```
$ mvn wildfly:deploy
```

Verification

1. In a browser, navigate to **http://localhost:8080/simple-oidc-example/secured**.

2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Forbidden
```

Because you have not assigned the required role to the user "jane_doe," jane_doe can't log in to your application. Only the users with the required role can log in.

To assign users the required roles, see Creating and assigning roles to users in Red Hat Single Sign-On .

9.3.5. Creating and assigning user roles in Red Hat Single Sign-On

Red Hat Single Sign-On is an identity and access management provider for securing your web applications with single sign-on (SSO). You can define users and assign roles in Red Hat Single Sign-On.

Prerequisites

- You have configured Red Hat Single Sign-On. For more information, see Configuring Red Hat Single Sign-On as an OpenID provider.

Procedure

1. Log in to the admin console at **http://localhost:<port>/auth/**. For example, **http://localhost:8180/auth/**.

2. Click the realm you use to connect with JBoss EAP. For example, *example_realm*.

3. Click **Clients**, then click the **client-name** you configured for JBoss EAP. For example, *my_jbeap*.

4. Click **Roles**, then **Add Role**.

5. Enter a role name, such as *example_role*, then click **Save**. This is the role name you configure in JBoss EAP for authorization.

6. Click **Users**, then **View all users**.

7. Click an ID to assign the role you created. For example, click the ID for *jane_doe*.

8. Click **Role Mappings**. In the **Client Roles** field, select the **client-name** you configured for JBoss EAP. For example, *my_jbeap*.

9. In **Available Roles**, select a role to assign. For example, *example_role*. Click **Add selected**.

Verification

1. In a browser, navigate to the application URL.

2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Secured Servlet
Current Principal '5cb0c4ca-0477-44c3-bdef-04db04d7e39d'
```

Users with the required role can log in to your application.

Additional resources

- Assigning permissions and access using roles and groups in Red Hat Single Sign-On

9.4. DEVELOPING JBOSS EAP BOOTABLE JAR APPLICATION WITH OPENID CONNECT

You can use OpenID Connect (OIDC) to delegate authentication to an external OpenID provider. The **elytron-oidc-client** galleon layer provides a native OIDC client in JBoss EAP bootable jar applications to connect with external OpenID providers.

To create an application secured with OpenID Connect using Red Hat Single Sign-On, follow these procedures:

- Configure Red Hat Single Sign-On as the OpenID provider
- Create a Maven project for your application
- Create a bootable jar application that uses OpenID Connect
- Restrict access to your application based on user roles
- Create and assign user roles in Red Hat Single Sign-On

9.4.1. Configuring Red Hat Single Sign-On as an OpenID provider

Red Hat Single Sign-On is an identity and access management provider for securing web applications with single sign-on (SSO). It supports OpenID Connect (an extension to OAuth 2.0).

Prerequisites

- You have installed the Red Hat Single Sign-On server. For more information, see Installing the Red Hat Single Sign-On server in the Red Hat Single Sign-On *Getting Started Guide*.
- You have created a user in your Red Hat Single Sign-On server instance. For more information, see Creating a user in the Red Hat Single Sign-On *Getting Started Guide*.

Procedure

1. Start the Red Hat Single Sign-On server at a port other than 8080 because JBoss EAP default port is 8080.

   Syntax

   ```
   $ RH_SSO_HOME/bin/standalone.sh -Djboss.socket.binding.port-offset=<offset-number>
   ```

   Example

   ```
   $ /home/servers/rh-sso-7.4/bin/standalone.sh -Djboss.socket.binding.port-offset=100
   ```

2. Log in to the Admin Console at **http://localhost:<port>/auth/**. For example, **http://localhost:8180/auth/**.
3. To create a realm, in the Admin Console, hover over **Master**, and click **Add realm**.
4. Enter a name for the realm. For example, **example_realm**. Ensure that **Enabled** is **ON** and click **Create**.
5. Click **Users**, then click **Add user** to add a user to the realm.
6. Enter a user name. For example, **jane_doe**. Ensure that User **Enabled** is **ON** and click **Save**.
7. Click **Credentials** to add a password to the user.
8. Set a password for the user. For example, **janedoep@$$**. Toggle **Temporary** to **OFF** and click **Set Password**. In the confirmation prompt, click **Set password**.
9. Click **Clients**, then click **Create** to configure a client connection.
10. Enter a client ID. For example, **my_jbeap**. Ensure that **Client Protocol** is set to **openid-connect**, and click **Save**.
11. Click **Installation**, then select **Keycloak OIDC JSON** as the **Format Option** to see the connection parameters.

    ```
    {
      "realm": "example_realm",
      "auth-server-url": "http://localhost:8180/auth/",
      "ssl-required": "external",
    ```

```
  "resource": "my_jbeap",
  "public-client": true,
  "confidential-port": 0
}
```

When configuring your JBoss EAP application to use Red Hat Single Sign-On as the identity provider, you use the parameters as follows:

```
"provider-url" : "http://localhost:8180/auth/realms/example_realm",
"ssl-required": "external",
"client-id": "my_jbeap",
"public-client": true,
"confidential-port": 0
```

12. Click **Clients**, click **Edit** next to **my_jbeap** to edit the client settings.

13. In **Valid Redirect URIs**, enter the URL where the page should redirect after authentication is successful.
    For this example, set this value to **http://localhost:8080/simple-oidc-layer-example/secured/***

Additional resources

- Configuring a Maven project for creating a secure application
- Creating a realm and a user

9.4.2. Configuring a Maven project for a bootable jar OIDC application

Create a Maven project with the required dependencies and the directory structure for creating a bootable jar application that uses OpenID Connect. The **elytron-oidc-client** galleon layer provides a native OpenID Connect (OIDC) client to connect with OpenID providers.

Prerequisites

- You have installed Maven. For more information, see Downloading Apache Maven.
- You have configured your Maven repository for the latest release. For more information, see Maven and the JBoss EAP microprofile Maven repository .

Procedure

1. Set up a Maven project using the **mvn** command. The command creates the directory structure for the project and the **pom.xml** configuration file.

   #### Syntax

   ```
   $ mvn archetype:generate \
   -DgroupId=${group-to-which-your-application-belongs} \
   -DartifactId=${name-of-your-application} \
   -DarchetypeGroupId=org.apache.maven.archetypes \
   -DarchetypeArtifactId=maven-archetype-webapp \
   -DinteractiveMode=false
   ```

   #### Example

```
$ mvn archetype:generate \
-DgroupId=com.example.oidc \
-DartifactId=simple-oidc-layer-example \
-DarchetypeGroupId=org.apache.maven.archetypes \
-DarchetypeArtifactId=maven-archetype-webapp \
-DinteractiveMode=false
```

2. Navigate to the application root directory.

Syntax

```
$ cd <name-of-your-application>
```

Example

```
$ cd simple-oidc-layer-example
```

3. Update the generated **pom.xml** file as follows:

a. Set the following repositories:

```
<repositories>
    <repository>
        <id>jboss</id>
        <url>https://maven.repository.redhat.com/ga</url>
        <snapshots>
            <enabled>false</enabled>
        </snapshots>
    </repository>
</repositories>
```

b. Set the following plugin repositories:

```
<pluginRepositories>
    <pluginRepository>
        <id>jboss</id>
        <url>https://maven.repository.redhat.com/ga</url>
        <snapshots>
            <enabled>false</enabled>
        </snapshots>
    </pluginRepository>
</pluginRepositories>
```

c. Set the following properties:

```
<properties>
    <maven.compiler.source>1.8</maven.compiler.source>
    <maven.compiler.target>1.8</maven.compiler.target>
    <bootable.jar.maven.plugin.version>6.1.2.Final-redhat-
00001</bootable.jar.maven.plugin.version>
    <jboss.xp.galleon.feature.pack.version>4.0.0.GA-redhat-
00002</jboss.xp.galleon.feature.pack.version>
</properties>
```

d. Set the following dependencies:

```
<dependencies>
   <dependency>
     <groupId>javax.servlet</groupId>
     <artifactId>javax.servlet-api</artifactId>
     <version>3.1.0.redhat-1</version>
     <scope>provided</scope>
   </dependency>
</dependencies>

<dependencyManagement>
   <dependencies>
      <dependency>
         <groupId>org.jboss.bom</groupId>
         <artifactId>jboss-eap-jakartaee8</artifactId>
         <version>7.3.4.GA</version>
         <type>pom</type>
         <scope>import</scope>
      </dependency>
      <dependency>
         <groupId>org.jboss.spec.javax.servlet</groupId>
         <artifactId>jboss-servlet-api_4.0_spec</artifactId>
         <scope>provided</scope>
      </dependency>
   </dependencies>
</dependencyManagement>
```

e. Set the following build configuration in the **<build>** element of the **pom.xml** file:

```
<finalName>${project.artifactId}</finalName>
<plugins>
   <plugin>
      <groupId>org.wildfly.plugins</groupId>
      <artifactId>wildfly-jar-maven-plugin</artifactId>  1
      <version>${bootable.jar.maven.plugin.version}</version>
      <configuration>
         <feature-pack-location>org.jboss.eap:wildfly-galleon-
pack:${jboss.xp.galleon.feature.pack.version}</feature-pack-location>
         <layers>
            <layer>jaxrs-server</layer>
            <layer>elytron-oidc-client</layer>  2
         </layers>
         <context-root>false</context-root>  3
      </configuration>
      <executions>
         <execution>
            <goals>
               <goal>package</goal>
            </goals>
         </execution>
      </executions>
   </plugin>
</plugins>
```

