



Red Hat Data Grid 8.4

Data Grid Operator 8.4 Release Notes

Get release information for Data Grid Operator 8.4

Red Hat Data Grid 8.4 Data Grid Operator 8.4 Release Notes

Get release information for Data Grid Operator 8.4

Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

Find out about features and enhancements in Data Grid Operator 8.4 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
CHAPTER 1. DATA GRID OPERATOR 8.4	7
1.1. DATA GRID OPERATOR 8.4.14	7
Enhancements to Gossip Router	7
1.2. DATA GRID OPERATOR 8.4.8	7
Improved monitoring with Cryostat and JFR recordings	7
ServiceMonitor target labels	7
Enhanced cross-site replication and Gossip Router configuration	7
Scheduling for Data Grid pods	7
1.3. DATA GRID OPERATOR 8.4.5	8
Allocate CPU and memory to ConfigListener	8
1.4. DATA GRID OPERATOR 8.4.2	8
Improved security for storing sensitive strings	8
Configuration enhancements to ConfigListener log level	8
Strategies to update Cache CR	9
1.5. DATA GRID OPERATOR 8.4 GA	9
Multi-operand support	9
Updated OpenShift bundle for Data Grid Operator deployments	9
Option to disable Gossip router	9
Federal Information Processing Standards (FIPS) compatibility	10
1.6. DATA GRID OPERATOR 8.4.X RELEASE INFORMATION	10
CHAPTER 2. KNOWN AND FIXED ISSUES	14
2.1. KNOWN ISSUES WITH DATA GRID OPERATOR DEPLOYMENTS	14
2.2. FIXED IN DATA GRID OPERATOR 8.4.14	14
2.3. FIXED IN DATA GRID OPERATOR 8.4.12	14
2.4. FIXED IN DATA GRID OPERATOR 8.4.11	14
2.5. FIXED IN DATA GRID OPERATOR 8.4.9	14
2.6. FIXED IN DATA GRID OPERATOR 8.4.8	14
2.7. FIXED IN DATA GRID OPERATOR 8.4.6	15
2.8. FIXED IN DATA GRID OPERATOR 8.4.5	15
2.9. FIXED IN DATA GRID OPERATOR 8.4.2	15
2.10. FIXED IN DATA GRID OPERATOR 8.4.0	16
CHAPTER 3. DATA GRID ON OPENSIFT	17
3.1. DATA GRID 8.4 IMAGES	17
Custom Data Grid Deployments	17
3.2. EMBEDDED CACHES ON OPENSIFT	17

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 8.4 Documentation](#)
- [Data Grid 8.4 Component Details](#)
- [Supported Configurations for Data Grid 8.4](#)
- [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. DATA GRID OPERATOR 8.4

Get version details for Data Grid Operator 8.4 and information about issues.

1.1. DATA GRID OPERATOR 8.4.14

What's new in 8.4.14.

Enhancements to Gossip Router

You can now configure the idle connection timeout in the Gossip Router, allowing the Gossip Router to force-close orphan TCP connection in case a Data Grid pod or a Gossip Router pod stops responding.

The Data Grid pod pings the Gossip Router at intervals defined in milliseconds. If Data Grid pod does not receive any message from the Gossip Router during the defined timeout period, Data Grid considers the connection orphan. The Gossip Router pod closes the connection if it does not receive a message or a ping from the Data Grid pod during the timeout period.

1.2. DATA GRID OPERATOR 8.4.8

What's new in 8.4.8.

Improved monitoring with Cryostat and JFR recordings

Data Grid now provides integration with Cryostat, that lets you monitor your Data Grid clusters that run on OpenShift using the JDK Flight Recorder (JFR). You can store and analyze your recordings using the integrated tools provided by Cryostat or export the recordings to an external monitoring application.

For more information, see [Setting up JFR recordings with Cryostat](#).

ServiceMonitor target labels

Data Grid Operator now lets you configure target labels for the **ServiceMonitor**. Use the Service labels to filter and aggregate the metrics collected from the monitored endpoints.

For more information see *Data Grid Operator Guide* [Configuring Service Monitor Target Labels](#).

