



Red Hat Data Grid 8.1

Data Grid Command Line Interface

Access data and manage Data Grid with the CLI

Red Hat Data Grid 8.1 Data Grid Command Line Interface

Access data and manage Data Grid with the CLI

Legal Notice

Copyright © 2021 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

Connect to Data Grid servers via the command line interface (CLI) to access data and perform management operations.

Table of Contents

RED HAT DATA GRID	7
DATA GRID DOCUMENTATION	8
DATA GRID DOWNLOADS	9
MAKING OPEN SOURCE MORE INCLUSIVE	10
CHAPTER 1. GETTING STARTED WITH THE DATA GRID CLI	11
1.1. CREATING AND MODIFYING USERS	11
1.2. CONNECTING TO DATA GRID SERVERS	11
1.3. NAVIGATING CLI RESOURCES	12
1.3.1. CLI Resources	13
1.4. SHUTTING DOWN DATA GRID SERVER	14
1.4.1. Restarting Data Grid Clusters	15
CHAPTER 2. PERFORMING CACHE OPERATIONS WITH THE DATA GRID CLI	16
2.1. CREATING CACHES WITH THE DATA GRID COMMAND LINE INTERFACE (CLI)	16
2.1.1. XML Configuration	16
2.1.2. JSON Configuration	17
2.2. ADDING CACHE ENTRIES	17
2.3. CLEARING CACHES AND DELETING ENTRIES	18
2.4. DELETING CACHES	18
CHAPTER 3. PERFORMING BATCH OPERATIONS	19
3.1. PERFORMING BATCH OPERATIONS WITH FILES	19
3.2. PERFORMING BATCH OPERATIONS INTERACTIVELY	19
CHAPTER 4. CONFIGURING THE DATA GRID CLI	21
4.1. SETTING DATA GRID CLI PROPERTIES AND PERSISTENT STORAGE	21
4.2. CREATING COMMAND ALIASES	21
4.3. TRUSTING DATA GRID SERVER CONNECTIONS	22
4.4. DATA GRID CLI STORAGE DIRECTORY	22
CHAPTER 5. WORKING WITH COUNTERS	24
5.1. CREATING COUNTERS	24
5.2. ADDING DELTAS TO COUNTERS	25
CHAPTER 6. QUERYING CACHES WITH PROTOBUF METADATA	26
6.1. CONFIGURING MEDIA TYPES	26
6.2. REGISTERING PROTOBUF SCHEMAS	27
6.3. QUERYING CACHES WITH PROTOBUF SCHEMAS	28
CHAPTER 7. PERFORMING CROSS-SITE REPLICATION OPERATIONS	31
7.1. BRINGING BACKUP LOCATIONS OFFLINE AND ONLINE	31
7.2. PUSHING STATE TO BACKUP LOCATIONS	31
CHAPTER 8. COMMAND REFERENCE	33
8.1. ADD(1)	33
8.1.1. NAME	33
8.1.2. SYNOPSIS	33
8.1.3. OPTIONS	33
8.1.4. EXAMPLES	33
8.1.5. SEE ALSO	33

8.2. ALIAS(1)	33
8.2.1. NAME	33
8.2.2. SYNOPSIS	33
8.2.3. EXAMPLES	34
8.2.4. SEE ALSO	34
8.3. CACHE(1)	34
8.3.1. NAME	34
8.3.2. SYNOPSIS	34
8.3.3. EXAMPLE	34
8.3.4. SEE ALSO	34
8.4. CAS(1)	34
8.4.1. NAME	34
8.4.2. SYNOPSIS	34
8.4.3. OPTIONS	34
8.4.4. EXAMPLE	35
8.4.5. SEE ALSO	35
8.5. CD(1)	35
8.5.1. NAME	35
8.5.2. DESCRIPTION	35
8.5.3. SYNOPSIS	35
8.5.4. EXAMPLE	35
8.5.5. SEE ALSO	35
8.6. CLEARCACHE(1)	35
8.6.1. NAME	35
8.6.2. SYNOPSIS	35
8.6.3. EXAMPLES	35
8.6.4. SEE ALSO	35
8.7. CONFIG(1)	36
8.7.1. NAME	36
8.7.2. SYNOPSIS	36
8.7.3. DESCRIPTION	36
8.7.4. COMMAND SYNOPSIS	36
8.7.5. COMMON OPTIONS	36
8.7.6. PROPERTIES	36
8.7.7. EXAMPLES	36
8.7.8. SEE ALSO	37
8.8. CONNECT(1)	37
8.8.1. NAME	37
8.8.2. DESCRIPTION	37
8.8.3. SYNOPSIS	37
8.8.4. OPTIONS	37
8.8.5. EXAMPLE	37
8.8.6. SEE ALSO	37
8.9. CONTAINER(1)	38
8.9.1. NAME	38
8.9.2. SYNOPSIS	38
8.9.3. EXAMPLE	38
8.9.4. SEE ALSO	38
8.10. COUNTER(1)	38
8.10.1. NAME	38
8.10.2. SYNOPSIS	38
8.10.3. EXAMPLE	38
8.10.4. SEE ALSO	38

8.11. CREATE(1)	38
8.11.1. NAME	38
8.11.2. SYNOPSIS	38
8.11.3. CREATE CACHE OPTIONS	39
8.11.4. CREATE COUNTER OPTIONS	39
8.11.5. EXAMPLES	39
8.11.6. SEE ALSO	39
8.12. DESCRIBE(1)	39
8.12.1. NAME	39
8.12.2. SYNOPSIS	39
8.12.3. EXAMPLES	39
8.12.4. SEE ALSO	40
8.13. DISCONNECT(1)	40
8.13.1. NAME	40
8.13.2. SYNOPSIS	40
8.13.3. EXAMPLE	40
8.13.4. SEE ALSO	40
8.14. DROP(1)	40
8.14.1. NAME	40
8.14.2. SYNOPSIS	40
8.14.3. EXAMPLES	40
8.14.4. SEE ALSO	41
8.15. ENCODING(1)	41
8.15.1. NAME	41
8.15.2. DESCRIPTION	41
8.15.3. SYNOPSIS	41
8.15.4. EXAMPLE	41
8.15.5. SEE ALSO	41
8.16. GET(1)	41
8.16.1. NAME	41
8.16.2. SYNOPSIS	41
8.16.3. OPTIONS	41
8.16.4. EXAMPLE	42
8.16.5. SEE ALSO	42
8.17. HELP(1)	42
8.17.1. NAME	42
8.17.2. SYNOPSIS	42
8.17.3. EXAMPLE	42
8.17.4. SEE ALSO	42
8.18. LOGGING(1)	42
8.18.1. NAME	42
8.18.2. SYNOPSIS	42
8.18.3. LOGGING SET OPTIONS	42
8.18.4. EXAMPLES	43
8.19. LS(1)	43
8.19.1. NAME	43
8.19.2. SYNOPSIS	43
8.19.3. EXAMPLES	43
8.19.4. SEE ALSO	43
8.20. MIGRATE(1)	43
8.20.1. NAME	43
8.20.2. SYNOPSIS	43
8.20.3. DESCRIPTION	43