1. JBoss EAP Maven plug-in to build the application as a bootable JAR
2. The **elytron-oidc-client** layer provides a native OpenID Connect (OIDC) client to connect with external OpenID providers.
3. Register the application in the **simple-oidc-layer-example** resource path. The servlet is then available at the URL **http://*server-url*/*application_name*/*servlet_path***, for example: **http://localhost:8080/simple-oidc-layer-example/secured**. By default, the application WAR file is registered under the root-context path, like **http://*server-url*/*servlet_path***, for example: **http://localhost:8080/secured**.

f. Set the application name, for example "simple-oidc-layer-example" in the **<build>** element of the **pom.xml** file.

```
<finalName>simple-oidc-layer-example</finalName>
```

Verification

- In the application root directory, enter the following command:

```
$ mvn install
```

You get an output similar to the following:

```
...
[INFO] ------------------------------------------------------------------------
[INFO] BUILD SUCCESS
[INFO] ------------------------------------------------------------------------
[INFO] Total time: 19.157 s
[INFO] Finished at: 2022-03-10T09:38:21+05:30
[INFO] ------------------------------------------------------------------------
```

You can now create a bootable jar application that uses OpenID Connect

9.4.3. Creating a bootable jar application that uses OpenID Connect

The following example demonstrates creating a servlet that prints a logged-in user's Principal. For an existing application, only those steps that are related to updating the deployment configuration are required.

In this example, the value of the Principal comes from the ID token from the OpenID provider. By default, the Principal is the value of the **"sub"** claim from the token. You can specify which claim value from the ID token to use as the Principal in one of the following:

- The **elytron-oidc-client** subsystem attribute **principal-attribute**.
- The **oidc.json** file.

<application_root> in the procedure denotes the **pom.xml** file directory. The **pom.xml** file contains your application's Maven configuration.

Prerequisites

- You have created a Maven project. For more information, see Configuring Maven project for creating a secure application .
- You have configured Red Hat Single Sign-On as the OpenID provider. For more information, see Configuring Red Hat Single Sign-On as an OpenID provider .

Procedure

1. Create a directory to store the Java files.

   Syntax

   ```
   $ mkdir -p <application_root>/src/main/java/com/example/oidc
   ```

   Example

   ```
   $ mkdir -p simple-oidc-layer-example/src/main/java/com/example/oidc
   ```

2. Navigate to the new directory.

   Syntax

   ```
   $ cd <application_root>/src/main/java/com/example/oidc
   ```

   Example

   ```
   $ cd simple-oidc-layer-example/src/main/java/com/example/oidc
   ```

3. Create a servlet "SecuredServlet.java" with the following content:

   ```
   package com.example.oidc;

   import java.io.IOException;
   import java.io.PrintWriter;
   import java.security.Principal;

   import javax.servlet.ServletException;
   import javax.servlet.annotation.WebServlet;
   import javax.servlet.http.HttpServlet;
   import javax.servlet.http.HttpServletRequest;
   import javax.servlet.http.HttpServletResponse;

   /**
    * A simple secured HTTP servlet.
    *
    */
   @WebServlet("/secured")
   public class SecuredServlet extends HttpServlet {

       @Override
       protected void doGet(HttpServletRequest req, HttpServletResponse resp) throws
   ServletException, IOException {
           try (PrintWriter writer = resp.getWriter()) {
               writer.println("<html>");
   ```

```
            writer.println("  <head><title>Secured Servlet</title></head>");
            writer.println("  <body>");
            writer.println("    <h1>Secured Servlet</h1>");
            writer.println("    <p>");
            writer.print(" Current Principal '");
            Principal user = req.getUserPrincipal();
            writer.print(user != null ? user.getName() : "NO AUTHENTICATED USER");
            writer.print("'");
            writer.println("    </p>");
            writer.println("  </body>");
            writer.println("</html>");
        }
    }

}
```

4. Add security rules for access to your application in the deployment descriptor **web.xml** file located in the **WEB-INF** directory of the application.

```
<?xml version="1.0" encoding="UTF-8"?>

<web-app version="2.5" xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd"
  metadata-complete="false">

   <security-constraint>
      <web-resource-collection>
         <web-resource-name>secured</web-resource-name>
         <url-pattern>/secured</url-pattern>
      </web-resource-collection>
      <auth-constraint>
         <role-name>*</role-name>
      </auth-constraint>
   </security-constraint>

   <security-role>
      <role-name>*</role-name>
   </security-role>
</web-app>
```

5. To secure the application with OpenID Connect, either update the deployment configuration or configure the **elytron-oidc-client** subsystem.

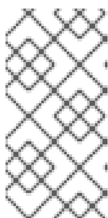

NOTE

If you configure OpenID Connect in both the deployment configuration and the **elytron-oidc-client** subsystem, the configuration in the **elytron-oidc-client** subsystem **secure-deployment** attribute takes precedence over the configuration in the application deployment descriptor.

- Updating the deployment configuration:

    i. Create a file **oidc.json** in the **WEB-INF** directory, like this:

```
{
  "provider-url" : "http://localhost:8180/auth/realms/example_realm",
  "ssl-required": "external",
  "client-id": "my_jbeap",
  "public-client": true,
  "confidential-port": 0
}
```

ii. Update the deployment descriptor **web.xml** file with the following text to declare that this application uses OIDC:

```
<login-config>
    <auth-method>OIDC</auth-method>
</login-config>
```

- Configuring the **elytron-oidc-client** subsystem:

i. Create a directory to store a CLI script in the application root directory:

Syntax

```
$ mkdir <application_root>/<cli_script_directory>
```

Example

```
$ mkdir simple-oidc-layer-example/scripts/
```

You can create the directory at any place that Maven can access, inside the application root directory.

ii. Create a CLI script, such as **configure-oidc.cli**, with the following content:

```
/subsystem=elytron-oidc-client/secure-deployment=simple-oidc-layer-
example.war:add(client-id=my_jbeap,provider-
url=http://localhost:8180/auth/realms/example_realm,public-client=true,ssl-
required=external)
```

The subsystem command defines the **simple-oidc-layer-example.war** resource as the deployment to secure in **elytron-oidc-client** subsystem.

iii. In the project **pom.xml** file, add the following configuration extract to the existing plug-in **<configuration>** element:

```
<cli-sessions>
    <cli-session>
        <script-files>
            <script>scripts/configure-oidc.cli</script>
        </script-files>
    </cli-session>
</cli-sessions>
```

6. In the application root directory, compile your application with the following command:

```
$ mvn package
```

7. Deploy the bootable jar application using the following command:

Syntax

```
$ java -jar <application_root>/target/simple-oidc-layer-example-bootable.jar
```

Example

```
$ java -jar simple-oidc-layer-example/target/simple-oidc-layer-example-bootable.jar
```

This starts JBoss EAP and deploys the application.