Enhanced cross-site replication and Gossip Router configuration

Several enhancements has been made to improve experience with Data Grid cross-site replication. With Data Grid Operator, you can configure the following:

- Allocate CPU and memory resources to tune Gossip Router pod
- Adjust the Gossip Router probe time to handle slower start up times
- Enable heartbeats between Data Grid pods and the Gossip Router pods to ensure that connections stay open
- Disable suspect events to rely only on heartbeats for failure detection to prevent unnecessary view change in case of temporary disconnections

For more information, see [Allocating CPU and memory for Gossip router pod](#).

Scheduling for Data Grid pods

You can now assign priority to the Data Grid pods to ensure that important workflows are scheduled first.

Use the **spec.scheduling** field, to assign priority to Data Grid pods.

Infinispan CR

```

kind: Infinispan
...
spec:
  scheduling:
    affinity:
      ...
    priorityClassName: "high-priority"
  ...

```

1.3. DATA GRID OPERATOR 8.4.5

What's new in 8.4.5.

Allocate CPU and memory to ConfigListener

You can allocate memory and CPU resources to **ConfigListener**. Although **ConfigListener** consumes minimal resources by default, setting limits might be required for some Data Grid deployments.

The **cpu** and **memory** fields have values in the format of **<limit>:<requests>**.

ConfigListener configuration

```

spec:
  configListener:
    enabled: true
    cpu: "2000m:1000m"
    memory: "2Gi:1Gi"

```

1.4. DATA GRID OPERATOR 8.4.2

What's new in 8.4.2.

Improved security for storing sensitive strings

Data Grid Operator lets you securely define and store custom Data Grid Server configuration. To protect sensitive text strings, such as passwords, add the entries in a credential store rather than directly in the Data Grid Server configuration.

For more information see *Data Grid Operator Guide* [Securing custom Data Grid configuration](#).

Configuration enhancements to ConfigListener log level

Data Grid Operator now lets you modify the log level of the **ConfigListener** pod. The default log level is **info**. You can change the log level to **debug** or **error**, if needed.

You can configure the log level of **ConfigListener** in your **Infinispan** CR as follows:

Infinispan CR

```

spec:
  configListener:
    enabled: true
    logging:
      level: info

```

Strategies to update Cache CR

Data Grid Operator lets you control how the **Cache** CR controller handles updates to the cache configuration in the **spec.template** field.

Cache CR update strategies

retain

Data Grid Operator updates the cache configuration on the Data Grid Server at runtime. This is the default strategy. If the update fails, the Operator updates the status of **Cache** CR to **Ready=False**.

recreate

Data Grid Operator updates the cache configuration on the Data Grid Server at runtime. If the update fails, the Operator deletes the cache from the Data Grid cluster and creates a new cache with the latest **spec.template** value.

Define the **Cache** CR update strategy as follows:

Cache CR

```
spec:
  updates:
    strategy: retain
```

1.5. DATA GRID OPERATOR 8.4 GA

What's new in 8.4 GA.

Multi-operand support

You can have a single Data Grid Operator installation that supports multiple versions of Data Grid.

Use the **spec.version** field in your **Infinispan** CR to upgrade between supported versions.

Data Grid Operator version 8.4 supports the following Data Grid versions:

- 8.4.0-1
- 8.3.1-1


For more information on upgrading Data Grid clusters see *Upgrading Data Grid clusters* in the [Data Grid Operator Guide](#).

Updated OpenShift bundle for Data Grid Operator deployments

The OpenShift bundle for Data Grid Operator deployments includes images designed for 64-bit ARM architecture.

Option to disable Gossip router

Data Grid Operator starts a Gossip router on each site, but you only need a single Gossip router to manage traffic between the Data Grid cluster members.



IMPORTANT

Although disabling Gossip router can save some resources, Data Grid recommends to keep the Gossip router enabled on all the remote sites. If you have a single Gossip router defined, and it becomes unavailable, the connection between the remote sites breaks.

For more information see *Disabling local Gossip router and service* in the [Data Grid Operator Guide](#).

Federal Information Processing Standards (FIPS) compatibility

From now you can run Data Grid instance on a FIPS-enabled OpenShift cluster without changing your configuration. Data Grid Operator 8.4 detects that it's running in a FIPS-enabled environment and it automatically enables FIPS mode for your **Infinispan** CR.

