



Red Hat Data Grid 8.2

Data Grid Operator 8.2 Release Notes

Data Grid Operator 8.2

Red Hat Data Grid 8.2 Data Grid Operator 8.2 Release Notes

Data Grid Operator 8.2

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

Get release information for Data Grid Operator 8.2 and learn how Red Hat supports Data Grid installations on Red Hat OpenShift.

Table of Contents

RED HAT DATA GRID	3
DATA GRID DOCUMENTATION	4
DATA GRID DOWNLOADS	5
MAKING OPEN SOURCE MORE INCLUSIVE	6
CHAPTER 1. UPGRADE TO THE LATEST DATA GRID VERSION	7
CHAPTER 2. DATA GRID OPERATOR 8.2	8
2.1. DATA GRID OPERATOR 8.2 GA	8
Backup CR and Restore CR	8
Batch CR	8
Usability improvements with user authentication	8
Extended configuration options for cross-site replication	8
Ephemeral storage types	9
Grafana dashboards	9
Custom code deployment	9
Cloud events	9
Native CLI	9
2.2. DATA GRID OPERATOR 8.2.X RELEASE INFORMATION	9
2.3. KNOWN ISSUES WITH OPENSIFT DEPLOYMENTS	11
CHAPTER 3. DATA GRID ON OPENSIFT	14
3.1. DATA GRID 8.2 IMAGES	14
Custom Data Grid Deployments	14
3.2. EMBEDDED CACHES ON OPENSIFT	14
CHAPTER 4. TECHNOLOGY PREVIEWS	16
4.1. TECHNOLOGY PREVIEW FEATURES	16

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 8.2 Documentation](#)
- [Data Grid 8.2 Component Details](#)
- [Supported Configurations for Data Grid 8.2](#)
- [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. UPGRADE TO THE LATEST DATA GRID VERSION

Red Hat recommends you upgrade any deployments from 8.2.x to the latest Data Grid 8 version as soon as possible. The Data Grid team regularly patch security vulnerabilities and actively fix issues on the latest version of the software.

You can find the latest Data Grid documentation at [Red Hat Data Grid Product Documentation](#) .

CHAPTER 2. DATA GRID OPERATOR 8.2

Get version details for Data Grid Operator 8.2 as well as information about issues.

2.1. DATA GRID OPERATOR 8.2 GA

Find out what's new with Data Grid Operator for Data Grid 8.2.

Backup CR and Restore CR

Data Grid Operator watches for custom resources (CR) that let you back up and restore Data Grid cluster state for disaster recovery or when migrating between Data Grid versions.

Backup CR

Archives Data Grid cluster content to a persistent volume.

Restore CR

Restores archived content to a Data Grid cluster.

Batch CR

Data Grid Operator provides a **Batch** CR that lets you create Data Grid resources in bulk. The **Batch** CR uses the Data Grid Operator command line interface (CLI) in batch mode to carry out sequences of operations.

Usability improvements with user authentication

Internal endpoint for the operator user

Data Grid Operator now separates the **operator** user that it uses for internal operations from application users and adds a new ***-generated-operator-secret** authentication secret.



NOTE

You no longer need to add credentials to a secret so that Data Grid Operator can access your cluster when creating caches. As of Data Grid 8.2, Data Grid Operator uses the **operator** user and corresponding password to perform cache operations.

Disabling authentication

You can disable authentication for application users to allow unrestricted access to caches.

Extended configuration options for cross-site replication

Cross-site replication with different cluster names

Data Grid Operator no longer requires Data Grid clusters to have the same name and run in matching namespaces for cross-site replication.

If cluster names or namespaces are different, you can specify them with the **spec.sites.locations.clusterName** and **spec.sites.locations.namespace** fields.

Manual cross-site configuration

You can specify static hosts and ports for Data Grid clusters to perform backups to Data Grid clusters running outside OpenShift or where access to the Kubernetes API is not available.

Configuring cross-site replication in the same OpenShift Container Platform cluster

For evaluation and demonstration purposes, you can configure Data Grid to back up between nodes in the same OpenShift cluster.

Ephemeral storage types

Data Grid Operator now lets you set the `spec.service.container.ephemeralStorage` field to define whether storage is ephemeral or permanent. Set the value to `true` to use ephemeral storage, which means all data in storage is deleted when clusters shut down or restart. The default value is `false`, which means storage is permanent.

Grafana dashboards


Data Grid Operator creates Grafana dashboards that let you visualize metrics to more effectively monitor Data Grid services.