8.20.4. COMMAND SYNOPSIS	44
8.20.5. COMMON OPTIONS	44
8.20.6. CLUSTER SYNCHRONIZE OPTIONS	44
8.20.7. CLUSTER DISCONNECT OPTIONS	44
8.21. PATCH(1)	44
8.21.1. NAME	44
8.21.2. DESCRIPTION	44
8.21.3. SYNOPSIS	44
8.21.4. PATCH LIST OPTIONS	45
8.21.5. PATCH INSTALL OPTIONS	45
8.21.6. PATCH DESCRIBE OPTIONS	45
8.21.7. PATCH ROLLBACK OPTIONS	45
8.21.8. PATCH CREATE OPTIONS	45
8.21.9. EXAMPLES	45
8.22. PUT(1)	46
8.22.1. NAME	46
8.22.2. DESCRIPTION	46
8.22.3. SYNOPSIS	46
8.22.4. OPTIONS	46
8.22.5. EXAMPLES	46
8.22.6. SEE ALSO	47
8.23. QUERY(1)	47
8.23.1. NAME	47
8.23.2. SYNOPSIS	47
8.23.3. OPTIONS	47
8.23.4. EXAMPLES	47
8.23.5. SEE ALSO	47
8.24. QUIT(1)	47
8.24.1. NAME	47
8.24.2. SYNOPSIS	47
8.24.3. EXAMPLE	47
8.24.4. SEE ALSO	48
8.25. REMOVE(1)	48
8.25.1. NAME	48
8.25.2. SYNOPSIS	48
8.25.3. OPTIONS	48
8.25.4. EXAMPLE	48
8.25.5. SEE ALSO	48
8.26. RESET(1)	48
8.26.1. NAME	48
8.26.2. SYNOPSIS	48
8.26.3. EXAMPLE	48
8.26.4. SEE ALSO	48
8.27. SCHEMA(1)	49
8.27.1. NAME	49
8.27.2. SYNOPSIS	49
8.27.3. OPTIONS	49
8.27.4. EXAMPLE	49
8.27.5. SEE ALSO	49
8.28. SHUTDOWN(1)	49
8.28.1. NAME	49
8.28.2. SYNOPSIS	49
8.28.3. EXAMPLES	49

8.28.4. SEE ALSO	49
8.29. SITE(1)	50
8.29.1. NAME	50
8.29.2. SYNOPSIS	50
8.29.3. OPTIONS	50
8.29.4. EXAMPLES	50
8.30. STATS(1)	51
8.30.1. NAME	51
8.30.2. SYNOPSIS	51
8.30.3. EXAMPLES	51
8.30.4. SEE ALSO	51
8.31. TASK(1)	51
8.31.1. NAME	51
8.31.2. SYNOPSIS	51
8.31.3. EXAMPLES	51
8.31.4. OPTIONS	51
8.31.5. SEE ALSO	52
8.32. UNALIAS(1)	52
8.32.1. NAME	52
8.32.2. SYNOPSIS	52
8.32.3. EXAMPLES	52
8.32.4. SEE ALSO	52
8.33. USER(1)	52
8.33.1. NAME	52
8.33.2. SYNOPSIS	52
8.33.3. DESCRIPTION	52
8.33.4. COMMAND SYNOPSIS	53
8.33.5. COMMON OPTIONS	53
8.33.6. USER CREATE/MODIFY OPTIONS	53
8.33.7. USER LS OPTIONS	53
8.33.8. USER ENCRYPT-ALL OPTIONS	54
8.34. VERSION(1)	54
8.34.1. NAME	54
8.34.2. SYNOPSIS	54
8.34.3. EXAMPLE	54
8.34.4. SEE ALSO	54

RED HAT DATA GRID

Data Grid is a high-performance, distributed in-memory data store.

Schemaless data structure

Flexibility to store different objects as key-value pairs.

Grid-based data storage

Designed to distribute and replicate data across clusters.

Elastic scaling

Dynamically adjust the number of nodes to meet demand without service disruption.

Data interoperability

Store, retrieve, and query data in the grid from different endpoints.

DATA GRID DOCUMENTATION

Documentation for Data Grid is available on the Red Hat customer portal.

- [Data Grid 8.1 Documentation](#)
- [Data Grid 8.1 Component Details](#)
- [Supported Configurations for Data Grid 8.1](#)
- [Data Grid 8 Feature Support](#)
- [Data Grid Deprecated Features and Functionality](#)

DATA GRID DOWNLOADS

Access the [Data Grid Software Downloads](#) on the Red Hat customer portal.



NOTE

You must have a Red Hat account to access and download Data Grid software.

MAKING OPEN SOURCE MORE INCLUSIVE

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

CHAPTER 1. GETTING STARTED WITH THE DATA GRID CLI

The command line interface (CLI) lets you remotely connect to Data Grid servers to access data and perform administrative functions.

1.1. CREATING AND MODIFYING USERS

Data Grid Server requires users to authenticate against a default property realm. Before you can access Data Grid Server, you must add credentials by creating at least one user and a password. You can also add and modify the security authorization groups to which users belong.

Procedure

1. Open a terminal in **\$RHDG_HOME**.
2. Create and modify Data Grid users with the **user** command.

TIP

Run **help user** for more details about using the command.

Creating users and passwords

- Linux

```
$ bin/cli.sh user create myuser -p "qwer1234!"
```

- Microsoft Windows

```
$ bin\cli.bat user create myuser -p "qwer1234!"
```

Creating users with group membership

- Linux

```
$ bin/cli.sh user create myuser -p "qwer1234!" -g supervisor,reader,writer
```

- Microsoft Windows

```
$ bin\cli.bat user create myuser -p "qwer1234!" -g supervisor,reader,writer
```

1.2. CONNECTING TO DATA GRID SERVERS

Establish CLI connections to Data Grid.

Prerequisites

Add user credentials and have at least one running Data Grid server instance.

Procedure

1. Open a terminal in **\$RHDG_HOME**.

2. Start the CLI.

- **Linux:**

```
$ bin/cli.sh
```

- **Microsoft Windows:**

```
$ bin\cli.bat
```

3. Run the **connect** command and enter your username and password when prompted.

- Data Grid Server on the default port of **11222**:

```
[disconnected]> connect
```

- Data Grid Server with a port offset of **100**:

```
[disconnected]> connect 127.0.0.1:11322
```

1.3. NAVIGATING CLI RESOURCES

The Data Grid CLI exposes a navigable tree that allows you to list, describe, and manipulate Data Grid cluster resources.

TIP

Press the tab key to display available commands and options. Use the **-h** option to display help text.

When you connect to a Data Grid cluster, it opens in the context of the default cache container.

```
[//containers/default]>
```

- Use **ls** to list resources.

```
[//containers/default]> ls  
caches  
counters  
configurations  
schemas  
tasks
```

- Use **cd** to navigate the resource tree.

```
[//containers/default]> cd caches
```

- Use **describe** to view information about resources.

```
[//containers/default]> describe  
{  
  "name" : "default",  
  "version" : "xx.x.x-FINAL",
```



```

"cluster_name" : "cluster",
"coordinator" : true,
"cache_configuration_names" : [ "org.infinispan.REPL_ASYNC", "__protobuf_metadata",
"org.infinispan.DIST_SYNC", "org.infinispan.LOCAL", "org.infinispan.INVALIDATION_SYNC",
"org.infinispan.REPL_SYNC", "org.infinispan.SCATTERED_SYNC",
"org.infinispan.INVALIDATION_ASYNC", "org.infinispan.DIST_ASYNC" ],
"physical_addresses" : "[192.0.2.0:7800]",
"coordinator_address" : "<hostname>",
"cache_manager_status" : "RUNNING",
"created_cache_count" : "1",
"running_cache_count" : "1",
"node_address" : "<hostname>",
"cluster_members" : [ "<hostname1>", "<hostname2>" ],
"cluster_members_physical_addresses" : [ "192.0.2.0:7800", "192.0.2.0:7801" ],
"cluster_size" : 2,
"defined_caches" : [ {
  "name" : "mycache",
  "started" : true
}, {
  "name" : "__protobuf_metadata",
  "started" : true
} ]
}

```

1.3.1. CLI Resources

The Data Grid CLI exposes different resources to:

- create, modify, and manage local or clustered caches.
- perform administrative operations for Data Grid clusters.

