



Red Hat Data Grid 8.3

Data Grid Operator 8.3 Release Notes

Get release information for Data Grid Operator 8.3

Red Hat Data Grid 8.3 Data Grid Operator 8.3 Release Notes

Get release information for Data Grid Operator 8.3

Legal Notice

Copyright © 2022 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

Find out about features and enhancements in Data Grid Operator 8.3 as well as known and resolved issues.

Table of Contents

RED HAT DATA GRID	3
DATA GRID DOCUMENTATION	4
DATA GRID DOWNLOADS	5
MAKING OPEN SOURCE MORE INCLUSIVE	6
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	7
CHAPTER 1. DATA GRID OPERATOR 8.3	8
1.1. DATA GRID OPERATOR 8.3.7	8
Disabling FIPS mode	8
Enhanced bidirectional reconciliation	8
Updated cluster view response	8
Validating and mutating webhooks	8
1.2. DATA GRID OPERATOR 8.3.4	8
Hot Rod rolling upgrades with Data Grid Operator	8
User-defined annotations for created resources	8
Performance improvement that avoids unnecessary iptable rules	8
OpenShift route as a expose type for cross-site replication	8
Data Grid Operator resource reconciliation	9
Configuration enhancements	9
1.3. DATA GRID OPERATOR 8.3.1	9
Updated OpenShift bundle for Data Grid Operator deployments	9
Enhancements to memory and CPU requests and limits	9
Readiness and liveness probes updated for better performance	9
1.4. DATA GRID OPERATOR 8.3 GA	9
Numerous improvements to the Data Grid Operator code base	10
Full support for native CLI installation and operation	10
Custom Data Grid Server configuration	10
Configurable number of relay nodes for cross-site replication	10
TLS security for cross-site connections	10
Cache service type deprecation	10
1.5. DATA GRID OPERATOR 8.3.X RELEASE INFORMATION	11
CHAPTER 2. KNOWN AND FIXED ISSUES	12
2.1. KNOWN ISSUES WITH DATA GRID OPERATOR DEPLOYMENTS	12
2.2. FIXED IN DATA GRID OPERATOR 8.3.0	12
2.3. FIXED IN DATA GRID OPERATOR 8.3.1	12
2.4. FIXED IN DATA GRID OPERATOR 8.3.4	12
2.5. FIXED IN DATA GRID OPERATOR 8.3.6	13
2.6. FIXED IN DATA GRID OPERATOR 8.3.7	13
CHAPTER 3. DATA GRID ON OPENSHIFT	14
3.1. DATA GRID 8.3 IMAGES	14
Custom Data Grid Deployments	14
3.2. EMBEDDED CACHES ON OPENSHIFT	14

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](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our technical content and encourage you to tell us what you think. If you'd like to add comments, provide insights, correct a typo, or even ask a question, you can do so directly in the documentation.



NOTE

You must have a Red Hat account and be logged in to the customer portal.

To submit documentation feedback from the customer portal, do the following:

1. Select the **Multi-page HTML** format.
2. Click the **Feedback** button at the top-right of the document.
3. Highlight the section of text where you want to provide feedback.
4. Click the **Add Feedback** dialog next to your highlighted text.
5. Enter your feedback in the text box on the right of the page and then click **Submit**.

We automatically create a tracking issue each time you submit feedback. Open the link that is displayed after you click **Submit** and start watching the issue or add more comments.

Thank you for the valuable feedback.

CHAPTER 1. DATA GRID OPERATOR 8.3

Get version details for Data Grid Operator 8.3 and information about issues.

1.1. DATA GRID OPERATOR 8.3.7

What's new in 8.3.7.

Disabling FIPS mode

You can disable FIPS mode in your Infinispan CR, so that any component that uses a JVM can ignore FIPS mode. This means that your Data Grid instance can run on any OpenShift cluster that has FIPS mode enabled.

Enhanced bidirectional reconciliation

Data Grid Operator 8.3.7 improves the bidirectional reconciliation mechanism, so that it includes idempotent creation and retrieval capabilities. The idempotent capability aligns bidirectional reconciliation with Kubernetes conventions.

