Red Hat Data Grid 8.5

Data Grid Operator 8.5 Release Notes

Get release information for Data Grid Operator 8.5

Last Updated: 2024-09-18

Red Hat Data Grid 8.5 Data Grid Operator 8.5 Release Notes

Get release information for Data Grid Operator 8.5

Legal Notice

Abstract

Find out about features and enhancements in Data Grid Operator 8.5 as well as known and resolved issues.

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 8.5 Documentation
- Data Grid 8.5 Component Details
- Supported Configurations for Data Grid 8.5
- 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 8.5

Get version details for Data Grid Operator 8.5 and information about issues.

1.1. DATA GRID OPERATOR 8.5.3

What's new in 8.5.3.

Automatic reloading of SSL/TLS certificates

Starting with Data Grid 8.5.1, Data Grid monitors keystore files for changes and automatically reloads them, without requiring a server or client restart, when certificates are renewed.

Therefore, with Data Grid Operator 8.5.3, **StatefulSet** rolling update is not triggered on key or truststore update in a server when managing Data Grid 8.5.1 Operands because it is not required.

1.2. DATA GRID OPERATOR 8.5.0

What's new in 8.5.0.

Ability to configure InitContainer resource

You can now configure the **InitContainer** resource. Previously, if a **LimitRange** was in effect for the deployment namespace, then the **InitContainer** would be restricted to these resource values causing issues such as OutOfMemoryError. You can configure InitContainer resource configuration in the Data Grid CR as follows:

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

Ability to define Batch resource CPU and memory request/limits

You can now define CPU and memory request/limits for Batch Job created by the Operator. You can define the resource request/limits in the Batch CR as follows:

```
apiVersion: infinispan.org/v2alpha1
kind: Batch
metadata:
  name: mybatch
spec:
  cluster: infinispan
  configMap: mybatch-config-map
  container:
    cpu: "2000m:1000m"
    memory: "2Gi:1Gi"
```

TLSv1.3 encryption for cross-site encryption

The default encryption protocol for cross-site is now TLSv1.3 instead of TLSv1.2.

Ability to define TopologyPodConstraints and Tolerations in StatefulSet

You can now configure more advanced high availability configurations by defining **TopologyPodConstraints** and **Tolerations** in **spec.statefulSet**.

Example

```
kind: Infinispan
...
spec:
  scheduling:
    affinity:
      ...
    tolerations:
      ...
    topologySpreadConstraints:
      ...
```

Cache service type removed

RHDG 8.5 removes the Cache service type cache. Instead, use the **DataGrid** service type to automate complex operations such as cluster upgrades and data migration.

Cloud events removed

RHDG 8.5 removes cloud events integration.

1.3. DATA GRID OPERATOR 8.5.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
8.5.3	8.5.1	8.5.1-1 8.5.0-3 8.5.0-2 8.5.0-1 8.4.8-1 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	Includes several bug fixes.

8.5.2	8.5.0	8.5.0-3 8.5.0-2 8.5.0-1 8.4.8-1 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	Includes several bug fixes.
8.5.1	8.5.0	8.5.0-2 8.5.0-1 8.4.8-1 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	Includes several bug fixes.
8.5.0	8.5.0	8.5.0-1 8.4.8-1 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	Includes several bug fixes.

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 8.5 release notes.

2.2. FIXED IN DATA GRID OPERATOR 8.5.3

Data Grid Operator 8.5.3 includes the following notable fixes:

- JDG-6764 **Updates to spec.image on existing Infinispan CR have no effect**

2.3. FIXED IN DATA GRID OPERATOR 8.5.0

Data Grid Operator 8.5.0 includes the following notable fixes:

- JDG-5000 **Gossip router pod generates lot of SSLHandshake warn messages**
- JDG-7063 **Nil pointer error on upgrade from dropped Operand version**
- JDG-7032 **Operator generates truststore from certificates using outdated algorithms**
- JDG-7093 **Operator may not reconcile Data Grid cluster properly after upgrade**
- JDG-5989 **Operator Configuration spec.autoscale should not be possible with a Data Grid service**

CHAPTER 3. DATA GRID ON OPENSHIFT

3.1. DATA GRID 8.5 IMAGES

Data Grid 8.5 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 8.5 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