Cache Resources

```

[//containers/default]> ls
caches
counters
configurations
schemas

```

caches

Data Grid cache instances. The default cache container is empty. Use the CLI to create caches from templates or **infinispan.xml** files.

counters

Strong or **Weak** counters that record the count of objects.

configurations

Data Grid configurations.

schemas

Protocol Buffers (Protobuf) schemas that structure data in the cache.

tasks

Remote tasks creating and managing Data Grid cache definitions.

Cluster Resources

```
[hostname@cluster/]> ls
containers
cluster
server
```

containers

Cache containers on the Data Grid cluster.

cluster

Lists Data Grid servers joined to the cluster.

server

Resources for managing and monitoring Data Grid servers.

1.4. SHUTTING DOWN DATA GRID SERVER

Stop individually running servers or bring down clusters gracefully.

Procedure


1. Create a CLI connection to Data Grid.
2. Shut down Data Grid Server in one of the following ways:
 - Stop all nodes in a cluster with the **shutdown cluster** command, for example:

```
[//containers/default]> shutdown cluster
```

This command saves cluster state to the **data** folder for each node in the cluster. If you use a cache store, the **shutdown cluster** command also persists all data in the cache.

- Stop individual server instances with the **shutdown server** command and the server hostname, for example:

```
[//containers/default]> shutdown server <my_server01>
```



IMPORTANT

The **shutdown server** command does not wait for rebalancing operations to complete, which can lead to data loss if you specify multiple hostnames at the same time.

TIP

Run **help shutdown** for more details about using the command.

Verification

Data Grid logs the following messages when you shut down servers:

```
ISPN080002: Data Grid Server stopping
ISPN000080: Disconnecting JGroups channel cluster
```

```
ISPN000390: Persisted state, version=<$version> timestamp=YYYY-MM-DDTHH:MM:SS  
ISPN080003: Data Grid Server stopped
```

1.4.1. Restarting Data Grid Clusters

When you bring Data Grid clusters back online after shutting them down, you should wait for the cluster to be available before adding or removing nodes or modifying cluster state.

If you shutdown clustered nodes with the **shutdown server** command, you must restart each server in reverse order.

For example, if you shutdown **server1** and then shutdown **server2**, you should first start **server2** and then start **server1**.

If you shutdown a cluster with the **shutdown cluster** command, clusters become fully operational only after all nodes rejoin.

You can restart nodes in any order but the cluster remains in DEGRADED state until all nodes that were joined before shutdown are running.

CHAPTER 2. PERFORMING CACHE OPERATIONS WITH THE DATA GRID CLI

The command line interface (CLI) lets you remotely connect to Data Grid servers to access data and perform administrative functions.

2.1. CREATING CACHES WITH THE DATA GRID COMMAND LINE INTERFACE (CLI)

Use the Data Grid CLI to add caches from templates or configuration files in XML or JSON format.

Prerequisites

Create a user and start at least one Data Grid server instance.

Procedure

1. Create a CLI connection to Data Grid.
2. Add cache definitions with the **create cache** command.
 - Add a cache definition from an XML or JSON file with the **--file** option.

```
[//containers/default]> create cache --file=configuration.xml mycache
```

- Add a cache definition from a template with the **--template** option.

```
[//containers/default]> create cache --template=org.infinispan.DIST_SYNC mycache
```

TIP

Press the tab key after the **--template=** argument to list available cache templates.

3. Verify the cache exists with the **ls** command.

```
[//containers/default]> ls caches  
mycache
```

4. Retrieve the cache configuration with the **describe** command.

```
[//containers/default]> describe caches/mycache
```

Reference

- [Creating Data Grid CLI Connections](#)
- [Performing Cache Operations with the Data Grid CLI](#)

2.1.1. XML Configuration

Data Grid configuration in XML format must conform to the schema and include:

- `<infinispan>` root element.
- `<cache-container>` definition.

Example XML Configuration

```
<infinispan>
  <cache-container>
    <distributed-cache name="myCache" mode="SYNC">
      <encoding media-type="application/x-protostream"/>
      <memory max-count="1000000" when-full="REMOVE"/>
    </distributed-cache>
  </cache-container>
</infinispan>
```

2.1.2. JSON Configuration

Data Grid configuration in JSON format:

- Requires the cache definition only.
- Must follow the structure of an XML configuration.
 - XML elements become JSON objects.
 - XML attributes become JSON fields.

Example JSON Configuration

```
{
  "distributed-cache": {
    "name": "myCache",
    "mode": "SYNC",
    "encoding": {
      "media-type": "application/x-protostream"
    },
    "memory": {
      "max-count": 1000000,
      "when-full": "REMOVE"
    }
  }
}
```

2.2. ADDING CACHE ENTRIES

Create **key:value** pair entries in the data container.

Prerequisites

Create a Data Grid cache that can store your data.

Procedure

1. Create a CLI connection to Data Grid.

2. Add entries into your cache as follows:

- Use the **put** command from the context of a cache:

```
[/containers/default/caches/mycache]> put hello world
```

- Use the **--cache=** with the **put** command:

```
[/containers/default]> put --cache=mycache hello world
```

3. Use the **get** command to verify entries.

```
[/containers/default/caches/mycache]> get hello  
world
```

2.3. CLEARING CACHES AND DELETING ENTRIES

Remove data from caches with the Data Grid CLI.

Procedure

1. Create a CLI connection to Data Grid.
2. Do one of the following:
 - Delete all entries with the **clearcache** command.

```
[/containers/default]> clearcache mycache
```

- Remove specific entries with the **remove** command.

```
[/containers/default]> remove --cache=mycache hello
```

2.4. DELETING CACHES

Drop caches to remove them and delete all data they contain.

Procedure

1. Create a CLI connection to Data Grid.
2. Remove caches with the **drop** command.

```
[/containers/default]> drop cache mycache
```

CHAPTER 3. PERFORMING BATCH OPERATIONS

Process operations in groups, either interactively or using batch files.

Prerequisites

- A running Data Grid cluster.

3.1. PERFORMING BATCH OPERATIONS WITH FILES

Create files that contain a set of operations and then pass them to the Data Grid CLI.

Procedure

1. Create a file that contains a set of operations.
For example, create a file named **batch** that creates a cache named **mybatch**, adds two entries to the cache, and disconnects from the CLI.

```
$ cat > batch<<EOF
create cache --template=org.infinispan.DIST_SYNC mybatch
put --cache=mybatch hello world
put --cache=mybatch hola mundo
disconnect
EOF
```

2. Run the CLI and specify the file as input.

```
$ bin/cli.sh -c localhost:11222 -f batch
```

3. Create a new Data Grid CLI connection and verify **mybatch**.

```
[/containers/default]> ls caches
__protobuf_metadata
mybatch
[/containers/default]> ls caches/mybatch
hola
hello
[/containers/default]> disconnect
[disconnected]>
```



NOTE

CLI batch files support system property expansion. Strings that use the **`\${property}** format are replaced with the value of the **property** system property.

3.2. PERFORMING BATCH OPERATIONS INTERACTIVELY

Use the standard input stream, **stdin**, to perform batch operations interactively.

Procedure

1. Start the Data Grid CLI in interactive mode.

```
$ bin/cli.sh -c localhost:11222 -f -
```



NOTE

If you do not use the **-c** flag, you must run the **connect** command.

```
$ bin/cli.sh -f -
connect
```

- Run batch operations, for example:

```
create cache --template=org.infinispan.DIST_SYNC mybatch
put --cache=mybatch hello world
put --cache=mybatch hola mundo
disconnect
quit
```

TIP

Use **echo** to add commands in interactive mode.