Verification

1. In a browser, navigate to **http://localhost:8080/simple-oidc-layer-example/secured**.

2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Secured Servlet
Current Principal '5cb0c4ca-0477-44c3-bdef-04db04d7e39d'
```

You can now log in to the application using the credentials you configured in the Red Hat Single Sign-On as the OpenID provider.

Additional resources

- OpenID Connect configuration in JBoss EAP
- Restricting access to applications based on user roles

9.4.4. Restricting access based on user roles in bootable jar OIDC applications

You can restrict access to all, or parts, of your application based on user roles. For example, you can let users with the "public" role have access to the parts of your application that aren't sensitive, and give users with the "admin" role access to those parts that are.

Prerequisites

- You have secured your application using OpenID Connect. For more information, see Creating a bootable jar application that uses OpenID Connect.

Procedure

1. Update the deployment descriptor **web.xml** file with the following text:

Syntax

```
<security-constraint>
    ...
    <auth-constraint>
        <role-name><allowed_role></role-name>
    </auth-constraint>
</security-constraint>
```

Example

```
<security-constraint>
    ...
    <auth-constraint>
        <role-name>example_role</role-name> (1)
    </auth-constraint>
</security-constraint>
```

(1) Allow only those users with the role **example_role** to access your application.

2. In the application root directory, recompile your application with the following command:

```
$ mvn package
```

3. Deploy the application.

```
$ java -jar simple-oidc-layer-example/target/simple-oidc-layer-example-bootable.jar
```

This starts JBoss EAP and deploys the application.

Verification

1. In a browser, navigate to **\localhost:8080/simple-oidc-layer-example/secured**.

2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Forbidden
```

Because you have not assigned the required role to the user "jane_doe," jane_doe can't log in to your application. Only the users with the required role can log in.

To assign users the required roles, see Creating and assigning roles to users in Red Hat Single Sign-On .

9.4.5. Creating and assigning user roles in Red Hat Single Sign-On

Red Hat Single Sign-On is an identity and access management provider for securing your web applications with single sign-on (SSO). You can define users and assign roles in Red Hat Single Sign-On.

Prerequisites

- You have configured Red Hat Single Sign-On. For more information, see Configuring Red Hat Single Sign-On as an OpenID provider.

Procedure

1. Log in to the admin console at **http://localhost:<port>/auth/**. For example, **http://localhost:8180/auth/**.
2. Click the realm you use to connect with JBoss EAP. For example, *example_realm*.
3. Click **Clients**, then click the **client-name** you configured for JBoss EAP. For example, *my_jbeap*.
4. Click **Roles**, then **Add Role**.
5. Enter a role name, such as *example_role*, then click **Save**. This is the role name you configure in JBoss EAP for authorization.
6. Click **Users**, then **View all users**.
7. Click an ID to assign the role you created. For example, click the ID for *jane_doe*.
8. Click **Role Mappings**. In the **Client Roles** field, select the **client-name** you configured for JBoss EAP. For example, *my_jbeap*.
9. In **Available Roles**, select a role to assign. For example, *example_role*. Click **Add selected**.

Verification

1. In a browser, navigate to the application URL.
2. Log in with your credentials. For example:

```
username: jane_doe
password: janedoep@$$
```

You get the following output:

```
Secured Servlet
Current Principal '5cb0c4ca-0477-44c3-bdef-04db04d7e39d'
```

Users with the required role can log in to your application.

Additional resources

- Assigning permissions and access using roles and groups in Red Hat Single Sign-On

CHAPTER 10. OBSERVABILITY IN JBOSS EAP

If you're a developer or system administrator, *observability* is a set of practices and technologies you can use to determine, based on certain signals from your application, the location and source of a problem in your application. The most common signals are metrics, events, and tracing. JBoss EAP uses OpenTelemetry for *observability*.

10.1. OPENTELEMETRY IN JBOSS EAP

OpenTelemetry is a set of tools, application programming interfaces (APIs), and software development kits (SDKs) you can use to instrument, generate, collect, and export telemetry data for your applications. Telemetry data includes metrics, logs, and traces. Analyzing an application's telemetry data helps you to improve your application's performance. JBoss EAP provides OpenTelemetry capability through the **opentelemetry** subsystem.

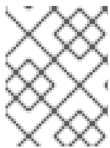

NOTE

Red Hat JBoss Enterprise Application Platform 7.4 provides only OpenTelemetry tracing capabilities.

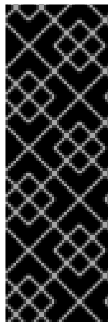

IMPORTANT

OpenTelemetry is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see https://access.redhat.com/support/offerings/techpreview .

Additional resources

- OpenTelemetry Documentation

10.2. OPENTELEMETRY CONFIGURATION IN JBOSS EAP

You configure a number of aspects of OpenTelemetry in JBoss EAP using the **opentelemetry** subsystem. These include exporter, span processor, and sampler.

exporter

To analyze and visualize traces and metrics, you export them to a collector such as Jaeger. You can configure JBoss EAP to use either Jaeger or any collector that supports the OpenTelemetry protocol (OTLP).

span processor

You can configure the span processor to export spans either as they are produced or in batches. You can also configure the number of traces to export.

sampler

You can configure the number of traces to record by configuring the sampler.

Example configuration

The following XML is an example of the full OpenTelemetry configuration, including default values. JBoss EAP does not persist the default values when you make changes, so your configuration might look different.

```
<subsystem xmlns="urn:wildfly:opentelemetry:1.0"
        service-name="example">
    <exporter
        type="jaeger"
        endpoint="http://localhost:14250"/>
    <span-processor
        type="batch"
        batch-delay="4500"
        max-queue-size="128"
        max-export-batch-size="512"
        export-timeout="45"/>
    <sampler
        type="on"/>
</subsystem>
```

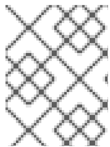

NOTE

You cannot use an OpenShift route object to connect with a Jaeger endpoint. Instead, use **http://*<ip_address>*:*<port>*** or **http://*<service_name>*:*<port>***.

Additional resources

- OpenTelemetry subsystem attributes

10.3. OPENTELEMETRY TRACING IN JBOSS EAP

JBoss EAP provides OpenTelemetry tracing to help you track the progress of user requests as they pass through different parts of your application. By analyzing traces, you can improve your application's performance and debug availability issues.

OpenTelemetry tracing consists of the following components:

Trace
: A collection of operations that a request goes through in an application.

Span
: A single operation within a trace. It provides request, error, and duration (RED) metrics and contains a span context.

Span context
: A set of unique identifiers that represents a request that the containing span is a part of.

JBoss EAP automatically traces REST calls to your Jakarta RESTful Web Services applications and container-managed Jakarta RESTful Web Services client invocations. JBoss EAP traces REST calls implicitly as follows:

- For each incoming request:
  - JBoss EAP extracts the span context from the request.
  - JBoss EAP starts a new span, then closes it when the request is completed.

- For each outgoing request:
  - JBoss EAP injects span context into the request.
  - JBoss EAP starts a new span, then closes it when the request is completed.

In addition to implicit tracing, you can create custom spans by injecting a **Tracer** instance into your application for granular tracing.

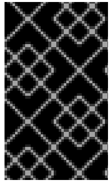

IMPORTANT

If you see duplicate traces exported for REST calls, disable the **microprofile-opentracing-smallrye** subsystem. For information about disabling the **microprofile-opentracing-smallrye**, see Removing the microprofile-opentracing-smallrye subsystem.

Additional resources

- Using Jaeger to observe the OpenTelemetry traces for an application
- OpenTelemetry application development in JBoss EAP

10.4. ENABLING OPENTELEMETRY TRACING IN JBOSS EAP

To use OpenTelemetry tracing in JBoss EAP you must first enable the **opentelemetry** subsystem.

Prerequisites

- You have installed JBoss EAP XP.

Procedure

1. Add the OpenTelemetry extension using the management CLI.

```
/extension=org.wildfly.extension.opentelemetry:add
```

2. Enable the **opentelemetry** subsystem using the management CLI.

```
/subsystem=opentelemetry:add
```

3. Reload JBoss EAP.

```
reload
```

Additional resources

- Configuring the **opentelemetry** subsystem

10.5. CONFIGURING THE OPENTELEMETRY SUBSYSTEM

You can configure the **opentelemetry** subsystem to set different aspects of tracing. Configure these based on the collector you use for observing the traces.

Prerequisites

- You have enabled the **opentelemetry** subsystem. For more information, see Enabling OpenTelemetry tracing in JBoss EAP.