When you upgrade Data Grid Operator from version 8.3.7 to version 8.4 the **-Dcom.redhat.fips=false** field is removed from the **Infinispan** CR and FIPS mode is no longer disabled.



NOTE

When you use Data Grid Server version 8.3.1, you must set the **-Dcom.redhat.fips=false** field in your **Infinispan** CR.

1.6. DATA GRID OPERATOR 8.4.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	Operand versions	Features
8.4.14	8.4.7	8.4.7-1 8.4.6-2 8.4.6-1 8.4.5-2 8.4.5-1 8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes.

Data Grid Operator version	Data Grid version	Operand versions	Features
8.4.13	8.4.6	8.4.6-2 8.4.6-1 8.4.5-2 8.4.5-1 8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes.
8.4.12	8.4.6	8.4.6-1 8.4.5-2 8.4.5-1 8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes.
8.4.11	8.4.5	8.4.5-2 8.4.5-1 8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	<ul style="list-style-type: none"> ● Fixes security vulnerabilities. ● Includes a bug fix. For details see Fixed in Data Grid Operator 8.4.11.
8.4.10	N/A	N/A	The 8.4.9 release included an error in version naming, when the release was mistakenly labeled as 8.4.10. To address this issue and maintain version continuity, the 8.4.10 release is not available, and the subsequent release is 8.4.11.

Data Grid Operator version	Data Grid version	Operand versions	Features
8.4.9	8.4.5	8.4.5-1 8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	<ul style="list-style-type: none"> Fixes security vulnerabilities. Includes a bug fix. For details see Fixed in Data Grid Operator 8.4.9.
8.4.8	8.4.4	8.4.4-1 8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes. For new features, see Data Grid Operator 8.4.8 .
8.4.7	8.4.3	8.4.3-2 8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Fixes security vulnerabilities.
8.4.6	8.4.3	8.4.3-1 8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes and logging enhancements.
8.4.5	8.4.2	8.4.2-1 8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes. For new features, see Data Grid Operator 8.4.5

Data Grid Operator version	Data Grid version	Operand versions	Features
8.4.4	8.4.1	8.4.1-3 8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Fixes security vulnerabilities.
8.4.3	8.4.1	8.4.1-2 8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Fixes security vulnerabilities.
8.4.2	8.4.1	8.4.1-1 8.4.0-2 8.4.0-1 8.3.1-1	Includes several bug fixes. For new features, see Data Grid Operator 8.4.2 .
8.4.1	8.4.0	8.4.0-2 8.4.0-1 8.3.1-1	Includes updates necessary for Open Source license compliance.
8.4.0	8.4.0	8.4.0-1 8.3.1-1	See Data Grid Operator 8.4 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 release does not include any known issues that affect Data Grid clusters that you manage with Data Grid Operator. For complete details about Data Grid, see the [Data Grid 8.4 release notes](#).

2.2. FIXED IN DATA GRID OPERATOR 8.4.14

Data Grid Operator 8.4.14 includes the following notable fixes:

- [JDG-6904](#) Unreadable log timestamps
- [JDG-6451](#) Updates to **ADDITIONAL_VARS** are not applied to existing clusters
- [JDG-6760](#) Backup CR does not work when xsite replication is enabled
- [JDG-6860](#) [Operator] Indexed HotRod Rolling Upgrades fail when cache name is sorted before `__protobuf_metadata` cache

2.3. FIXED IN DATA GRID OPERATOR 8.4.12

Data Grid Operator 8.4.12 includes the following notable fixes:

- [JDG-6544](#) Changing the service type from **Cache** to **DataGrid** does not result in automatic reconciliation
- [JDG-6574](#) Setting up a custom **NodePort** value for cross-site replication has no effect
- [JDG-6573](#) Deadlock occurs when creating caches with a zero-capacity node and a single stateful node
- [JDG-6425](#) Updates to **podTargetLabel** are not reflected in pods

2.4. FIXED IN DATA GRID OPERATOR 8.4.11

Data Grid Operator 8.4.11 includes the following notable fixes:

- [JDG-6584](#) **Infinispan** CR should fail if Encryption keystore/certs missing in Secret

2.5. FIXED IN DATA GRID OPERATOR 8.4.9

Data Grid Operator 8.4.9 includes the following notable fixes:

- [JDG-6549](#) Data Grid Server HEAD requests failing with End Of File (EOF)

2.6. FIXED IN DATA GRID OPERATOR 8.4.8

Data Grid Operator 8.4.8 includes the following notable fixes:

- [JDG-6412](#) Data Grid Operator crashes when **Cache** CR is missing the template definition

- [JDG-6373](#) **FileAlreadyExistsException** during dependency extraction after container restart
- [JDG-6304](#) Data Grid webhook allows incompatible TLS configuration

2.7. FIXED IN DATA GRID OPERATOR 8.4.6

Data Grid Operator 8.4.6 includes the following notable fixes:

- [JDG-6128](#) Data Grid Operator logs error message with stacktrace multiple times after cluster restart
- [JDG-6127](#) Data Grid Operator repeatedly logs error messages indicating missing secrets while waiting for OpenShift to create a secret
- [JDG-6207](#) **spec.Image** field is overwritten by Operand image for CVE releases
- [JDG-6204](#) **status.Operand.Image** field not updated when defining **spec.Image**
- [JDG-6107](#) Creating caches with authorization fails and produces output issues
- [JDG-6055](#) Changing **memory** or **cpu** values for **spec.configListener** has no effect on the ConfigListener deployment
- [JDG-5835](#) Scaling cluster down and up with **purge-on-startup=false** with one or more file-stores might result in stale entries

2.8. FIXED IN DATA GRID OPERATOR 8.4.5

Data Grid Operator 8.4.5 includes the following notable fixes:

- [JDG-6063](#) ConfigListener ignores the Cache CR metadata name when discovering existing Cache CRs
- [JDG-5986](#) DNS discovery fails when pods are not marked as ready
- [JDG-5935](#) Outdated names for metering labels
- [JDG-5936](#) and [JDG-5931](#) Metering labels are not updated after Data Grid Operator upgrade
- [JDG-5939](#) Removal of the GracefulShutdownTask from Data Grid Operator

2.9. FIXED IN DATA GRID OPERATOR 8.4.2

Data Grid Operator 8.4.2 includes the following notable fixes:

- [JDG-5623](#) Gossip router fails to start with TLS configured on FIPS enabled OpenShift
- [JDG-5681](#) Incorrect **Cache CR** status when immutable fields are modified
- [JDG-5577](#) **Cache CR** created by **listener** is not removed upon disabling the listener
- [JDG-5756](#) Users cannot configure LDAP because the authentication mechanisms of the default endpoint cannot be modified

- [JDG-5789](#) **Infinispan.status.podStatus** field shows incorrect pod names for existing deployment topology
- [JDG-5791](#) Data Grid Operator modifies the content of **spec.template** in the **Cache** CR after its creation
- [JDG-5794](#) **ConfigListener** log level cannot be modified
- [JDG-5818](#) OpenShift rolling upgrades with **RollingMigration** strategy result in **Infinispan** CR status as **Pending**
- [JDG-5820](#) Data Grid Operator fails when upgrading Data Grid Server from 8.3.1-1 to 8.4.0-x using the Hot Rod rolling migration strategy
- [JDG-5836](#) **ConfigListener** stale CR check only compares resource names

2.10. FIXED IN DATA GRID OPERATOR 8.4.0

Data Grid Operator 8.4.0 includes the following notable fixes:

- [JDG-5680](#) Default Anti-affinity strategy configuration with the Data Grid Operator is not valid
- [JDG-5650](#) **configListener** breaks with non-yaml **Cache** CRs template and large strings
- [JDG-5461](#) Server image doesn't enable Garbage collection (GC) logging by default
- [JDG-5459](#) Zero controller execute can hang indefinitely if Zero Pod is not immediately ready

CHAPTER 3. DATA GRID ON OPENSHIFT

3.1. DATA GRID 8.4 IMAGES

Data Grid 8.4 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 8.4 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](#)