Custom code deployment

Data Grid Operator allows you to add custom code, such as scripts and event listeners, to your Data Grid clusters.

Cloud events

Configure Data Grid as a Knative source by sending **CloudEvents** to Apache Kafka topics.



NOTE

Sending cloud events with Red Hat OpenShift Serverless is currently available as a [Technology Preview](#).

Native CLI

Data Grid 8.2 adds a native CLI that you can run on Linux, macOS, or Windows and use as an **oc** client plugin.

1. Download the native CLI from the Red Hat customer portal at [Data Grid Software Downloads](#).
2. Open the **README** included with the distribution for installation instructions and example usage.



NOTE

The native CLI is currently available as a [Technology Preview](#).

2.2. DATA GRID OPERATOR 8.2.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 more frequent.



IMPORTANT

If you upgrade Data Grid clusters manually and have upgraded the channel for your Data Grid Operator subscription from **8.1.x** to **8.2.x** you should apply the upgrade for the latest Data Grid 8.2.x version as soon as possible to avoid potential data loss that can result from an issue in 8.2.0.

Data Grid Operator version	Data Grid version	Features
8.2.8	8.2.3	<p>Fixes the following security vulnerabilities CVE-2021-44832, CVE-2021-45046, and CVE-2021-45105 which affect the Apache Log4j logging library.</p> <p>Red Hat recommends you upgrade your deployment to this version as soon as possible. If you cannot upgrade, Red Hat recommends that you follow the mitigation steps that are included in the security advisory page for each of the Log4j vulnerabilities mentioned above.</p>
8.2.7	8.2.2	<p>Fixes security vulnerabilities, including CVE-2021-44228, which affects the Apache Log4j logging library.</p> <p>Red Hat recommends you upgrade your deployment to this version as soon as possible. If you cannot upgrade, Red Hat recommends that you follow the steps to mitigate this vulnerability in the RHSB-2021-009 Log4Shell - Remote Code Execution security bulletin.</p>
8.2.6	8.2.1	Fixes security vulnerabilities.
8.2.5	8.2.1	Fixes security vulnerabilities.
8.2.4	8.2.1	<ul style="list-style-type: none"> * Upgrades Data Grid Operator to Level 4 - Deep Insights capabilities. * Improves event logging to enhance deployment monitoring. * Updates the API version of the Data Grid CRDs. This change removes usage of deprecated OpenShift APIs that will no longer be available as of OpenShift 4.9. * Adds support for loading external dependencies via HTTP or FTP. * Fixes security vulnerabilities.
8.2.3	8.2.1	Fixes security vulnerabilities.
8.2.2	8.2.1	<ul style="list-style-type: none"> * Add configurable ports for Load Balancer services with the spec.expose.port and spec.service.sites.local.expose.port fields. * Fixes bugs and security vulnerabilities.

Data Grid Operator version	Data Grid version	Features
8.2.1	8.2.0	<ul style="list-style-type: none"> * Adds support for security authorization (Role Based Access Control or RBAC). * Adds support for client certificate authentication. * Fixes bugs and security vulnerabilities.
8.2.0	8.2.0	See Data Grid Operator 8.2 GA

2.3. KNOWN ISSUES WITH OPENSIFT DEPLOYMENTS

This section describes issues that affect Data Grid clusters running on Red Hat OpenShift. For complete details about Data Grid, you should refer to the [Data Grid 8.2 release notes](#).

Data Grid pods crash after upgrade or restart on Red Hat OpenShift Container Platform 4.9

Issue: [JDG-5026](#)

Description: After upgrading Data Grid, or when restarting a cluster after a graceful shutdowns, some pods do not start running. As a result the Data Grid cluster cannot successfully restart and restore to the correct state.

Data Grid Server nodes running in pods do not log any messages when this occurs. The affected nodes encounter a fatal error and silently crash.



NOTE

This issue affects installations on Red Hat OpenShift version 4.9 only and is related to changes in how Red Hat OpenShift handles **StatefulSet** objects.

Workaround: If you install Data Grid Operator on Red Hat OpenShift 4.9, you should use the **Backup** and **Restore** CRs to recreate Data Grid clusters. Before you upgrade the Data Grid version or perform a graceful shutdown create a backup. You can then create a new cluster and restore its state.

Alternatively you can scale your cluster to one node **replicas=1** before you upgrade or restart. However that node must have capacity for the entire data set.