Procedure

1. Set the exporter type for the traces.

   Syntax

   ```
   /subsystem=opentelemetry:write-attribute(name=exporter-type, value=<exporter_type>)
   ```

   Example

   ```
   /subsystem=opentelemetry:write-attribute(name=exporter-type, value=jaeger)
   ```

2. Set the endpoint at which to export the traces.

   Syntax

   ```
   /subsystem=opentelemetry:write-attribute(name=endpoint, value=<URL:port>)
   ```

   Example

   ```
   /subsystem=opentelemetry:write-attribute(name=endpoint, value=http:localhost:14250)
   ```

3. Set the service name under which the traces are exported.

   Syntax

   ```
   /subsystem=opentelemetry:write-attribute(name=service-name, value=<service_name>)
   ```

   Example

   ```
   /subsystem=opentelemetry:write-attribute(name=service-name,
   value=exampleOpenTelemetryService)
   ```

Additional resources

- Using Jaeger to observe the OpenTelemetry traces for an application

10.6. USING JAEGER TO OBSERVE THE OPENTELEMETRY TRACES FOR AN APPLICATION

JBoss EAP automatically and implicitly traces REST calls to Jakarta RESTful Web Services applications. You do not need to add any configuration to your Jakarta RESTful Web Services application or configure the **opentelemetry** subsystem. The following procedure demonstrates how to observe traces for the **helloworld-rs** quickstart in the Jaeger console.

Prerequisites

- You have installed Docker. For more information, see Get Docker.
- You have downloaded the **helloworld-rs** quickstart. The quickstart is available at helloworld-rs.
- You have configured the the **opentelemetry** subsystem. For more information, see Configuring the **opentelemetry** subsystem.

Procedure

1. Start the Jaeger console using its Docker image.

```
$ docker run -d --name jaeger \
  -e COLLECTOR_ZIPKIN_HOST_PORT=:9411 \
  -p 5775:5775/udp \
  -p 6831:6831/udp \
  -p 6832:6832/udp \
  -p 5778:5778 \
  -p 16686:16686 \
  -p 14268:14268 \
  -p 14250:14250 \
  -p 9411:9411 \
  jaegertracing/all-in-one:1.29
```

2. Use Maven to deploy the **helloworld-rs** quickstart from its root directory.

```
$ mvn clean install wildfly:deploy
```

3. In a web browser, access the quickstart at **http://localhost:8080/helloworld-rs/**, then click any link.
4. In a web browser, open the Jaeger console at **http://localhost:16686/search**. **hello-world.rs** is listed under **Service**.
5. Select **hello-world.rs** and click **Find Traces**. The details of the trace for **hello-world.rs** are listed.

Additional resources

- OpenTelemetry application development in JBoss EAP

10.7. OPENTELEMETRY TRACING APPLICATION DEVELOPMENT

Although JBoss EAP automatically and implicitly traces REST calls to Jakarta RESTful Web Services applications, you can create custom spans from your application for granular tracing. A *span* is a single operation within a trace. You can create a span when, for example, a resource is defined, a method is called, and so on, in your application. You create custom traces in your application by injecting a **Tracer** instance.

10.7.1. Configuring a Maven project for OpenTelemetry tracing

For creating an OpenTelemetry tracing application, create a Maven project with the required dependencies and directory structure.

Prerequisites

- You have installed Maven. For more information, see Downloading Apache Maven.
- You have configured your Maven repository for the latest release. For information about installing the latest Maven repository patch, see Maven and the JBoss EAP microprofile maven repository.

Procedure

1. In the CLI, use the **mvn** command to set up a Maven project. This command creates the directory structure for the project and the **pom.xml** configuration file.

   #### Syntax

   ```
   $ mvn archetype:generate \
   -DgroupId=<group-to-which-your-application-belongs> \
   -DartifactId=<name-of-your-application> \
   -DarchetypeGroupId=org.apache.maven.archetypes \
   -DarchetypeArtifactId=maven-archetype-webapp \
   -DinteractiveMode=false
   ```

   #### Example

   ```
   $ mvn archetype:generate \
   -DgroupId=com.example.opentelemetry \
   -DartifactId=simple-tracing-example \
   -DarchetypeGroupId=org.apache.maven.archetypes \
   -DarchetypeArtifactId=maven-archetype-webapp \
   -DinteractiveMode=false
   ```

2. Navigate to the application root directory.

   #### Syntax

   ```
   $ cd <name-of-your-application>
   ```

   #### Example

   ```
   $ cd simple-tracing-example
   ```

3. Update the generated **pom.xml** file.

   a. Set the following properties:

   ```
   <properties>
       <maven.compiler.source>1.8</maven.compiler.source>
       <maven.compiler.target>1.8</maven.compiler.target>
       <failOnMissingWebXml>false</failOnMissingWebXml>
       <version.server.bom>4.0.0.GA</version.server.bom>
       <version.wildfly-jar.maven.plugin>6.1.1.Final</version.wildfly-jar.maven.plugin>
   </properties>
   ```

   b. Set the following dependencies:

```
<dependencies>
    <dependency>
        <groupId>jakarta.enterprise</groupId>
        <artifactId>jakarta.enterprise.cdi-api</artifactId>
        <version>2.0.2</version>
        <scope>provided</scope>
    </dependency>

    <dependency>
        <groupId>org.jboss.spec.javax.ws.rs</groupId>
        <artifactId>jboss-jaxrs-api_2.1_spec</artifactId>
        <version>2.0.2.Final</version>
        <scope>provided</scope>
    </dependency>

    <dependency>
        <groupId>io.opentelemetry</groupId>
        <artifactId>opentelemetry-api</artifactId>
        <version>1.5.0</version>
        <scope>provided</scope>
    </dependency>
</dependencies>
```

c. Set the following build configuration to use **mvn widlfy:deploy** to deploy the application:

```
<build>
  <!-- Set the name of the archive -->
  <finalName>${project.artifactId}</finalName>
  <plugins>
    <!-- Allows to use mvn wildfly:deploy -->
    <plugin>
      <groupId>org.wildfly.plugins</groupId>
      <artifactId>wildfly-maven-plugin</artifactId>
    </plugin>
  </plugins>
</build>
```

Verification

- In the application root directory, enter the following command:

```
$ mvn install
```

You get an output similar to the following:

```
[INFO] ------------------------------------------------------------------------
[INFO] BUILD SUCCESS
[INFO] ------------------------------------------------------------------------
[INFO] Total time: 1.440 s
[INFO] Finished at: 2021-12-27T14:45:12+05:30
[INFO] ------------------------------------------------------------------------
```

You can now create an OpenTelemetry tracing application.

Additional resources

- Creating applications that create custom spans

10.7.2. Creating applications that create custom spans

The following procedure demonstrates how to create an application that can create two custom spans like these:

- **prepare-hello** - When the method **getHello()** in the application is called.
- **process-hello** - When the value **hello** is assigned to a new **String** object **hello**.

This procedure also demonstrates how to view these spans in a Jaeger console. **<application_root>** in the procedure denotes the directory that contains the **pom.xml** file, which contains the Maven configuration for your application.

Prerequisites

- You have installed Docker. For more information, see Get Docker.
- You have created a Maven project. For more information, see Configuring Maven project for OpenTelemetry tracing.
- You have configured the the **opentelemetry** subsystem. For more information, see Configuring the **opentelemetry** subsystem.

Procedure

1. In the *<application_root>*, create a directory to store the Java files.

   Syntax

   ```
   $ mkdir -p src/main/java/com/example/opentelemetry
   ```

   Example

   ```
   $ mkdir -p src/main/java/com/example/opentelemetry
   ```

2. Navigate to the new directory.

   Syntax

   ```
   $ cd src/main/java/com/example/opentelemetry
   ```

   Example

   ```
   $ cd src/main/java/com/example/opentelemetry
   ```

3. Create a **JakartaRestApplication.java** file with the following content. This **JakartaRestApplication** class declares the application as a Jakarta RESTful Web Services application.

```
package com.example.opentelemetry;

import javax.ws.rs.ApplicationPath;
import javax.ws.rs.core.Application;

@ApplicationPath("/")
public class JakartaRestApplication extends Application {
}
```

4. Create an **ExplicitlyTracedBean.java** file with the following content for the class **ExplicitlyTracedBean**. This class creates custom spans by injecting a **Tracer** class.

```
package com.example.opentelemetry;

import javax.enterprise.context.RequestScoped;
import javax.inject.Inject;
import io.opentelemetry.api.trace.Span;
import io.opentelemetry.api.trace.Tracer;

@RequestScoped
public class ExplicitlyTracedBean {

    @Inject
    private Tracer tracer; (1)

    public String getHello() {
        Span prepareHelloSpan = tracer.spanBuilder("prepare-hello").startSpan(); (2)
        prepareHelloSpan.makeCurrent();

        String hello = "hello";

        Span processHelloSpan = tracer.spanBuilder("process-hello").startSpan(); (3)
        processHelloSpan.makeCurrent();

        hello = hello.toUpperCase();

        processHelloSpan.end();
        prepareHelloSpan.end();

        return hello;
    }
}
```

(1) Inject a **Tracer** class to create custom spans.