The following example shows how to use **echo describe** to get cluster information:

```
$ echo describe|bin/cli.sh -c localhost:11222 -f -
{
  "name" : "default",
  "version" : "10.0.0-SNAPSHOT",
  "coordinator" : false,
  "cache_configuration_names" : [ "org.infinispan.REPL_ASYNC", "__protobuf_metadata",
  "org.infinispan.DIST_SYNC", "qcache", "org.infinispan.LOCAL", "dist_cache_01",
  "org.infinispan.INVALIDATION_SYNC", "org.infinispan.REPL_SYNC",
  "org.infinispan.SCATTERED_SYNC", "mycache", "org.infinispan.INVALIDATION_ASYNC",
  "mybatch", "org.infinispan.DIST_ASYNC" ],
  "cluster_name" : "cluster",
  "physical_addresses" : "[192.168.1.7:7800]",
  "coordinator_address" : "thundercat-34689",
  "cache_manager_status" : "RUNNING",
  "created_cache_count" : "4",
  "running_cache_count" : "4",
  "node_address" : "thundercat-47082",
  "cluster_members" : [ "thundercat-34689", "thundercat-47082" ],
  "cluster_members_physical_addresses" : [ "10.36.118.25:7801", "192.168.1.7:7800" ],
  "cluster_size" : 2,
  "defined_caches" : [ {
    "name" : "__protobuf_metadata",
    "started" : true
  }, {
    "name" : "mybatch",
    "started" : true
  } ]
}
```


CHAPTER 4. CONFIGURING THE DATA GRID CLI

Define configuration properties for the Data Grid CLI.

4.1. SETTING DATA GRID CLI PROPERTIES AND PERSISTENT STORAGE

Configure Data Grid CLI startup operations and customize the location for persistent storage.

Prerequisites

Create at least one Data Grid user.

Procedure

1. Optionally set a custom path to the Data Grid CLI storage directory in one of the following ways:

- Using the **cli.dir** system property:

```
$ bin/cli.sh -Dcli.dir=/path/to/cli/storage ...
```

- Using the **ISPN_CLI_DIR** environment variable:

```
export ISPN_CLI_DIR=/path/to/cli/storage
$ bin/cli.sh ...
```

2. Set values for configuration properties with the **config set** command.
3. Verify configuration properties with the **config get** command.

TIP

Run **help config** to review available configuration properties and get example usage.

4.2. CREATING COMMAND ALIASES

Create aliases for Data Grid CLI commands to define custom shortcuts.

Procedure

1. Create aliases with the **alias <alias>=<command>** command.
For example, set **q** as an alias for the **quit** command:

```
[//containers/default]> alias q=quit
```

2. Run the **alias** command to check the defined aliases.

```
[//containers/default]> alias
alias q='quit'
```

3. Delete aliases with the **unalias** command, for example:

```
[//containers/default]> unalias q
```

4.3. TRUSTING DATA GRID SERVER CONNECTIONS

Secure Data Grid CLI connections to Data Grid Server with SSL/TLS certificates. If you create a key store as an SSL identity for Data Grid Server, the CLI can validate server certificates to verify the identity.

Prerequisites

- Set up an SSL identity for Data Grid Server.
- Create at least one Data Grid user.

Procedure

1. Specify the location of the server key store, as in the following example:

```
$ bin/cli.sh config set truststore /home/user/my-trust-store.jks
```

2. Define the key store password, if necessary, as follows:

```
$ bin/cli.sh config set truststore-password secret
```

3. Verify your CLI configuration.

```
$ bin/cli.sh config get truststore
truststore=/home/user/my-trust-store.jks

$ bin/cli.sh config get truststore-password
truststore-password=secret
```

Reference

[Setting Up SSL Identities for Data Grid Server](#)

4.4. DATA GRID CLI STORAGE DIRECTORY

Data Grid CLI stores configuration in the following default directory:

Operating System	Default Path
Linux/Unix	\$HOME/.config/red_hat_data_grid
Microsoft Windows	%APPDATA%/Sun/Java/red_hat_data_grid
Mac OS	\$HOME/Library/Java/red_hat_data_grid

This directory contains the following files:

cli.properties

Stores values for CLI configuration properties.

aliases

Stores command aliases.

history

Stores CLI history.

CHAPTER 5. WORKING WITH COUNTERS

Counters provide atomic increment and decrement operations that record the count of objects.

Prerequisites

- Start the Data Grid CLI.
- Connect to a running Data Grid cluster.

5.1. CREATING COUNTERS

Create strong and weak counters with the Data Grid CLI.

Procedure

1. Create a CLI connection to Data Grid.
2. Run the **create counter** command with the appropriate arguments.

- a. Create **my-weak-counter**.

```
[//containers/default]> create counter --concurrency-level=1 --initial-value=5 --  
storage=PERSISTENT --type=weak my-weak-counter
```

- b. Create **my-strong-counter**.

```
[//containers/default]> create counter --initial-value=3 --storage=PERSISTENT --  
type=strong my-strong-counter
```

3. List available counters.

```
[//containers/default]> ls counters  
my-strong-counter  
my-weak-counter
```

4. Verify counter configurations.

- a. Describe **my-weak-counter**.

```
[//containers/default]> describe counters/my-weak-counter  
  
{  
  "weak-counter":{  
    "initial-value":5,  
    "storage":"PERSISTENT",  
    "concurrency-level":1  
  }  
}
```

- b. Describe **my-strong-counter**.

```
[//containers/default]> describe counters/my-strong-counter
```

```
{
  "strong-counter":{
    "initial-value":3,
    "storage":"PERSISTENT",
    "upper-bound":5
  }
}
```

5.2. ADDING DELTAS TO COUNTERS

Increment or decrement counters with arbitrary values.

Procedure

1. Select a counter.

```
[//containers/default]> counter my-weak-counter
```

2. List the current count.

```
[//containers/default/counters/my-weak-counter]> ls
5
```

3. Increment the counter value by **2**.

```
[//containers/default/counters/my-weak-counter]> add --delta=2
```

4. Decrement the counter value by **-4**.

```
[//containers/default/counters/my-weak-counter]> add --delta=-4
```



NOTE

Strong counters return values after the operation is applied. Use **--quiet=true** to hide the return value.

For example, **my-strong-counter]> add --delta=3 --quiet=true**.

Weak counters return empty responses.

CHAPTER 6. QUERYING CACHES WITH PROTOBUF METADATA

Data Grid supports using Protocol Buffers (Protobuf) to structure data in the cache so that you can query it.

Prerequisites

- Start the Data Grid CLI.
- Connect to a running Data Grid cluster.

6.1. CONFIGURING MEDIA TYPES

Encode cache entries with different media types to store data in a format that best suits your requirements.

For example, the following procedure shows you how to configure the **application/x-protostream** media type.

