



# Red Hat Data Grid 7.3

## Red Hat Data Grid 7.3 Feature Support

Data Grid Documentation





## Legal Notice

Copyright © 2019 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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 feature support information for Data Grid 7.3.

---

## Table of Contents

<b>1. DATA GRID</b> .....	<b>2</b>
1.1. Data Grid Usage Modes	2
<b>2. SUPPORTED DATA GRID FEATURES</b> .....	<b>2</b>
2.1. Supported Features in Data Grid 7.3	2
2.2. Features in Remote Client-Server Mode (or Server Mode)	4
2.3. Features in Library Mode (or Embedded Mode)	5
2.4. Technology Previews	5
2.5. Deprecated Features	5
<b>3. SUPPORTED HOT ROD FEATURES</b> .....	<b>5</b>

## 1. 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.

### 1.1. Data Grid Usage Modes

Red Hat Data Grid supports two usage modes:

#### Remote Client-Server Mode (or Server Mode)

Run Data Grid as managed and distributed cluster that you can scale to meet demand. Client applications remotely access Data Grid through endpoints such as **Hot Rod** or **REST**.

#### Library Mode (or Embedded Mode)

Build and deploy Data Grid as a single node in your application runtime, where the application process and cache share the same JVM memory. In Library Mode, Data Grid can remotely access nodes hosted in other JVMs across your environment.

Data Grid supports Library Mode in application containers such as JBoss Enterprise Application Platform (EAP) or as standalone applications. Find out more about supported application containers in the [Supported Configurations](#).

## 2. SUPPORTED DATA GRID FEATURES

### 2.1. Supported Features in Data Grid 7.3

- File Cache Store and Loading
- JDBC Cache Store and Loading
- RocksDB Cache Store and Loading (Replaces LevelDB)
- Cache Passivation
- Remote Cache Store
- Cluster Cache Store
- Asynchronous Store

- Cluster Configuration Using UDP
- Cluster Configuration Using TCP
- Mortal and Immortal Data
- Eviction Strategy
- Expiration
- Unscheduled Write-behind Cache Store
- Write-through Cache Store
- Local Cache Mode (local)
- Clustering Mode (replicated)
- Clustering Mode (invalidation)
- Clustering Mode (distribution)
- Asynchronous Clustering Modes
- Marshalling
- JMX Management
- Cross-Datacenter Replication and State Transfer
- Role-based Access Control
- Node Authentication and Authorization
- Encrypted Communication Within the Cluster
- Per Invocation Flags
- Customizable Network Partition Handling
- Spring Integration
- Querying (by values)
- Continuous Queries
- Clustered Listeners and Notifications for Cache Events
- Near Caching
- JSR-107 Support
- CDI
- Asynchronous API
- Distributed Streams

Distributed Streams are available in Remote Client-Server Mode via Remote Task Execution.

- Off Heap Cache Storage
- Ickle Query Language
- EAP Integration
- Cassandra Cache Store and Loading
- Command Line Interface (CLI)
- Apache Camel Component for JBoss Fuse



#### NOTE

Red Hat Fuse 7.3 and later provide a **camel-infinispan** component that replaces the Data Grid Apache Camel component, **jboss-datagrid-7.3-camel-library**.

You should download and use the **camel-infinispan** component that Red Hat Fuse provides instead of **jboss-datagrid-7.3-camel-library**.

Refer to the *Infinispan Component* section in the *Apache Camel Component Reference* document for more information. You can find the *Apache Camel Component Reference* in the [Red Hat Fuse documentation](#).

## 2.2. Features in Remote Client-Server Mode (or Server Mode)

- Deploy custom cache store to RHDG Server
- Datasources with JDBC Cache Stores
- REST Interface
- Hot Rod Java client
- Hot Rod C++ Client
- Hot Rod .NET Client
- Hot Rod Node.js Client
- Data Compatibility Between Client-server Protocols
- Data Compatibility Between Hot Rod Java and C++ Client
- Rolling Upgrades for Hot Rod Cluster
- Controlled Shutdown and Restart of Cluster
- Hot Rod Client as a JBoss EAP Module
- Externalizing HTTP sessions from JBoss EAP 7 to remote RHDG cluster
- Externalizing HTTP sessions from JBoss Web Server to remote RHDG cluster
- Remote Task Execution
- Apache Spark with Scala 2.11



- Apache Hadoop Integration
- Administration Console

### 2.3. Features in Library Mode (or Embedded Mode)

- READ\_COMMITTED and REPEATABLE\_READ Isolation Modes
- Lazy Deserialization
- Configuration format (**infinispan.xml**) is shared between Library Mode and Remote Client-Server Mode.
- Grouping API
- Java Transactional API (JTA) Support and Configuration
- Java Transactional API (JTA) Deadlock Detection
- Transaction Recovery
- Transaction and Batching
- Key Affinity
- RHDG as Lucene Directory
- JPA Cache Store
- Data Grid as a JBoss EAP Module



#### NOTE

Data Grid now provides Library and Java Hot Rod Client modules for EAP as a single package.

### 2.4. Technology Previews

[Technology Preview Features in Data Grid 7.3](#)

### 2.5. Deprecated Features

[Features and Functionality Deprecated in Data Grid 7.3](#)

## 3. SUPPORTED HOT ROD FEATURES

Table 1. Hot Rod Client Feature Support by Language

	Java	C++	C#	Node.js
--	------	-----	----	---------

	Java	C++	C#	Node.js
Transactions	Supported with <b>FULL_XA</b> and <b>NON_XA</b> transaction modes.	Supported with <b>NON_XA</b> transactions only.	Supported with <b>NON_XA</b> transactions only.	N/A
Clustered Counters	Supported	Tech Preview	Tech Preview	N/A
Statistics via JMX	Supported	N/A	N/A	N/A
Cross-site Failover	Supported	Supported	Supported	Supported
Authentication	Supported	Supported	Supported	Supported
TLS-based Encryption	Supported	Supported	Supported	Supported
Server Name Indication (SNI)	Supported	Supported	Supported	Supported
Asynchronous API	Supported	Supported	Supported	N/A
Near Caching	Supported	Supported	Supported	N/A
Continuous Queries	Supported	Supported	Supported	N/A
Remote Event Listeners	Supported	Supported	Supported	Supported
Remote Querying	Supported	Supported	Supported	N/A
Remote Execution	Supported	Supported	Supported	N/A