See [Resource reconciliation](#) (Data Grid Operator)

Updated cluster view response

The Data Grid Operator 8.3.7 can now update the **status** field to **False** when it cannot create a cache from a **Cache** custom resource (CR).

Validating and mutating webhooks

Data Grid Operator 8.3.7 includes validating and mutating webhooks. A webhook changes the behavior of the Kubernetes API, so that this API can immediately reject any CR with invalid specifications.

1.2. DATA GRID OPERATOR 8.3.4

What's new in 8.3.4.

Hot Rod rolling upgrades with Data Grid Operator

From version 8.3.4 onward, you can upgrade Data Grid clusters when new versions become available without service downtime and data loss. The shutdown upgrade remains the default upgrade type. For information on how to enable Hot Rod rolling upgrades see [Upgrading Data Grid clusters](#).



IMPORTANT

Hot Rod rolling upgrades are available as a technology preview feature. For more information, see [Red Hat Technology Preview Features Support Scope](#).

User-defined annotations for created resources

This release introduces configurable annotations that you can use to organize and monitor Data Grid resources.

Performance improvement that avoids unnecessary iptable rules

The Data Grid Operator endpoint is now available as a headless service. Headless services avoid the creation of several **iptables** rules that can impact the performance of the cluster.

OpenShift route as a expose type for cross-site replication

This release lets you use OpenShift **Route** as a new expose type for cross-site replication. You can use an OpenShift **Route** to handle network traffic for backup operations between Data Grid clusters.

Data Grid Operator resource reconciliation

Data Grid Operator 8.3.4 introduces bidirectional reconciliation that enables syncing Data Grid caches with OpenShift Custom Resource (CR) definitions. Bidirectional reconciliation synchronizes your CRs with changes that you make to Data Grid resources through the Data Grid Console, command line interface (CLI), or other client application and vice versa.

To perform bidirectional reconciliation, you must have **configListener** enabled in the Data Grid service.

Configuration enhancements

When configuring cross-site replication, you no longer need to specify locations in the Data Grid CR. Data Grid Operator configures Data Grid services without site locations configured.

1.3. DATA GRID OPERATOR 8.3.1

What's new in 8.3.1.

Updated OpenShift bundle for Data Grid Operator deployments

The OpenShift bundle for Data Grid Operator deployments includes images built with OpenJDK 11 to support multiple system architectures including:

- x86 (x86_64)
- s390x (IBM Z)
- ppc64le (IBM Power Systems)

OpenJ9 deprecation

Any **OpenJ9** images for IBM Z and IBM Power Systems will be deprecated.

For more information see [Java Change in Power and Z OpenShift Images](#) .

Enhancements to memory and CPU requests and limits

This release increases the default values for memory and CPU resources that Data Grid Operator requests from the OpenShift scheduler when creating Data Grid pods.

- **1Gi** of memory.
- CPU requests are unbounded.

Additionally you can now allocate resources in the format of **<limit>:<requests>** in your **Infinispan** CR as follows:

```
spec:
  container:
    cpu: "2000m:1000m"
    memory: "2Gi:1Gi"
```

Readiness and liveness probes updated for better performance

Settings for the readiness and liveness probes are updated so that Data Grid clusters become available sooner.

1.4. DATA GRID OPERATOR 8.3 GA

What's new in 8.3 GA.

Numerous improvements to the Data Grid Operator code base

The Data Grid team have invested a lot of time and effort into internal improvements to the Data Grid Operator code base to enable future enhancements.

Full support for native CLI installation and operation

To improve the installation and operation experience, Data Grid 8.3 offers full support for the native CLI as an **oc** client plugin. Installing the native CLI extends your **oc** client with the following commands:

Command	Description
oc infinispn install	Creates Data Grid Operator subscriptions and installs into the global namespace by default.
oc infinispn create cluster	Creates Data Grid clusters.
oc infinispn get clusters	Displays running Data Grid clusters.
oc infinispn shell	Starts an interactive remote shell session on a Data Grid cluster.
oc infinispn delete cluster	Removes Data Grid clusters.
oc infinispn uninstall	Removes Data Grid Operator installations and all managed resources.