Procedure

1. Create a Data Grid configuration file that adds a distributed cache named **qcache** and configures the media type, for example:

```
<infinispan>
  <cache-container>
    <distributed-cache name="qcache">
      <encoding>
        <key media-type="application/x-protostream"/>
        <value media-type="application/x-protostream"/>
      </encoding>
    </distributed-cache>
  </cache-container>
</infinispan>
```

2. Create **qcache** from **pcache.xml** with the **--file=** option.

```
[//containers/default]> create cache --file=pcache.xml pcache
```

3. Verify **pcache**.

```
[//containers/default]> ls caches
pcache
__protobuf_metadata
[//containers/default]> describe caches/pcache
{
  "distributed-cache" : {
    "mode" : "SYNC",
    "encoding" : {
      "key" : {
        "media-type" : "application/x-protostream"
      },

```

```

    "value" : {
      "media-type" : "application/x-protostream"
    }
  },
  "transaction" : {
    "mode" : "NONE"
  }
}
}
}

```

4. Add an entry to **pcache** and check the encoding.

```

[//containers/default]> put --cache=pcache good morning
[//containers/default]> cd caches/pcache
[//containers/default/caches/pcache]> get good
{
  "_type" : "string",
  "_value" : "morning"
}

```

6.2. REGISTERING PROTOBUF SCHEMAS

Protobuf schemas contain data structures known as messages in **.proto** definition files.

Procedure

1. Create a schema file named **person.proto** with the following messages:

```

package org.infinispan.rest.search.entity;

message Address {
  required string street = 1;
  required string postCode = 2;
}

message PhoneNumber {
  required string number = 1;
}

message Person {
  optional int32 id = 1;
  required string name = 2;
  required string surname = 3;
  optional Address address = 4;
  repeated PhoneNumber phoneNumbers = 5;
  optional uint32 age = 6;
  enum Gender {
    MALE = 0;
    FEMALE = 1;
  }
  optional Gender gender = 7;
}

```

2. Register **person.proto**.

```
[[/containers/default]> schema --upload=person.proto person.proto
```

3. Verify **person.proto**.

```
[[/containers/default]> cd caches/___protobuf_metadata
[[/containers/default/caches/___protobuf_metadata]> ls
person.proto
[[/containers/default/caches/___protobuf_metadata]> get person.proto
```

6.3. QUERYING CACHES WITH PROTOBUF SCHEMAS

Data Grid automatically converts JSON to Protobuf so that you can read and write cache entries in JSON format and use Protobuf schemas to query them.

For example, consider the following JSON documents:

lukecage.json

```
{
  "_type": "org.infinispan.rest.search.entity.Person",
  "id": 2,
  "name": "Luke",
  "surname": "Cage",
  "gender": "MALE",
  "address": {"street": "38th St", "postCode": "NY 11221"},
  "phoneNumbers": [{"number": "4444"}, {"number": "5555"}]
}
```

jessicajones.json

```
{
  "_type": "org.infinispan.rest.search.entity.Person",
  "id": 1,
  "name": "Jessica",
  "surname": "Jones",
  "gender": "FEMALE",
  "address": {"street": "46th St", "postCode": "NY 10036"},
  "phoneNumbers": [{"number": "1111"}, {"number": "2222"}, {"number": "3333"}]
}
```

matthewmurdock.json

```
{
  "_type": "org.infinispan.rest.search.entity.Person",
  "id": 3,
  "name": "Matthew",
  "surname": "Murdock",
  "gender": "MALE",
  "address": {"street": "57th St", "postCode": "NY 10019"},
  "phoneNumbers": []
}
```


Each of the preceding JSON documents contains:

- A **_type** field that identifies the Protobuf message to which the JSON document corresponds.
- Several fields that correspond to datatypes in the **person.proto** schema.

Procedure

1. Navigate to the **pcache** cache.

```
[//containers/default/caches]> cd pcache
```

2. Add each JSON document as an entry to the cache, for example:

```
[//containers/default/caches/pcache]> put --encoding=application/json --file=jessicajones.json
jessicajones
[//containers/default/caches/pcache]> put --encoding=application/json --
file=matthewmurdock.json matthewmurdock
[//containers/default/caches/pcache]> put --encoding=application/json --file=lukecage.json
lukecage
```

3. Verify that the entries exist.

```
[//containers/default/caches/pcache]> ls
lukecage
matthewmurdock
jessicajones
```

4. Query the cache to return entries from the Protobuf **Person** entity where the gender datatype is **MALE**.

```
[//containers/default/caches/pcache]> query "from org.infinispan.rest.search.entity.Person p
where p.gender = 'MALE'"
{
  "total_results" : 2,
  "hits" : [ {
    "hit" : {
      "_type" : "org.infinispan.rest.search.entity.Person",
      "id" : 2,
      "name" : "Luke",
      "surname" : "Cage",
      "gender" : "MALE",
      "address" : {
        "street" : "38th St",
        "postCode" : "NY 11221"
      },
      "phoneNumbers" : [ {
        "number" : "4444"
      }, {
        "number" : "5555"
      } ]
    }
  }, {
    "hit" : {
      "_type" : "org.infinispan.rest.search.entity.Person",
```

```
"id" : 3,  
"name" : "Matthew",  
"surname" : "Murdock",  
"gender" : "MALE",  
"address" : {  
  "street" : "57th St",  
  "postCode" : "NY 10019"  
}  
}  
}]  
}
```

CHAPTER 7. PERFORMING CROSS-SITE REPLICATION OPERATIONS

Data Grid clusters running in different locations can discover and communicate with each other to backup data.

Prerequisites

- Start the Data Grid CLI.
- Connect to a running Data Grid cluster.

7.1. BRINGING BACKUP LOCATIONS OFFLINE AND ONLINE

Take backup locations offline manually and bring them back online.

Procedure

1. Create a CLI connection to Data Grid.
2. Check if backup locations are online or offline with the **site status** command:

```
//containers/default]> site status --cache=cacheName --site=NYC
```



NOTE

--site is an optional argument. If not set, the CLI returns all backup locations.

3. Manage backup locations as follows:
 - Bring backup locations online with the **bring-online** command:

```
//containers/default]> site bring-online --cache=customers --site=NYC
```

- Take backup locations offline with the **take-offline** command:

```
//containers/default]> site take-offline --cache=customers --site=NYC
```

For more information and examples, run the **help site** command.

7.2. PUSHING STATE TO BACKUP LOCATIONS

Transfer cache state to remote backup locations.

Procedure

1. Create a CLI connection to Data Grid.
2. Use the **site** command to push state transfer, as in the following example:

```
//containers/default]> site push-site-state --cache=cacheName --site=NYC
```

For more information and examples, run the **help site** command.

CHAPTER 8. COMMAND REFERENCE

Review manual pages for Data Grid CLI commands.

TIP

Use **help** command to access manual pages directly from your CLI session.

For example, to view the manual page for the **get** command do the following:

```
$ help get
```

8.1. ADD(1)

8.1.1. NAME

add - increments and decrements counters with arbitrary values.

8.1.2. SYNOPSIS

```
add ['OPTIONS'] ['COUNTER_NAME']
```

8.1.3. OPTIONS

--delta='nnn'

Sets a delta to increment or decrement the counter value. Defaults to **1**.

-q, --quiet=[true|false]

Hides return values for strong counters. The default is **false**.

8.1.4. EXAMPLES

```
add --delta=10 cnt_a
```

Increments the value of **cnt_a** by **10**.

```
add --delta=-5 cnt_a
```

Decrements the value of **cnt_a** by **5**.

8.1.5. SEE ALSO

`cas(1)`, `reset(1)`

8.2. ALIAS(1)

8.2.1. NAME

alias - creates or displays aliases.

8.2.2. SYNOPSIS

```
alias ['ALIAS-NAME']='COMMAND']
```

8.2.3. EXAMPLES

alias q=quit

Creates **q** as an alias for the **quit** command.

alias

Lists all defined aliases.

8.2.4. SEE ALSO

config(1), unalias(1)

8.3. CACHE(1)

8.3.1. NAME

cache - selects the default cache for subsequent commands.