(2) Create a span called **prepare-hello** to indicate that the method **getHello()** was called.

(3) Create a span called **process-hello** to indicate that the value **hello** was assigned to a new **String** object called **hello**.

5. Create a **TracedResource.java** file with the following content for **TracedResource** class. This file injects the **ExplicitlyTracedBean** class and declares two endpoints: **traced** and **cdi-trace**.

```
package com.example.opentelemetry;
```

```
import javax.enterprise.context.RequestScoped;
import javax.inject.Inject;
import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import javax.ws.rs.core.MediaType;

@Path("/hello")
@RequestScoped
public class TracedResource {
    @Inject
    private ExplicitlyTracedBean tracedBean;

    @GET
    @Path("/traced")
    @Produces(MediaType.TEXT_PLAIN)
    public String hello() {
        return "hello";
    }

    @GET
    @Path("/cdi-trace")
    @Produces(MediaType.TEXT_PLAIN)
    public String cdiHello() {
        return tracedBean.getHello();
    }
}
```

6. Navigate to the application root directory.

Syntax

```
$ cd <path_to_application_root>/<application_root>
```

Example

```
$ cd ~/applications/simple-tracing-example
```

7. Compile and deploy the application with the following command:

```
$ mvn clean package wildfly:deploy
```

8. Start the Jaeger console.

```
$ docker run -d --name jaeger \
  -e COLLECTOR_ZIPKIN_HOST_PORT=:9411 \
  -p 5775:5775/udp \
  -p 6831:6831/udp \
  -p 6832:6832/udp \
  -p 5778:5778 \
  -p 16686:16686 \
  -p 14268:14268 \
```

```
-p 14250:14250 \
-p 9411:9411 \
jaegertracing/all-in-one:1.29
```

9. In a browser, navigate to **\localhost:8080/simple-tracing-example/hello/cdi-trace**.

10. In a browser, open the Jaeger console at **http://localhost:16686/search**.

11. In the Jaeger console, select **JBoss EAP XP** and click **Find Traces**.

12. Click **3 Spans**.

13. The Jaeger console displays the following traces:

```
|GET /hello/cdi-trace (1)
-
 | prepare-hello (2)
  -
   | process-hello (3)
```

(1) This is the span for the automatic implicit trace.

(2) The custom span **prepare-hello** indicates that the method **getHello()** was called. It is the child of span for the automatic implicit trace.

(3) The custom span **process-hello** indicates that the value **hello** was assigned to a new **String** object **hello**. It is the child of the **prepare-hello** span.

Whenever you access the application endpoint at **http://localhost:16686/search**, a new trace is created with all the child spans.

Additional resources

- OpenTelemetry tracing in JBoss EAP

CHAPTER 11. REFERENCE

11.1. MICROPROFILE CONFIG REFERENCE

11.1.1. Default MicroProfile Config attributes

The MicroProfile Config specification defines three **ConfigSources** by default.

ConfigSources are sorted according to their ordinal number. If a configuration must be overwritten for a later deployment, the lower ordinal **ConfigSource** is overwritten before a higher ordinal **ConfigSource**.

Table 11.1. Default MicroProfile Config attributes

ConfigSource	Ordinal
System properties	**400**
Environment variables	**300**
Property files **META-INF/microprofile-config.properties** found on the classpath	**100**

11.1.2. MicroProfile Config SmallRye ConfigSources

The **microprofile-config-smallrye** project defines more **ConfigSources** you can use in addition to the default MicroProfile Config **ConfigSources**.

Table 11.2. Additional MicroProfile Config attributes

ConfigSource	Ordinal
config-source in the Subsystem	**100**
ConfigSource from the Directory	**100**
ConfigSource from Class	**100**

An explicit ordinal is not specified for these **ConfigSources**. They inherit the default ordinal value found in the MicroProfile Config specification.

11.2. MICROPROFILE FAULT TOLERANCE REFERENCE

11.2.1. MicroProfile Fault Tolerance configuration properties

SmallRye Fault Tolerance specification defines the following properties in addition to the properties defined in the MicroProfile Fault Tolerance specification.

Table 11.3. MicroProfile Fault Tolerance configuration properties

Property	Default value	Description
io.smallrye.faulttolerance.mainThreadPoolSize	**100**	Maximum number of threads in the thread pool.
io.smallrye.faulttolerance.mainThreadPoolQueueSize	**-1 (unbounded)**	Size of the queue that the thread pool should use.

11.3. MICROPROFILE JWT REFERENCE

11.3.1. MicroProfile Config JWT standard properties

The **microprofile-jwt-smallrye** subsystem supports the following MicroProfile Config standard properties.

Table 11.4. MicroProfile Config JWT standard properties

Property	Default	Description
mp.jwt.verify.publickey	NONE	String representation of the public key encoded using one of the supported formats. Do not set if you have set **mp.jwt.verify.publickey.location**.
mp.jwt.verify.publickey.location	NONE	The location of the public key, may be a relative path or URL. Do not be set if you have set **mp.jwt.verify.publickey**.
mp.jwt.verify.issuer	NONE	The expected value of any **iss** claim of any JWT token being validated.

Example **microprofile-config.properties** configuration:

```
mp.jwt.verify.publickey.location=META-INF/public.pem
mp.jwt.verify.issuer=jwt-issuer
```

11.4. MICROPROFILE OPENAPI REFERENCE

11.4.1. MicroProfile OpenAPI configuration properties

In addition to the standard MicroProfile OpenAPI configuration properties, JBoss EAP supports the following additional MicroProfile OpenAPI properties. These properties can be applied in both the global and the application scope.

Table 11.5. MicroProfile OpenAPI properties in JBoss EAP

Property	Default value	Description
mp.openapi.extensions.enabled	**true**	Enables or disables registration of an OpenAPI endpoint. When set to **false**, disables generation of OpenAPI documentation. You can set the value globally using the config subsystem, or for each application in a configuration file such as **/META-INF/microprofile-config.properties**. You can parameterize this property to selectively enable or disable **microprofile-openapi-smallrye** in different environments, such as production or development. You can use this property to control which application associated with a given virtual host should generate a MicroProfile OpenAPI model.
mp.openapi.extensions.path	**/openapi**	You can use this property for generating OpenAPI documentation for multiple applications associated with a virtual host. Set a distinct **mp.openapi.extensions.path** on each application associated with the same virtual host.

Property	Default value	Description
mp.openapi.extensions.servers.relative	**true**	Indicates whether auto-generated server records are absolute or relative to the location of the OpenAPI endpoint. Server records are necessary to ensure, in the presence of a non-root context path, that consumers of an OpenAPI document can construct valid URLs to REST services relative to the host of the OpenAPI endpoint. The value **true** indicates that the server records are relative to the location of the OpenAPI endpoint. The generated record contains the context path of the deployment. When set to **false**, JBoss EAP XP generates server records including all the protocols, hosts, and ports at which the deployment is accessible.

11.5. MICROPROFILE REACTIVE MESSAGING REFERENCE

11.5.1. MicroProfile reactive messaging connectors for integrating with external messaging systems

The following is a list of reactive messaging property key prefixes required by the MicroProfile Config specification:

- **mp.messaging.incoming.[channel-name].[attribute]=[value]**
- **mp.messaging.outgoing.[channel-name].[attribute]=[value]**
- **mp.messaging.connector.[connector-name].[attribute]=[value]**

Note that **channel-name** is either the **@Incoming.value()** or the **@Outgoing.value()**. For clarification, look at this example of a pair of connector methods:

```
@Outgoing("to")
public int send() {
  int i = // Randomly generated...
  return i;
}
```

```
@Incoming("from")
public void receive(int i) {
  // Process payload
}
```

In this example, the required property prefixes are as follows:

- **mp.messaging.incoming.from**. This defines the **receive()** method.
- **mp.messaging.outgoing.to**. This defines the **send()** method.

Remember that this is an example. Because different connectors recognize different properties, the prefixes you indicate depend on the connector you want to configure.

