



Red Hat Data Grid 8.3

Data Grid Code Tutorials

Learn how to use Data Grid capabilities

Red Hat Data Grid 8.3 Data Grid Code Tutorials

Learn how to use Data Grid capabilities

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

Run code tutorials for remote caches and embedded caches that demonstrate various Data Grid capabilities and usage patterns.

Table of Contents

RED HAT DATA GRID	3
DATA GRID DOCUMENTATION	4
DATA GRID DOWNLOADS	5
MAKING OPEN SOURCE MORE INCLUSIVE	6
CHAPTER 1. REMOTE CACHES	7
1.1. REMOTE CACHE TUTORIALS	7
1.2. HOT ROD JAVA CLIENT TUTORIALS	7
Data Grid documentation	8
CHAPTER 2. EMBEDDED CACHES	9
2.1. EMBEDDED CACHE TUTORIALS	9
Data Grid documentation	9
CHAPTER 3. SPRING AND SPRING BOOT	11
3.1. SPRING AND SPRING BOOT TUTORIALS	11
Data Grid documentation	12

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 8.3 Documentation](#)
- [Data Grid 8.3 Component Details](#)
- [Supported Configurations for Data Grid 8.3](#)
- [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. REMOTE CACHES

Deploy multiple Data Grid Server instances to create remote cache clusters that give you a fault-tolerant and scalable data tier with high-speed access from Hot Rod and REST clients.

1.1. REMOTE CACHE TUTORIALS

To run these tutorials you need at least one locally running instance of Data Grid Server.

You can [download](#) the distribution and run the following commands:

```

$ ./bin/cli.sh user create admin -p "password"
$ ./bin/server.sh

```



NOTE

Data Grid Server enables authentication and authorization by default. Creating a user named **admin** gives you administrative access to Infinispan Server.

Building and running remote cache tutorials

You can build and run remote cache tutorials directly in your IDE or from the command line as follows:

```
$ mvn -s /path/to/maven-settings.xml clean package exec:exec
```

1.2. HOT ROD JAVA CLIENT TUTORIALS

- Hot Rod Java clients require JDK 8 or later. However, Data Grid recommends using Java 11 at a minimum.

Tutorial link	Description
Remote cache use example	The simplest code example that demonstrates how a remote distributed cache works.
Per cache configuration	Demonstrates how to configure caches dynamically when we connect to the Data Grid Server.
Near caches	Demonstrates how configure near caching to improve the read performance in remote caches.
Cache Admin API	Demonstrates how to use the Administration API to create caches and cache templates dynamically.
Encoding	Demonstrates how encoding of caches work.
Client listeners	Detect when data changes in a remote cache with Client Listeners.

Tutorial link	Description
Query	Demonstrates how to query remote cache values.
Continuous query	Demonstrates how to use Continuous Query and remote caches.
Transactions	Demonstrates how remote transactions work.
Secured caches	Demonstrates how to configure caches that have authorization enabled.
TLS authorization	Demonstrates how to connect to Data Grid Server with TLS authorization.
Counters	Demonstrates how remote counters work.
Multimap	Demonstrates how remote multimap works.
Task execution	Demonstrates how to register server tasks and how to execute them from the Hot Rod client.
JUnit 5 and Testcontainers	Demonstrates how to use the Data Grid and JUnit 5 extension.
Persistence	Demonstrates how to use the Data Grid and persistent caches.

Data Grid documentation

You can find more resources for Hot Rod Java clients in our documentation at:

- [Hot Rod Java client guide](#)
- [Marshalling and Encoding Data Guide](#)
- [Querying Data Grid caches](#)
- [Developer guide](#)
- [REST API](#)

CHAPTER 2. EMBEDDED CACHES

Add Data Grid as a dependency to your Java project and use embedded caches that increase application performance and give you capabilities to handle complex use cases.

2.1. EMBEDDED CACHE TUTORIALS

You can run embedded cache tutorials directly in your IDE or from the command line as follows:

```
$ mvn -s /path/to/maven-settings.xml clean package exec:exec
```

Tutorial link	Description
Distributed caches	Demonstrates how Distributed Caches work.
Replicated caches	Demonstrates how Replicated Caches work.
Invalidated caches	Demonstrates how Invalidated Caches work.
Transactions	Demonstrates how transactions work.
Streams	Demonstrates how Distributed Streams work.
JCache integration	Demonstrates how JCache works.
Functional Maps	Demonstrates how Functional Map API works.
Map API	Demonstrates how the Map API works with Data Grid caches.
Multimap	Demonstrates how to use Multimap.
Queries	Uses Data Grid Query to perform full-text queries on cache values.
Clustered Listeners	Detects when data changes in an embedded cache with Clustered Listeners.
Counters	Demonstrates how to use an embedded Clustered Counter.
Clustered Locks	Demonstrates how to use an embedded Clustered Lock.
Clustered execution	Demonstrates how to use an embedded Clustered Counter.


[Data Grid documentation](#)

You can find more resources about embedded caches in our documentation at:

- [Embedding Data Grid Caches](#)
- [Data Grid Developer Guide](#)
- [Querying Data Grid caches](#)

CHAPTER 3. SPRING AND SPRING BOOT

3.1. SPRING AND SPRING BOOT TUTORIALS



NOTE

These code tutorials use Data Grid Server and require at least one running instance.

Run Spring examples

```
$ mvn -s /path/to/maven-settings.xml package exec:exec
```

Run Spring Boot examples

```
$ mvn -s /path/to/maven-settings.xml spring-boot:run
```

Displaying actuator statistics

Navigate to <http://localhost:8080/actuator/metrics> in your browser to display a list of available metrics. Cache metrics are prefixed with "cache." Display each metric for each cache using tags. For example for the 'puts' stats in the basque-names cache:

<http://localhost:8080/actuator/metrics/cache.puts?tag=name:basque-names>

Collecting statistics with Prometheus

The `prometheus.yml` file in this project contains a `host.docker.internal` binding that allows Prometheus to scrap metrics that the Spring actuator exposes.

Change the **YOUR_PATH** value in the following command to the directory where Prometheus is running and then run:

Podman

```
$ podman run -d --name=prometheus -p 9090:9090 -v YOUR_PATH/integrations/spring-
boot/prometheus.yml:/etc/prometheus/prometheus.yml prom/prometheus --
config.file=/etc/prometheus/prometheus.yml
```

Tutorial link	Description
Spring Boot and Spring Cache remote mode	Demonstrates how to use Spring Caches with Spring Boot and the Data Grid Server.
Spring Boot and Spring Session remote mode	Demonstrates how to use Spring Session with Spring Boot and the Data Grid Server.
Spring Boot and Spring Cache embedded mode	Demonstrates how to use Spring Caches with Spring Boot and Data Grid Embedded.

Tutorial link	Description
Spring Boot and Spring Session embedded mode	Demonstrates how to use Spring Session with Spring Boot and Data Grid Embedded.
Spring cache embedded without Spring Boot	Demonstrates how to use Spring Cache and Data Grid Embedded without Spring Boot.
Spring session embedded without Spring Boot	Demonstrates how to use Spring Session and Data Grid Embedded without Spring Boot.

Data Grid documentation

You can find more resources in our documentation at:

- [Using Data Grid with Spring](#)
- [Data Grid Spring Boot Starter](#)