Clients cannot connect to remote caches that use TLS/SSL encryption

Issue: [JDG-4763](#)

Description: Clients cannot connect to remote caches and Data Grid logs print a **WARN** log message related to SSL.

See the following Red Hat knowledge base article for full details about log messages: [Clients are not able to connect a server after update to RHDG 8.2.1](#)

Workaround: Modify your **Infinispan** CR to use Java TLS/SSL libraries instead of OpenSSL as follows:

```
spec:  
  container:  
    extraJvmOpts: "-Dorg.infinispan.openssl=false"
```

Data Grid Operator upgrades Data Grid clusters only if all pods are in the `Ready` state

Issue: [JDG-4724](#)

Description: Data Grid cluster upgrades fail if Data Grid Operator detects any pods are not fully running and in the `Ready` state.

Workaround: There is no workaround for this issue. When upgrades fail, Data Grid clusters do not roll back to the previous version. In this case you should delete the cluster and then create a new one with the desired version.

Unexpected pod restarts during upgrade can lead to deployment failure

Issue: [JDG-4737](#)

Description: If Data Grid pods restart during OLM upgrade from an earlier version, and the restart does not originate from Data Grid Operator, then Data Grid cluster can fail to deploy successfully.

Workaround: There is no workaround for this issue.

Data Grid Server does not configure a property realm for authentication with the `Validate` strategy

Issue: [JDG-4722](#)

Description: If you configure Data Grid to validate client certificates, **spec.security.endpointEncryption.clientCert: `Validate`**, then Data Grid Server disables credentials authentication.

Workaround: There is no workaround for this issue.

Hot Rod clients cannot connect to Data Grid clusters that validate client certificates

Issue: [JDG-4688](#)

Description: If you configure Data Grid to validate client certificates, **spec.security.endpointEncryption.clientCert: `Validate`**, Hot Rod clients cannot connect to Data Grid clusters using the `EXTERNAL` authentication mechanism and the following message is written to logs:

```
Caused by: java.lang.SecurityException: ISPN004031: The selected authentication mechanism  
'EXTERNAL' is not among the supported server mechanisms:
```

Workaround: There is no workaround for this issue. If you require client certificate authentication and use Hot Rod clients, you should configure Data Grid to authenticate client certificates, **spec.security.endpointEncryption.clientCert: `Authenticate`**.

Hot Rod clients cannot connect to Data Grid clusters through OpenShift Routes when using client certificate authentication

Issue: [JDG-4689](#)

Description: If you expose Data Grid to clients through an OpenShift Route and enable client certificate authentication, Hot Rod clients cannot successfully connect and **org.infinispan.client.hotrod.exceptions.TransportException** errors are written to logs.

Workaround: There is no workaround for this issue.

Data Grid on OpenShift continually restarts after OOM exceptions

Issue: [JDG-3991](#)

Description: If out of memory exceptions cause Data Grid Server to terminate on OpenShift, the nodes cannot restart. 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.

Native CLI running as an OpenShift client plugin cannot use encrypted connections

Issue: [JDG-4566](#)

Description: Running the native executable of the Data Grid command line interface (CLI) as an **oc** client plugin, which is currently a technology preview feature, is not possible to use the **--trustall** argument when connecting to Data Grid clusters that use endpoint encryption.

Workaround: There is no workaround for this issue.

CHAPTER 3. DATA GRID ON OPENSIFT

3.1. DATA GRID 8.2 IMAGES

Data Grid 8.2 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.



NOTE

Red Hat supports Data Grid 8.2 on OpenShift only through Data Grid Operator subscriptions.

Custom Data Grid Deployments

Red Hat does not support customization of any 8.2 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
- Encryption mechanisms (SYM_ENCRYPT or ASYM_ENCRYPT)
- Persistent datasources

Additional resources

- [Data Grid Container Images](#)

3.2. EMBEDDED CACHES ON OPENSIFT

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.

Reference

- [DNS_PING](#)

CHAPTER 4. TECHNOLOGY PREVIEWS

Data Grid releases offer technology preview features. Find out more about Red Hat support for these capabilities.

4.1. TECHNOLOGY PREVIEW FEATURES

Technology preview features or capabilities are not supported with Red Hat production service-level agreements (SLAs) and might not be functionally complete.

Red Hat does not recommend using technology preview features or capabilities for production. These features provide early access to upcoming product features, which enables you to test functionality and provide feedback during the development process.

For more information, see [Red Hat Technology Preview Features Support Scope](#) .