11.5.2. Example of the data exchange between reactive messaging streams and user-initialized code

The following is an example of data exchange between reactive messaging streams and code that a user triggered through the **@Channel** and **Emitter** constructs:

```
@Path("/")
@ApplicationScoped
class MyBean {
    @Inject @Channel("my-stream")
    Emitter<String> emitter; (1)

    Publisher<String> dest;

    public MyBean() { (2)
    }

    @Inject
    public MyBean(@Channel("my-stream") Publisher<String> dest) {
        this.dest = subscribeAndAllowMultipleSubscriptions(dest);
    }

    private Publisher subscribeAndAllowMultipleSubscriptions(Publisher delegate) {
    } (3) (4) (5)

    @POST
    public PublisherBuilder<String> publish(@FormParam("value") String value) {
        return emitter.send(value);
    }

    @GET
    public Publisher poll() {
        return dest;
    }

    @PreDestroy
    public void close() { (6)

    }
}
```

In-line details:

1. Wraps the constructor-injected publisher.
2. You need this empty constructor to satisfy the Contexts and Dependency Injection (CDI) for Java specification.
3. Subscribe to the delegate.
4. Wrap the delegate in a publisher that can handle multiple subscriptions.
5. The wrapping publisher forwards data from the delegate.
6. Unsubscribe from the reactive messaging-provided publisher.

In this example, MicroProfile Reactive Messaging is listening to the **my-stream** memory stream, so messages sent through the **Emitter** are received on this injected publisher. Note, though, that the following conditions must be true for this data exchange to succeed:

1. There must be an active subscription on the channel before you call **Emitter.send()**. In this example, notice that the **subscribeAndAllowMultipleSubscriptions()** method called by the constructor ensures that there's an active subscription by the time the bean is available for user code calls.

2. You can have only one **Subscription** on the injected **Publisher**. If you want to expose the receiving publisher with a REST call, where each call to the **poll()** method results in a new subscription to the **dest** publisher, you have to implement your own publisher to broadcast data from the injected to each client.

11.5.3. The Apache Kafka user API

You can use the Apache Kafka user API to get more information about messages Kafka received, and to influence how Kafka handles messages. This API is stored in the **io/smallrye/reactive/messaging/kafka/api** package, and it consists of the following classes:

- **IncomingKafkaRecordMetadata**. This metadata contains the following information:
  - The Kafka record **key**, represented by a **Message**.
  - The Kafka **topic** and **partition** used for the **Message**, and the **offset** within those.
  - The **Message timestamp** and **timestampType**.
  - The **Message headers**. These are pieces of information that the application can attach on the producing side, and receive on the consuming side.
- **OutgoingKafkaRecordMetadata**. With this metadata, you can specify or override how Kafka handles messages. It contains the following information:
  - The **key**. which Kafka treats as the message key.
  - The **topic** you want Kafka to use.
  - The **partition**.
  - The **timestamp**, if you don't want the one that Kafka generates.

  - **headers**.

- KafkaMetadataUtil contains utility methods to write **OutgoingKafkaRecordMetadata** to a **Message**, and to read **IncomingKafkaRecordMetadata** from a **Message**.

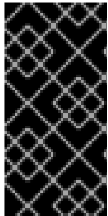

IMPORTANT

If you write **OutgoingKafkaRecordMetadata** to a **Message** sent to a channel that's not mapped to Kafka, the reactive messaging framework ignores it. Conversely, if you read **IncomingKafkaRecordMetadata** from a **Message** from a channel that's not mapped to Kafka, that message returns as **null**.

Example of how to write and read a message key

```
@Inject
@Channel("from-user")
Emitter<Integer> emitter;

@Incoming("from-user")
@Outgoing("to-kafka")
public Message<Integer> send(Message<Integer> msg) {
    // Set the key in the metadata
    OutgoingKafkaRecordMetadata<String> md =
            OutgoingKafkaRecordMetadata.<String>builder()
                .withKey("KEY-" + i)
                .build();
    // Note that Message is immutable so the copy returned by this method
    // call is not the same as the parameter to the method
    return KafkaMetadataUtil.writeOutgoingKafkaMetadata(msg, md);
}

@Incoming("from-kafka")
public CompletionStage<Void> receive(Message<Integer> msg) {
    IncomingKafkaRecordMetadata<String, Integer> metadata =
        KafkaMetadataUtil.readIncomingKafkaMetadata(msg).get();

    // We can now read the Kafka record key
    String key = metadata.getKey();

    // When using the Message wrapper around the payload we need to explicitly ack
    // them
    return msg.ack();
}
```

Example of Kafka mapping in a microprofile-config.properties file

```
kafka.bootstrap.servers=kafka:9092

mp.messaging.outgoing.to-kafka.connector=smallrye-kafka
mp.messaging.outgoing.to-kafka.topic=some-topic
mp.messaging.outgoing.to-
kafka.value.serializer=org.apache.kafka.common.serialization.IntegerSerializer
mp.messaging.outgoing.to-
kafka.key.serializer=org.apache.kafka.common.serialization.StringSerializer
```

```
mp.messaging.incoming.from-kafka.connector=smallrye-kafka
mp.messaging.incoming.from-kafka.topic=some-topic
mp.messaging.incoming.from-
kafka.value.deserializer=org.apache.kafka.common.serialization.IntegerDeserializer
mp.messaging.incoming.from-
kafka.key.deserializer=org.apache.kafka.common.serialization.StringDeserializer
```

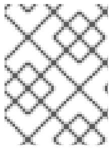

NOTE

You must specify the **key.serializer** for the outgoing channel and the **key.deserializer** for the incoming channel.

11.5.4. Example MicroProfile Config properties file for the Kafka connector

This is an example of a simple **microprofile-config.properties** file for a Kafka connector. Its properties correspond to the properties in the example in "MicroProfile reactive messaging connectors for integrating with external messaging systems."

```
kafka.bootstrap.servers=kafka:9092

mp.messaging.outgoing.to.connector=smallrye-kafka
mp.messaging.outgoing.to.topic=my-topic
mp.messaging.outgoing.to.value.serializer=org.apache.kafka.common.serialization.IntegerSerializer

mp.messaging.incoming.from.connector=smallrye-kafka
mp.messaging.incoming.from.topic=my-topic
mp.messaging.incoming.from.value.deserializer=org.apache.kafka.common.serialization.IntegerDeseria
zer
```

Table 11.6. Discussion of entries

Entry	Description
to, **from**	These are "channels."
send, **receive**	These are "methods." Note that the **to** channel is on the **send()** method and the **from** channel is on the **receive()** method.
kafka.bootstrap.servers=kafka:9092	This specifies the URL of the Kafka broker that the application must connect to. You can also specify a URL at the channel level, like this: **mp.messaging.outgoing.to.bootstrap.servers=kafka:9092**
mp.messaging.outgoing.to.connector=smallrye-kafka	This indicates that you want the **to** channel to receive messages from Kafka. SmallRye reactive messaging is a framework for building applications. Note that the **smallrye-kafka** value is SmallRye reactive messaging-specific. If you're provisioning your own server using Galleon, you can enable the Kafka integration by including the **microprofile-reactive-messaging-kafka** Galleon layer.

Entry	Description
mp.messaging.outgoing.to.topic=my-topic	This indicates that you want to send data to a Kafka topic called **my-topic**. A Kafka "topic" is a category or feed name that messages are stored on and published to. All Kafka messages are organized into topics. Producer applications write data *to* topics and consumer applications read data *from* topics.
mp.messaging.outgoing.to.value.serializer=org.apache.kafka.common.serialization.IntegerSerializer	This tells the connector to use **IntegerSerializer** to serialize the values that the **send()** method outputs when it writes to a topic. Kafka provides serializers for standard Java types. You can implement your own serializer by writing a class that implements **org.apache.kafka.common.serialization.Serializer**, and then include that class in your deployment.
mp.messaging.incoming.from.connector=smallrye-kafka	This indicates that you want to use the **from** channel to receive messages from Kafka. Again, the **smallrye-kafka** value is SmallRye reactive messaging-specific.
mp.messaging.incoming.from.topic=my-topic	This indicates that your connector should read data from the Kafka topic called **my-topic**.
mp.messaging.incoming.from.value.deserializer=org.apache.kafka.common.serialization.IntegerDeserializer	This tells the connector to use **IntegerDeserializer** to deserialize the values from the topic before calling the **receive()** method. You can implement your own deserializer by writing a class that implements **org.apache.kafka.common.serialization.Deserializer**, and then include that class in your deployment.

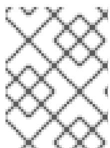

NOTE

This list of properties is not comprehensive. See the SmallRye Reactive Messaging Apache Kafka documentation for more information.