Custom Data Grid Server configuration

You can now add Data Grid Server configuration to a **ConfigMap** and make it available to Data Grid Operator. Data Grid Operator can then apply the custom configuration to your Data Grid cluster.

Configurable number of relay nodes for cross-site replication

Data Grid clusters now use router pods, which are GossipRouter instances, to coordinate RELAY messages for cross-site replication. You can also configure the number of pods in each cluster that can send RELAY messages to router pods with the **sites.local.maxRelayNodes** field.

TLS security for cross-site connections

Add keystores, and optional trust stores, to encrypt RELAY messages and secure cross-site replication traffic between Data Grid clusters.

Cache service type deprecation

The **Cache** service type was designed to provide a convenient way to create a low-latency data store with minimal configuration. Additional development on the **Infinispn** CRD has shown that the **Cache** CR offers a better approach to achieving this goal, ultimately giving users more choice and less deployment overhead. For this reason, the **Cache** service type is planned for removal in the next version of the **Infinispn** CRD and is no longer under active development.

Red Hat recommends configuring the **DataGrid** service type for clusters. The **DataGrid** service type continues to benefit from new features and improved tooling to automate complex operations such as cluster upgrades and data migration.

You can create **DataGrid** service clusters as follows:

- Set **spec.service.type: DataGrid** in your **Infinispn** CR.

- Use the `-Pservice.type=DataGrid` argument with the native CLI plugin.

1.5. DATA GRID OPERATOR 8.3.X RELEASE INFORMATION

The following table provides detailed version information for Data Grid Operator.



NOTE

Data Grid Operator versions do not always directly correspond to Data Grid versions because the release schedule is different.

Data Grid Operator version	Data Grid version	Features
8.3.8	8.3.1	Fixes security vulnerabilities.
8.3.7	8.3.1	Includes several bug fixes. For new features, see Data Grid Operator 8.3.7 .
8.3.6	8.3.1	Includes several bug fixes.
8.3.5	8.3.1	Fixes security vulnerabilities.
8.3.4	8.3.1	See Data Grid Operator 8.3.4
8.3.3	8.3.0	Fixes security vulnerabilities.
8.3.2	8.3.0	Fixes security vulnerabilities.
8.3.1	8.3.0	<ul style="list-style-type: none"> * Updates bundle for Data Grid Operator deployments. * Increases default memory and CPU limits. * Adds separate settings for requests and limits to the Infinispan CRD. * Updates readiness and liveness probes for better performance.
8.3.0	8.3.0	See Data Grid Operator 8.3 GA

CHAPTER 2. KNOWN AND FIXED ISSUES

Learn about known issues for Data Grid Operator and find out which issues are fixed.

2.1. KNOWN ISSUES WITH DATA GRID OPERATOR DEPLOYMENTS

This section describes issues that affect Data Grid clusters that you manage with Data Grid Operator. For complete details about Data Grid, you should refer to the [Data Grid 8.3 release notes](#).

Unexpected behavior occurs after OOM exceptions

Issue: [JDG-3991](#)

Description: If out of memory exceptions cause Data Grid Server to terminate on OpenShift, unexpected behavior can occur. In some cases nodes cannot restart and the **org.infinispan.LOCKS** enters degraded mode. The following exception is written to the pod log file:

```
FATAL (main) [org.infinispan.SERVER] ISPN080028: Red Hat Data Grid Server failed to start
java.util.concurrent.ExecutionException:
org.infinispan.manager.EmbeddedCacheManagerStartupException:
org.infinispan.commons.CacheException: Initial state transfer timed out for cache
org.infinispan.LOCKS on <pod-name-id>
```

Workaround: There is no workaround for this issue. You should configure eviction in Data Grid caches to help avoid OOM exceptions.