8.3.2. SYNOPSIS

```
cache ['CACHE_NAME']
```

8.3.3. EXAMPLE

cache mycache

Selects **mycache** and is the same as navigating the resource tree using **cd caches/mycache**.

8.3.4. SEE ALSO

cd(1), clear(1), container(1), get(1), put(1), remove(1)

8.4. CAS(1)

8.4.1. NAME

cas - performs 'compare-and-swap' operations on strong counters.

8.4.2. SYNOPSIS

```
cas ['OPTIONS'] ['COUNTER_NAME']
```

8.4.3. OPTIONS

--expect='nnn'

Specifies the expected value of the counter.

--value='nnn'

Sets a new value for the counter.

-q, --quiet='[true|false]'

Hides return values. The default is false.

8.4.4. EXAMPLE

cas --expect=10 --value=20 cnt_a

Sets the value of **cnt_a** to **20** only if the current value is **10**

8.4.5. SEE ALSO

add(1), cas(1), reset(1)

8.5. CD(1)

8.5.1. NAME

cd - navigates the server resource tree.

8.5.2. DESCRIPTION

PATH can be absolute or relative to the current resource. **../** specifies parent resources.

8.5.3. SYNOPSIS

cd ['PATH']

8.5.4. EXAMPLE

cd caches

Changes to the **caches** path in the resource tree.

8.5.5. SEE ALSO

cache(1), ls(1), container(1)

8.6. CLEARCACHE(1)

8.6.1. NAME

clearcache - removes all entries from a cache.

8.6.2. SYNOPSIS

clearcache ['CACHE_NAME']

8.6.3. EXAMPLES

clearcache mycache

Removes all entries from **mycache**.

8.6.4. SEE ALSO

cache(1), drop(1), remove(1)

8.7. CONFIG(1)

8.7.1. NAME

config - manages CLI configuration properties.

8.7.2. SYNOPSIS

config

config set 'name' 'value'

config get 'name'

8.7.3. DESCRIPTION

Manage (list, set, get) CLI configuration properties.

8.7.4. COMMAND SYNOPSIS

config

Lists all configuration properties that are set.

config set 'name' ['value']

Sets the value of a specific property. If you do not specify a value, the property is not set.

config get 'name'

Retrieves the value of a specific property.

8.7.5. COMMON OPTIONS

These options apply to all commands:

-h, --help

Displays a help page for the command or sub-command.

8.7.6. PROPERTIES

autoconnect-url

Specifies the URL to which the CLI automatically connects on startup.

autoexec

Specifies the path of a CLI batch file to execute on startup.

trustall

Specifies whether to trust all server certificates. Values are **false** (default) and **true**.

truststore

Defines the path to a keystore that contains a certificate chain that verifies server identity.

truststore-password

Specifies a password to access the keystore.

8.7.7. EXAMPLES

config set autoconnect-url <http://192.0.2.0:11222>

Connects to a server at a custom IP address when you start the CLI.

config get autoconnect-url

Returns the value for the **autoconnect-url** configuration property.

config set autoexec /path/to/mybatchfile

Runs a batch file named "mybatchfile" when you start the CLI.

config set trustall true

Trusts all server certificates.

config set truststore /home/user/my-trust-store.jks

Specifies the path of a keystore named "my-trust-store.jks".

config set truststore-password secret

Sets the keystore password, if required.

8.7.8. SEE ALSO

[alias\(1\)](#), [unalias\(1\)](#)

8.8. CONNECT(1)

8.8.1. NAME

`connect` - connects to running Data Grid servers.

8.8.2. DESCRIPTION

Defaults to <http://localhost:11222> and prompts for credentials if authentication is required.

8.8.3. SYNOPSIS

```
connect ['OPTIONS'] ['SERVER_LOCATION']
```

8.8.4. OPTIONS

-u, --username='USERNAME'

Specifies a username to authenticate with Data Grid servers.

-p, --password='PASSWORD'

Specifies passwords.

8.8.5. EXAMPLE

```
connect 127.0.0.1:11322 -u test -p changeme
```

Connects to a locally running server using a port offset of **100** and example credentials.

8.8.6. SEE ALSO

[disconnect\(1\)](#)

8.9. CONTAINER(1)

8.9.1. NAME

container - selects the container for running subsequent commands.

8.9.2. SYNOPSIS

```
container ['CONTAINER_NAME']
```

8.9.3. EXAMPLE

container default

Selects the default container and is the same as navigating the resource tree using **cd containers/default**.

8.9.4. SEE ALSO

cd(1), clear(1), container(1), get(1), put(1), remove(1)

8.10. COUNTER(1)

8.10.1. NAME

counter - selects the default counter for subsequent commands.

8.10.2. SYNOPSIS

```
counter ['COUNTER_NAME']
```

8.10.3. EXAMPLE

counter cnt_a

Selects **cnt_a** and is the same as navigating the resource tree using **cd counters/cnt_a**.

8.10.4. SEE ALSO

add(1), cas(1)

8.11. CREATE(1)

8.11.1. NAME

create - creates caches and counters on Data Grid servers.

8.11.2. SYNOPSIS

```
create cache ['OPTIONS'] CACHE_NAME
```

```
create counter ['OPTIONS'] COUNTER_NAME
```

8.11.3. CREATE CACHE OPTIONS

-f, --file='FILE'

Specifies a configuration file in JSON or XML format.

-t, --template='TEMPLATE'

Specifies a configuration template. Use tab autocompletion to see available templates.

-v, --volatile='[true|false]'

Specifies whether the cache is persistent or volatile. The default is false.

8.11.4. CREATE COUNTER OPTIONS

-t, --type='[weak|strong]'

Specifies if the counter is weak or strong.

-s, --storage='[PERSISTENT|VOLATILE]'

Specifies whether the counter is persistent or volatile.

-c, --concurrency-level='nnn'

Sets the concurrency level of the counter.

-i, --initial-value='nnn'

Sets the initial value of the counter.

-l, --lower-bound='nnn'

Sets the lower bound of a **strong** counter.

-u, --upper-bound='nnn'

Sets the upper bound of a **strong** counter.

8.11.5. EXAMPLES

create cache --template=org.infinispan.DIST_SYNC mycache

Creates a cache named **mycache** from the **DIST_SYNC** template.

create counter --initial-value=3 --storage=PERSISTENT --type=strong cnt_a

Creates a strong counter named **cnt_a**.

8.11.6. SEE ALSO

drop(1)

8.12. DESCRIBE(1)

8.12.1. NAME

describe - displays information about resources.

8.12.2. SYNOPSIS

describe ['PATH']

8.12.3. EXAMPLES

describe //containers/default

Displays information about the default container.

describe //containers/default/caches/mycache

Displays information about the **mycache** cache.

describe //containers/default/caches/mycache/k1

Displays information about the **k1** key.

describe //containers/default/counters/cnt1

Displays information about the **cnt1** counter.

8.12.4. SEE ALSO

cd(1), ls(1)

8.13. DISCONNECT(1)

8.13.1. NAME

disconnect - ends CLI sessions with Data Grid servers.

8.13.2. SYNOPSIS

disconnect

8.13.3. EXAMPLE

disconnect

Ends the current CLI session.

8.13.4. SEE ALSO

connect(1)

8.14. DROP(1)

8.14.1. NAME

drop - deletes caches and counters.

8.14.2. SYNOPSIS

drop cache **CACHE_NAME**

drop counter **COUNTER_NAME**

8.14.3. EXAMPLES

drop cache mycache

Deletes the **mycache** cache.

drop counter cnt_a

Deletes the **cnt_a** counter.

8.14.4. SEE ALSO

create(1), clearcache(1)

8.15. ENCODING(1)**8.15.1. NAME**

encoding - displays and sets the encoding for cache entries.