Mandatory MicroProfile Reactive Messaging prefixes

The MicroProfile Reactive Messaging specification requires the following method property key prefixes for Kafka:

- **mp.messaging.incoming.[channel-name].[attribute]=[value]`**
- **mp.messaging.outgoing.[channel-name].[attribute]=[value]`**
- **mp.messaging.connector.[connector-name].[attribute]=[value]`**

Note that **channel-name** is either the **@Incoming.value()** or the **@Outgoing.value()**.

Now consider the following method pair example:

```
@Outgoing("to")
public int send() {
    int i = // Randomly generated...
```

```
    return i;
}

@Incoming("from")
public void receive(int i) {
    // Process payload
}
```

In this method pair example, note the following required property prefixes:

- **mp.messaging.incoming.from**. This prefix selects the property as your configuration of the **receive()** method.
- **mp.messaging.outgoing.to**. This prefix selects the property as your configuration of the **send()** method.

11.6. OPENID CONNECT REFERENCE

11.6.1. elytron-oidc-client subsystem attributes

The **elytron-oidc-client** subsystem provides attributes to configure its behavior.

Table 11.7. elytron-oidc-client subsystem attributes

Attribute	Description
provider	Configuration for an OpenID Connect provider.
secure-deployment	A deployment secured by an OpenID Connect provider.
realm	Configuration for a Red Hat Single Sign-On realm. This is provided for convenience. You can copy the configuration in the keycloak client adapter and use it here. Using the **provider** is recommended instead.

IMPORTANT

Do not use the following **provider**, **realm**, and **secure-deployment** attributes in your configuration as they are not supported at present:

- **autodetect-bearer-only**
- **bearer-only**

Do not use the following **secure-deployment** attributes in your configuration as it is not supported at present

- **enable-basic-auth**

Use the three **elytron-oidc-client** attributes for the following purposes:

- **provider**: For configuring the OpenID Connect provider. For more information, see **provider** attributes.
- **secure-deployment**: For configuring the deployment secured by an OpenID Connect. For more information, see **secure-deployment** attributes
- **realm**: For configuring Red Hat Single Sign-On. For more information, see **realm** attributes. The use of **realm** is not recommended. It is provided for convenience. You can copy the configuration in the keycloak client adapter and use it here. Using the **provider** attribute is recommended instead.

Table 11.8. **provider** attributes

Attribute	Default value	Description
allow-any-hostname	**false**	If you set the value to **true**, hostname verification is skipped when communicating with the OpenID provider. This is useful when testing. Do not set this to **ture** in a production environment.
always-refresh-token		If set to **true**, JBoss EAP refreshes tokens on every web request.
auth-server-url		The base URL of the Red Hat Single Sign-On realm authorization server. If you use this attribute, you must also define the **realm** attribute. You can alternatively use the **provider-url** attribute to provide both base URL and the realm in a single attribute.
client-id		The client-id of JBoss EAP registered with the OpenID provider.
client-key-password		If you specify **client-keystore**, specify it's password in this attribute.
client-keystore		If your application communicates with the OpenID provider over HTTPS, set the path to the client keystore in this attribute.
client-keystore-password		If you specify the **client keystore**, provide the password for accessing it in this attribute.
confidential-port	**8443**	Specify the confidential port (SSL/TLS) used by the OpenID provider.
connection-pool-size		Specify the connection pool size to be used when communicating with the OpenID provider.
connection-timeout-millis		Specify the timeout for establishing a connection with the remote host in milliseconds. The minimum is **-1L**, and the maximum **2147483647L**.**-1L** indicates that the value is undefined, which is the default.

Attribute	Default value	Description
connection-ttl-millis		Specify the amount of time in milliseconds for the connection to be kept alive. The minimum is **-1L**, and the maximum **2147483647L**. **-1L** indicates that the value is undefined, which is the default.
cors-allowed-headers		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Allow-Headers** header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-allowed-methods		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the Access-Control-Allow-Methods header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-exposed-headers		If CORS is enabled, this sets the value of the Access-Control-Expose-Headers header. This should be a comma-separated string. This is optinal. If not set, this header is not returned in CORS responses.
cors-max-age		Set the value for Cross-Origin Resource Sharing (CORS) Max-Age header. The value can be between **-1L** and **2147483647L**. This attribute only takes effect if **enable-cors** is set to **true**.
disable-trust-manager		Specify whether or not to make use of a trust manager when communicating with the OpenID provider over HTTPS.
enable-cors	**false**	Enable Red Hat Single Sign-On Cross-Origin Resource Sharing (CORS) support.
expose-token	**false**	If set to **true**, an authenticated browser client can obtain the signed access token, through a Javascript HTTP invocation, via the URL **root/k_query_bearer_token**. This is optional. This is specific to Red Hat Single Sign-On.
ignore-oauth-query-parameter	**false**	Disable query parameter parsing for access_token.
principal-attribute		Specify which claim value from the ID token to use as the principal for the identity
provider-url		Specify the OpenID provider URL.
proxy-url		Specify the URL for the HTTP proxy if you use one.
realm-public-key		Specify the public key of the realm.

Attribute	Default value	Description
register-node-at-startup	**false**	If set to **true**, a registration request is sent to Red Hat Single Sign-On. This attribute is useful only when your application is clustered.
register-node-period		Specify how often to re-register the node.
socket-timeout-millis		Specify the timeout for socket waiting for data in milliseconds.
ssl-required	**external**	Specify whether communication with the OpenID provider should be over HTTPS. The value can be one of the following: • **all** - all communication happens over HTTPS. • **external** - Only the communication with external clients happens over HTTPs. • **none** - HTTPs is not used.
token-signature-algorithm	**RS256**	Specify the token signature algorithm used by the OpenID provider. The supported algorithms are: • **RS256** • **RS384** • **RS512** • **ES256** • **ES384** • **ES512**
token-store		Specify cookie or session storage for auth-session data.
truststore		Specify the truststore used for client HTTPS requests.
truststore-password		Specify the truststore password.
verify-token-audience	**false**	If set to **true**, then during bearer-only authentication,verify if token contains this client name (**resource**) as an audience.

Table 11.9. **secure-deployment** attributes

Attribute	Default value	Description

Attribute	Default value	Description
allow-any-hostname	**false**	If you set the value to **true**, hostname verification is skipped when communicating with the OpenID provider. This is useful when testing. Do not set this to **ture** in a production environment.
always-refresh-token		If set to **true**, JBoss EAP refreshes tokens on every web request.
auth-server-url		The base URL of the Red Hat Single Sign-On realm authorization server You can alternatively use the **provider-url** attribute.
client-id		The client-id of JBoss EAP registered with the OpenID provider.
client-key-password		If you specify **client-keystore**, specify it’s password in this attribute.
client-keystore		If your application communicates with the OpenID provider over HTTPS, set the path to the client keystore in this attribute.
client-keystore-password		If you specify the **client keystore**, provide the password for accessing it in this attribute.
confidential-port	**8443**	Specify the confidential port (SSL/TLS) used by OpenID provider.
connection-pool-size		Specify the connection pool size to be used when communicating with the OpenID provider.

Attribute	Default value	Description
connection-timeout-millis		Specify the timeout for establishing a connection with the remote host in milliseconds. The minimum is **-1L**, and the maximum **2147483647L**. **-1L** indicates that the value is undefined, which is the default.
connection-ttl-millis		Specify the amount of time in milliseconds for the connection to be kept alive. The minimum is **-1L**, and the maximum **2147483647L**. -1L indicates that the value is undefined, which is the default.
cors-allowed-headers		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Allow-Headers** header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-allowed-methods		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Allow-Methods** header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-exposed-headers		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Expose-Headers** header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-max-age		Set the value for Cross-Origin Resource Sharing (CORS) Max-Age header. The value can be between **-1L** and **2147483647L**. This attribute only takes effect if `enable-

Attribute	Default value	Description
credential		Specify the credential to use to communicate with the OpenID provider.
disable-trust-manager		Specify whether or not to make use of a trust manager when communicating with the OpenID provider over HTTPS.
enable-cors	**false**	Enable Red Hat Single Sign-On Cross-Origin Resource Sharing (CORS) support.
expose-token	**false**	If set to **true**, an authenticated browser client can obtain the signed access token, through a Javascript HTTP invocation, via the URL **root/k_query_bearer_token**. This is optional.This is specific to Red Hat Single Sign-On.
ignore-oauth-query-parameter	**false**	Disable query parameter parsing for access_token.
min-time-between-jwks-requests		If adapter recognizes a token signed by an unknown public key, JBoss EAP tries to download new public key from the **elytron-oidc-client** server. However, JBoss EAP deosn't try to download new public key if it has already tried it in less than the value, in seconds, that you set for this attribute. The value can be between **-1L** and **2147483647L**.
principal-attribute		Specify which claim value from the ID token to use as the principal for the identity
provider		Specify the OpenID provider.
provider-url		Specify the OpenID provider URL.
proxy-url		Specify the URL for the HTTP proxy if you use one.

