Red Hat Data Grid 8.1

Data Grid Operator 8.1 Release Notes

Data Grid Operator 8.1

Last Updated: 2023-11-23

Red Hat Data Grid 8.1 Data Grid Operator 8.1 Release Notes

Data Grid Operator 8.1

Legal Notice

Abstract

Get release information for Data Grid Operator 8.1 and learn how Red Hat supports Data Grid installations on Red Hat OpenShift.

Table of Contents

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 8.1 Documentation
- Data Grid 8.1 Component Details
- Supported Configurations for Data Grid 8.1
- 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 8.1.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.

1.1. APACHE LOG4J SECURITY VULNERABILITIES

Data Grid includes Apache Log4j components in:

- Data Grid Server distribution
- Data Grid Server container image for Red Hat OpenShift

Data Grid 8.2.2 fixes CVE-2021-44228 which is a critical security vulnerability. Data Grid 8.2.3 fixes several Log4j CVEs of moderate severity. For more information, see the Data Grid 8.2 release notes or the Data Grid Operator 8.2 release notes. You can find the latest Data Grid documentation at Red Hat Data Grid Product Documentation.

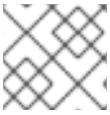

NOTE

Data Grid Operator version 8.2.7 corresponds to Data Grid 8.2.2.

If you cannot upgrade your Data Grid deployment, you must follow the steps to mitigate the critical vulnerability as outlined in the RHSB-2021-009 Log4Shell - Remote Code Execution security bulletin.

CHAPTER 2. DATA GRID OPERATOR 8.1

Learn about new features and get version details for Data Grid Operator 8.1.

2.1. VERSION DETAILS

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

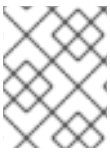

NOTE

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

Data Grid Operator version	Data Grid version	Features
8.1.6	8.1.1	* Bug fixes. * Configurable **StorageClass** objects for persistent volume claims. * Documentation improvements for creating routes, Hot Rod client configuration, and cache creation. * Fixes security vulnerabilities.
8.1.5	8.1.1	* Support for custom labels. * Metering labels for Red Hat subscriptions. * Fixes security vulnerabilities.
8.1.4	8.1.1	* Anti-affinity settings. * Ability to disable encryption. * Multiple namespace installation.
8.1.3	8.1.0	Fixes security vulnerabilities.
8.1.2	8.1.0	Fixes security vulnerabilities.
8.1.1	8.1.0	Fixes security vulnerabilities.

Data Grid Operator version	Data Grid version	Features
8.1.0	8.1.0	* Cross-site replication. * Automatic scaling. * Ability to expose services via routes. * Automatic encryption with the OpenShift service CA. * **Cache** CR. * Configurable number of owners for Cache service nodes. * Data Grid 8.1.0

2.2. DATA GRID OPERATOR 8.1 GA

Find out what's new with Data Grid Operator for Data Grid 8.1.

2.2.1. Cross-Site Replication

Data Grid Operator improves setup and management of cross-site replication capabilities for Data Grid clusters running on OpenShift.

See Cross-Site Replication with Data Grid Operator .

2.2.2. Expose Services via Routes

Data Grid Operator updates the **spec.expose** resource so you can create OpenShift Routes with passthrough encryption to make Data Grid clusters available on the network.

See Exposing Data Grid Through Routes .

2.2.3. Automatic Scaling

Data Grid Operator can automatically scale the default cache on Cache service nodes up or down based on memory usage.

See Configuring Automatic Scaling.

2.2.4. Automatic Encryption with the OpenShift Service CA

By default, if the Red Hat OpenShift service CA is available, Data Grid Operator generates TLS certificates, signed by the Red Hat OpenShift service CA, to encrypt client connections.

NOTE

You must use encryption if the OpenShift service CA is present.

If you are upgrading from 8.0, you should retrieve the generated **tls.crt** certificate and add it to a client trust store.

See Securing Data Grid Connections.

2.2.5. Cache Custom Resource

You can now create caches with Data Grid service nodes with the **Cache** CR.

IMPORTANT

Creating caches with Data Grid Operator is available as a technology preview.

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 them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

See Creating Data Grid Caches.

2.2.6. Number of Owners

Cache service nodes now use a default value of **2** for the number of owners, which replicates each entry across the cluster. You can modify the number of owners with the **spec.service.replicationFactor** resource in your **Infinispan** CR.

See Configuring the Number of Owners.

CHAPTER 3. DATA GRID ON OPENSHIFT

3.1. DATA GRID 8.1 IMAGES

Data Grid 8.1 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 8.1 on OpenShift only through Data Grid Operator subscriptions.

Custom Data Grid Deployments

Red Hat does not support customization of any 8.1 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. DATA GRID LIBRARY MODE ON OPENSHIFT

Embedding Data Grid in custom applications, also referred to as Library Mode, is intended for specific uses only when running on OpenShift:

- 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