8.15.2. DESCRIPTION

Sets a default encoding for **put** and **get** operations on a cache. If no argument is specified, the **encoding** command displays the current encoding.

Valid encodings use standard MIME type (IANA media types) naming conventions, such as the following:

- **text/plain**
- **application/json**
- **application/xml**
- **application/octet-stream**

8.15.3. SYNOPSIS

encoding ['ENCODING']

8.15.4. EXAMPLE

encoding application/json

Configures the currently selected cache to encode entries as **application/json**.

8.15.5. SEE ALSO

get(1), put(1)

8.16. GET(1)**8.16.1. NAME**

get - retrieves entries from a cache.

8.16.2. SYNOPSIS

get ['OPTIONS'] KEY

8.16.3. OPTIONS

-c, --cache='NAME'

Specifies the cache from which to retrieve entries. Defaults to the currently selected cache.

8.16.4. EXAMPLE

get hello -c mycache

Retrieves the value of the key named **hello** from **mycache**.

8.16.5. SEE ALSO

query(1), put(1)

8.17. HELP(1)

8.17.1. NAME

help - prints manual pages for commands.

8.17.2. SYNOPSIS

help ['COMMAND']

8.17.3. EXAMPLE

help get

Prints the manual page for the **get** command.

8.17.4. SEE ALSO

version(1)

8.18. LOGGING(1)

8.18.1. NAME

logging - inspects and manipulates the Data Grid server runtime logging configuration.

8.18.2. SYNOPSIS

logging list-loggers

logging list-appenders

logging set ['OPTIONS'] [LOGGER_NAME]

logging remove **LOGGER_NAME**


8.18.3. LOGGING SET OPTIONS

-l, --level='OFF|TRACE|DEBUG|INFO|WARN|ERROR|ALL'

Specifies the logging level for the specific logger.

-a, --appender='APPENDER'

Specifies an appenders to set on the specific logger. The option can be repeated for multiple appenders.



NOTE

calling **logging set** without a logger name will modify the root logger.

8.18.4. EXAMPLES

logging list-loggers

Lists all available loggers

logging set --level=DEBUG --appenders=FILE org.infinispan

Sets the log level for the **org.infinispan** logger to **DEBUG** and configures it to use the **FILE** appender.

8.19. LS(1)

8.19.1. NAME

ls - lists resources for the current path or a given path.

8.19.2. SYNOPSIS

ls ['PATH']

8.19.3. EXAMPLES

ls caches

Lists the available caches.

ls ../

Lists parent resources.

8.19.4. SEE ALSO

cd(1)

8.20. MIGRATE(1)

8.20.1. NAME

migrate - migrates data from one version of Data Grid to another.

8.20.2. SYNOPSIS

migrate cluster synchronize

migrate cluster disconnect

8.20.3. DESCRIPTION

Use the **migrate** command to migrate data from one version of Data Grid to another.

8.20.4. COMMAND SYNOPSIS

Migrate clusters

migrate cluster synchronize

Synchronize data between the source cluster and the target cluster.

migrate cluster disconnect

Disconnects the target cluster from the source cluster.

8.20.5. COMMON OPTIONS

These options apply to all commands:

-h, --help

Displays a help page for the command or sub-command.

8.20.6. CLUSTER SYNCHRONIZE OPTIONS

-c, --cache='name'

The name of the cache to synchronize.

-b, --read-batch='num'

The amount of entries to process in a batch. Defaults to 10000.

-t, --threads='num'

The number of threads to use. Defaults to the number of cores on the server.

8.20.7. CLUSTER DISCONNECT OPTIONS

-c, --cache='name'

The name of the cache to disconnect from the source.

8.21. PATCH(1)

8.21.1. NAME

patch - manages server patches.

8.21.2. DESCRIPTION

List, describe, install, rollback, and create server patches.

Patches are zip archive files that contain artifacts to upgrade servers and resolve issues or add new features. Patches can apply target versions to multiple server installations with different versions.

8.21.3. SYNOPSIS

patch ls

patch install 'patch-file'

patch describe 'patch-file'

patch rollback

patch create 'patch-file' 'target-server' 'source-server-1' ['source-server-2'...]

8.21.4. PATCH LIST OPTIONS

--server='path/to/server'

Sets the path to a target server outside the current server home directory.

-v, --verbose

Shows the content of each installed patch, including information about individual files.

8.21.5. PATCH INSTALL OPTIONS

--dry-run

Shows the operations that the patch performs without applying any changes.

--server='path/to/server'

Sets the path to a target server outside the current server home directory.

8.21.6. PATCH DESCRIBE OPTIONS

-v, --verbose

Shows the content of the patch, including information about individual files

8.21.7. PATCH ROLLBACK OPTIONS

--dry-run

Shows the operations that the patch performs without applying any changes.

--server='path/to/server'

Sets the path to a target server outside the current server home directory.

8.21.8. PATCH CREATE OPTIONS

-q, --qualifier='name'

Specifies a descriptive qualifier string for the patch; for example, 'one-off for issue nnnn'.

8.21.9. EXAMPLES

patch ls

Lists the patches currently installed on a server in order of installation.

patch install mypatch.zip

Installs "mypatch.zip" on a server in the current directory.

patch install mypatch.zip --server=/path/to/server/home

Installs "mypatch.zip" on a server in a different directory.

patch describe mypatch.zip

Displays the target version and list of source versions for "mypatch.zip".

patch create mypatch.zip 'target-server' 'source-server-1' ['source-server-2'...]

Creates a patch file named "mypatch.zip" that uses the version of the target server and applies to the source server versions.

patch rollback

Rolls back the last patch that was applied to a server and restores the previous version.

8.22. PUT(1)

8.22.1. NAME

put - adds or updates cache entries.

8.22.2. DESCRIPTION

Creates entries for new keys. Replaces values for existing keys.

8.22.3. SYNOPSIS

```
put ['OPTIONS'] KEY [VALUE]
```

8.22.4. OPTIONS

-c, --cache='NAME'

Specifies the name of the cache. Defaults to the currently selected cache.

-e, --encoding='ENCODING'

Sets the media type for the value.

-f, --file='FILE'

Specifies a file that contains the value for the entry.

-l, --ttl='TTL'

Sets the number of seconds before the entry is automatically deleted (time-to-live). Defaults to the value for **lifespan** in the cache configuration if **0** or not specified. If you set a negative value, the entry is never deleted.

-i, --max-idle='MAXIDLE'

Sets the number of seconds that the entry can be idle. If a read or write operation does not occur for an entry after the maximum idle time elapses, the entry is automatically deleted. Defaults to the value for **maxidle** in the cache configuration if **0** or not specified. If you set a negative value, the entry is never deleted.

-a, --if-absent=[true|false]

Puts an entry only if it does not exist.

8.22.5. EXAMPLES

```
put -c mycache hello world
```

Adds the **hello** key with a value of **world** to the **mycache** cache.

```
put -c mycache -f myfile -i 500 hola
```

Adds the **hola** key with the value from the contents of **myfile**. Also sets a maximum idle of **500** seconds.

8.22.6. SEE ALSO

get(1), remove(1)

8.23. QUERY(1)

8.23.1. NAME

query - retrieves entries that match lckle query strings.

8.23.2. SYNOPSIS

query ['OPTIONS'] **QUERY_STRING**

8.23.3. OPTIONS

-c, --cache='NAME'

Specifies the cache to query. Defaults to the currently selected cache.

--max-results='MAX_RESULTS'

Sets the number of results to return. The default is **10**.

-o, --offset='OFFSET'

Specifies the index of the first result to return. The default is **0**.