Attribute	Default value	Description
public-client	**false**	If set to **true**, no client credentials are sent when communicating with the OpenID provider. This is optional.
realm		The realm with which to connect in Red Hat Single Sign-On.
realm-public-key		Specify the public key of the realm.
redirect-rewrite-rule		Specify the rewrite rule to apply to the redirect URI.
register-node-at-startup	**false**	If set to **true**, a registration request is sent to Red Hat Single Sign-On. This attribute is useful only when your application is clustered.
register-node-period		Specify how often to re-register the node.
resource		Specify the name of the application you are securing with OIDC. Alternatively, you can specify the **client-id**.
socket-timeout-millis		Specify the timeout for socket waiting for data in milliseconds.
ssl-required	**external**	Specify whether communication with the OpenID provider should be over HTTPS. The value can be one of the following: • **all** - all communication happens over HTTPS. • **external** - Only the communication with external clients happens over HTTPs. • **none** - HTTPs is not used.

Attribute	Default value	Description
token-minimum-time-to-live		The adapter refreshes the token if the current token is expired or is to expire within the amount of time you set in seconds.
token-signature-algorithm	**RS256**	Specify the token signature algorithm used by the OpenID provider. The supported algorithms are: • **RS256** • **RS384** • **RS512** • **ES256** • **ES384** • **ES512**
token-store		Specify cookie or session storage for auth-session data.
truststore		Specify the truststore used for adapter client HTTPS requests.
truststore-password		Specify the truststore password.
turn-off-change-session-id-on-login	**false**	The session id is changed by default on a successful login. Set the value to **true** to turn this off.
use-resource-role-mappings	**false**	Use resource-level permissions obtained from token.
verify-token-audience	**false**	If set to **true**, then during bearer-only authentication, the adapter verifies if token contains this client name (**resource**) as an audience.

Table 11.10. **realm** attributes

Attribute	Default value	Description

Attribute	Default value	Description
allow-any-hostname	**false**	If you set the value to **true**, hostname verification is skipped when communicating with the OpenID provider. This is useful when testing. Do not set this to **ture** in a production environment.
always-refresh-token		If set to **true**, JBoss EAP refreshes tokens on every web request.
auth-server-url		The base URL of the Red Hat Single Sign-On realm authorization server You can alternatively use the **provider-url** attribute.
client-key-password		If you specify **client-keystore**, specify it's password in this attribute.
client-keystore		If your application communicates with the OpenID provider over HTTPS, set the path to the client keystore in this attribute.
client-keystore-password		If you specify the **client keystore**, provide the password for accessing it in this attribute.
confidential-port	**8443**	Specify the confidential port (SSL/TLS) used by Red Hat Single Sign-On.
connection-pool-size		Specify the connection pool size to be used when communicating with Red Hat Single Sign-On.
connection-timeout-millis		Specify the timeout for establishing a connection with the remote host in milliseconds. The minimum is **-1L**, and the maximum **2147483647L**. **-1L** indicates that the value is undefined, which is the default.

Attribute	Default value	Description
connection-ttl-millis		Specify the amount of time in milliseconds for the connection to be kept alive. The minimum is **-1L**, and the maximum **2147483647L**. **-1L** indicates that the value is undefined, which is the default.
cors-allowed-headers		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the Access-Control-Allow-Headers header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-allowed-methods		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Allow-Methods header**. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-exposed-headers		If Cross-Origin Resource Sharing (CORS) is enabled, this sets the value of the **Access-Control-Expose-Headers** header. This should be a comma-separated string. This is optional. If not set, this header is not returned in CORS responses.
cors-max-age		Set the value for Cross-Origin Resource Sharing (CORS) Max-Age header. The value can be between **-1L** and **2147483647L**. This attribute only takes effect if **enable-cors** is set to **true**.
disable-trust-manager		Specify whether or not to make use of a trust manager when communicating with the OpenID provider over HTTPS._
enable-cors	**false**	Enable {RHProductShortName} Cross-Origin Resource Sharing (CORS) support.

Attribute	Default value	Description
expose-token	**false**	If set to **true**, an authenticated browser client can obtain the signed access token, through a Javascript HTTP invocation, via the URL **root/k_query_bearer_token**. This is optional.
ignore-oauth-query-parameter	**false**	Disable query parameter parsing for access_token.
principal-attribute		Specify which claim value from the ID token to use as the principal for the identity
provider-url		Specify the OpenID provider URL.
proxy-url		Specify the URL for the HTTP proxy if you use one.
realm-public-key		Specify the public key of the realm.
register-node-at-startup	**false**	If set to **true**, a registration request is sent to Red Hat Single Sign-On. This attribute is useful only when your application is clustered.
register-node-period		Specify how often to re-register the node.
socket-timeout-millis		Specify the timeout for socket waiting for data in milliseconds.
ssl-required	**external**	Specify whether communication with the OpenID provider should be over HTTPS. The value can be one of the following: • **all** - all communication happens over HTTPS. • **external** - Only the communication with external clients happens over HTTPs. • **none** - HTTPs is not used.

Attribute	Default value	Description
token-signature-algorithm	**RS256**	Specify the token signature algorithm used by the OpenID provider. The supported algorithms are: • **RS256** • **RS384** • **RS512** • **ES256** • **ES384** • **ES512**
token-store		Specify cookie or session storage for auth-session data.
truststore		Specify the truststore used for client HTTPS requests.
truststore-password		Specify the truststore password.
verify-token-audience	**false**	If set to **true**, then during bearer-only authentication, the adapter verifies if token contains this client name (resource) as an audience.

Additional resources

- OpenID Connect configuration in JBoss EAP
- Securing applications using OpenID Connect with Red Hat Single Sign-On

11.7. OPENTELEMETRY REFERENCE

11.7.1. OpenTelemetry subsystem attributes

You can modify **opentelemetry** subsystem attributes to configure its behavior. The attributes are grouped by the aspect they configure: exporter, sampler, and span processor.

Table 11.11. Exporter attribute group

Attribute	Description	Default value

Attribute	Description	Default value
endpoint	The URL to which OpenTelemetry pushes traces. Set this to the URL where your exporter listens.	http://localhost:14250/
exporter-type	The exporter to which traces are sent. It can be one of the following: • **jaeger**. The exporter you use is Jaeger. • **otlp**. The exporter you use works with the OpenTelemetry protocol.	**jaeger**

Table 11.12. Sampler attribute group

Attribute	Description	Default value
ratio	The ratio of traces to export. The value must be between **0.0** and **1.0**. For example, to export one trace in every 100 traces created by an application, set the value to **0.01**. This attribute takes effect only if you set the attribute **sampler-type** as **ratio**.	

Table 11.13. Span processor attribute group

Attribute	Description	Default value
batch-delay	The interval in milliseconds between two consecutive exports by JBoss EAP. This attribute only takes effect if you set the attribute **span-processor-type** as **batch**.	**5000**
export-timeout	The maximum amount of time in milliseconds to allow for an export to complete before being cancelled.	**30000**

Attribute	Description	Default value
max-export-batch-size	The maximum number of traces that are published in each batch. This number should be should be lesser or equal to the value of **max-queue-size**. You can set this attribute only if you set the attribute **span-processor-type** as **batch**.	**512**
max-queue-size	The maximum number of traces to queue before exporting. If an application creates more traces, they are not recorded. This attribute only takes effect if you set the attribute **span-processor-type** as **batch**.	**2048**

span-processor-type	The type of span processor to use. The value can be one of the following:
	• **batch**: JBoss EAP exports traces in batches that are defined using the following attributes:
	  ◦ **batch-delay**
	  ◦ **max-export-batch-size**
	  ◦ **max-queue-size**
	• **simple**: JBoss EAP exports traces are as soon as they finish.
		batch

Additional resources

- OpenTelemetry in JBoss EAP