2.2. FIXED IN DATA GRID OPERATOR 8.3.0

Data Grid Operator 8.3.0 includes the following notable fixes:

- [JDG-4682](#) Removing **Infinispan** CR deletes user created secrets
- [JDG-4763](#) Clients cannot connect to remote caches that use TLS/SSL encryption
- [JDG-4572](#) Native CLI **oc plugin infinispan delete** command does not fail without subcommand
- [JDG-4568](#) Native CLI **oc plugin** does not interpret arrow keys correctly
- [JDG-4574](#) Native CLI **oc plugin install** command does not work
- [JDG-5026](#) Cluster is unable to start after graceful shutdown

2.3. FIXED IN DATA GRID OPERATOR 8.3.1

Data Grid Operator 8.3.1 includes the following notable fixes:

- [JDG-5078](#) Ensure compatibility with sidecar injection

2.4. FIXED IN DATA GRID OPERATOR 8.3.4

Data Grid Operator 8.3.4 includes the following notable fixes:

- [JDG-5171](#) Operator version is undefined

- [JDG-5108](#) Data Grid restore procedure fails when the Data Grid Operator receives an upgrade
- [JDG-5119](#) Deadlock occurs when deleting Infinispan CR with propagationPolicy = foreground
- [JDG-4304](#) Data Grid Operator does not reflect EndpointEncryption field changes

2.5. FIXED IN DATA GRID OPERATOR 8.3.6

Data Grid Operator 8.3.6 includes the following notable fixes:

- [JDG-5393](#) Operator exposes clusterwide ServiceAccount privileges
- [JDG-5370](#) Graceful shutdown fails with cross-site replication

2.6. FIXED IN DATA GRID OPERATOR 8.3.7

Data Grid Operator 8.3.7 includes the following notable fixes:

- [JDG-5425](#) Specifying multiple authorization permmissions for a role might lead to server startup failure.
- [JDG-5410](#) Specifying `tls.key` and `tls.crt` keys for a certificate secret causes server startup failure.
- [JDG-5405](#) Initiating a graceful shutdown of a cluster with no pods causes issues with the Pod Status diagram on OpenShift.

CHAPTER 3. DATA GRID ON OPENSHIFT

3.1. DATA GRID 8.3 IMAGES

Data Grid 8.3 includes two container images, the Data Grid Operator image and Data Grid Server image.

Data Grid images are hosted on the Red Hat Container Registry, where you can find health indexes for the images along with information about each tagged version.

Custom Data Grid Deployments

Red Hat does not support customization of any 8.3 images from the Red Hat Container Registry through the Source-to-Image (S2I) process or **ConfigMap** API.

As a result it is not possible to use custom:

- Discovery protocols
- JGroups **SYM_ENCRYPT** or **ASYM_ENCRYPT** encryption mechanisms

Additional resources

- [Data Grid Container Images](#)

3.2. EMBEDDED CACHES ON OPENSHIFT

Using embedded Data Grid caches in applications running on OpenShift, which was referred to as Library Mode in previous releases, is intended for specific uses only:

- Using local or distributed caching in custom Java applications to retain full control of the cache lifecycle. Additionally, when using features that are available only with embedded Data Grid such as distributed streams.
- Reducing network latency to improve the speed of cache operations.

The Hot Rod protocol provides near-cache capabilities that achieve equivalent performance to a standard client-server architecture.

Requirements

Embedding Data Grid in applications running on OpenShift requires you to use a discovery mechanism so Data Grid nodes can form clusters to replicate and distribute data.

Red Hat supports only `DNS_PING` as the cluster discovery mechanism.

`DNS_PING` exposes a port named **ping** that Data Grid nodes use to perform discovery and join clusters. TCP is the only supported protocol for the **ping** port, as in the following example for a pod on OpenShift:

```
spec:
  ...
  ports:
    - name: ping
```

port: 8888
protocol: TCP
targetPort: 8888

Limitations

Embedding Data Grid in applications running on OpenShift also has some specific limitations:

- Persistent cache stores are not currently supported.
- UDP is not supported with embedded Data Grid.

Custom caching services

Red Hat highly discourages embedding Data Grid to build custom caching servers to handle remote client requests. To benefit from regular, automatic updates with performance improvements and fix security issues, you should create Data Grid clusters with the Data Grid Operator instead.

Additional resources

- [Embedding Data Grid in Java Applications](#)