--query-mode='QUERY_MODE'

Specifies how the server executes the query. Values are **FETCH** and **BROADCAST**. The default is **FETCH**.

8.23.4. EXAMPLES

query "from org.infinispan.rest.search.entity.Person p where p.gender = 'MALE'"

Queries the currently selected cache to return entries from a Protobuf **Person** entity where the gender datatype is **MALE**.

8.23.5. SEE ALSO

schema(1)

8.24. QUIT(1)

8.24.1. NAME

quit - exits the command line interface.

8.24.2. SYNOPSIS

quit

8.24.3. EXAMPLE

quit

Exits the CLI.

8.24.4. SEE ALSO

disconnect(1), shutdown(1)

8.25. REMOVE(1)**8.25.1. NAME**

remove - deletes entries from a cache.

8.25.2. SYNOPSIS

remove **KEY** ['OPTIONS']

8.25.3. OPTIONS

--cache='NAME'

Specifies the cache from which to remove entries. Defaults to the currently selected cache.

8.25.4. EXAMPLE

remove --cache=mycache hola

Deletes the **hola** entry from the **mycache** cache.

8.25.5. SEE ALSO

cache(1), drop(1), clearcache(1)

8.26. RESET(1)**8.26.1. NAME**

reset - restores the initial values of counters.

8.26.2. SYNOPSIS

reset ['COUNTER_NAME']

8.26.3. EXAMPLE

reset cnt_a

Resets the **cnt_a** counter.

8.26.4. SEE ALSO

add(1), cas(1), drop(1)

8.27. SCHEMA(1)

8.27.1. NAME

schema - uploads and registers protobuf schemas.

8.27.2. SYNOPSIS

```
schema ['OPTIONS'] SCHEMA_NAME
```

8.27.3. OPTIONS

```
-u, --upload='FILE'
```

Uploads a file as a protobuf schema with the given name.

8.27.4. EXAMPLE

```
schema --upload=person.proto person.proto
```

Registers a **person.proto** Protobuf schema.

8.27.5. SEE ALSO

query(1)

8.28. SHUTDOWN(1)

8.28.1. NAME

shutdown - stops running servers or brings clusters down gracefully.

8.28.2. SYNOPSIS

```
shutdown server ['SERVERS']
```

```
shutdown cluster
```

8.28.3. EXAMPLES

```
shutdown server
```

Stops the server to which the CLI is connected.

```
shutdown server my_server01
```

Stops the server with hostname **my_server01**.

```
shutdown cluster
```

Stores cluster state, persists entries if you use a cache store, and stops all nodes.

8.28.4. SEE ALSO

connect(1), disconnect(1), quit(1)

8.29. SITE(1)

8.29.1. NAME

site - manages backup locations and performs cross-site replication operations.

8.29.2. SYNOPSIS

site status ['OPTIONS']

site bring-online ['OPTIONS']

site take-offline ['OPTIONS']

site push-site-state ['OPTIONS']

site cancel-push-state ['OPTIONS']

site cancel-receive-state ['OPTIONS']

site push-site-status ['OPTIONS']

8.29.3. OPTIONS

--cache='CACHE_NAME'

Specifies a cache.

--site='SITE_NAME'

Specifies a backup location.

8.29.4. EXAMPLES

site status --cache=mycache

Returns the status of all backup locations for **mycache**.

site status --cache=mycache --site=NYC

Returns the status of **NYC** for **mycache**.

site bring-online --cache=mycache --site=NYC

Brings the site **NYC** online for **mycache**.

site take-offline --cache=mycache --site=NYC

Takes the site **NYC** offline for **mycache**.

site push-site-state --cache=mycache --site=NYC

Backs up caches to remote backup locations.

site push-site-status --cache=mycache

Displays the status of the operation to backup **mycache**.

site cancel-push-state --cache=mycache --site=NYC

Cancels the operation to backup **mycache** to **NYC**.

site cancel-receive-state --cache=mycache --site=NYC

Cancels the operation to receive state from **NYC**.

site clear-push-state-status --cache=myCache

Clears the status of the push state operation for **mycache**.

8.30. STATS(1)

8.30.1. NAME

stats - displays statistics about resources.

8.30.2. SYNOPSIS

stats ['PATH']

8.30.3. EXAMPLES

stats //containers/default

Displays statistics about the default container.

stats //containers/default/caches/mycache

Displays statistics about the **mycache** cache.

8.30.4. SEE ALSO

cd(1), ls(1), describe(1)

8.31. TASK(1)

8.31.1. NAME

task - executes and uploads server-side tasks and scripts

8.31.2. SYNOPSIS

task upload --file='script' 'TASK_NAME'

task exec ['TASK_NAME']

8.31.3. EXAMPLES

task upload --file=hello.js hello

Uploads a script from a **hello.js** file and names it **hello**.

task exec @@cache@names

Runs a task that returns available cache names.

task exec hello -Pgreetee=world

Runs a script named **hello** and specifies the **greetee** parameter with a value of **world**.

8.31.4. OPTIONS

-P, --parameters='PARAMETERS'

Passes parameter values to tasks and scripts.

-f, --file='FILE'

Uploads script files with the given names.

8.31.5. SEE ALSO

ls(1)

8.32. UNALIAS(1)

8.32.1. NAME

unalias - deletes aliases.

8.32.2. SYNOPSIS

unalias 'ALIAS-NAME'

8.32.3. EXAMPLES

unalias q

Deletes the **q** alias.

8.32.4. SEE ALSO

config(1), alias(1)

8.33. USER(1)

8.33.1. NAME

user - manages Data Grid users in property security realms.

8.33.2. SYNOPSIS

user ls

user create 'username'

user describe 'username'

user remove 'username'

user password 'username'

user groups 'username'

user encrypt-all

8.33.3. DESCRIPTION

Manage (list, create, describe, remove, modify) users stored in a property security realm. Note: You can use this command only with property realms.

8.33.4. COMMAND SYNOPSIS

user ls

Lists the users or groups which are present in the property file.

user create 'username'

Creates a user after prompting for a password.

user describe 'username'

Describes a user, including its username, realm and any groups it belongs to.

user remove 'username'

Removes the specified user from the property file.

user password 'username'

Changes the password for a user.

user groups 'username'

Sets the groups to which a user belongs.

user encrypt-all

Encrypt all passwords in a plain-text user property file.

8.33.5. COMMON OPTIONS

These options apply to all commands:

-h, --help

Displays a help page for the command or sub-command.

-s, --server-root='path-to-server-root'

The path to the server root. Defaults to **server**.

-f, --users-file='users.properties'

The name of the property file which contains the user passwords. Defaults to **users.properties**.

-w, --groups-file='groups.properties'

The name of the property file which contains the user to groups mapping. Defaults to **groups.properties**.

8.33.6. USER CREATE/MODIFY OPTIONS

-a, --algorithms

Specifies the algorithms used to hash the password.

-g, --groups='group1,group2,...'

Specifies the groups to which the user belongs.

-p, --password='password'

Specifies the user's password.

-r, --realm='realm'

Specifies the realm name.

--plain-text

Whether passwords should be stored in plain-text (not recommended).

8.33.7. USER LS OPTIONS

--groups

Shows a list of groups instead of the users.

8.33.8. USER ENCRYPT-ALL OPTIONS

-a, --algorithms

Specifies the algorithms used to hash the password.

8.34. VERSION(1)

8.34.1. NAME

version - displays the server version and CLI version.

8.34.2. SYNOPSIS

version

8.34.3. EXAMPLE

version

Returns the version for the server and the CLI.

8.34.4. SEE ALSO

help(1)