



Red Hat OpenShift Data Foundation 4.9

4.9 Release notes

Release notes for feature and enhancements, known issues, and other important release information.

Red Hat OpenShift Data Foundation 4.9 4.9 Release notes

Release notes for feature and enhancements, known issues, and other important release information.

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

The release notes for Red Hat OpenShift Data Foundation 4.9 summarizes all new features and enhancements, notable technical changes, and any known bugs upon general availability.

Table of Contents

CHAPTER 1. OVERVIEW	3
1.1. ABOUT THIS RELEASE	3
CHAPTER 2. NEW FEATURES	4
CHAPTER 3. ENHANCEMENTS	6
CHAPTER 4. TECHNOLOGY PREVIEWS	7
CHAPTER 5. DEVELOPER PREVIEWS	8
CHAPTER 6. BUG FIXES	9
CHAPTER 7. KNOWN ISSUES	13
CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES	16
8.1. RHBA-2022:0865 OPENSIFT DATA FOUNDATION 4.9.4 BUG FIXES AND SECURITY UPDATES	16
8.2. RHBA-2022:0684 OPENSIFT DATA FOUNDATION 4.9.3 BUG FIXES AND SECURITY UPDATES	16
8.3. RHBA-2022:0346 OPENSIFT DATA FOUNDATION 4.9.2 BUG FIXES AND SECURITY UPDATES	16
8.4. RHSA-2022:0032 OPENSIFT DATA FOUNDATION 4.9.1 BUG FIXES AND SECURITY UPDATES	16

CHAPTER 1. OVERVIEW

Red Hat OpenShift Data Foundation is software-defined storage that is optimized for container environments. It runs as an operator on OpenShift Container Platform to provide highly integrated and simplified persistent storage management for containers.

Red Hat OpenShift Data Foundation is integrated into the latest Red Hat OpenShift Container Platform to address platform services, application portability, and persistence challenges. It provides a highly scalable backend for the next generation of cloud-native applications, built on a new technology stack that includes Red Hat Ceph Storage, the Rook.io Operator, and NooBaa's Multicloud Object Gateway technology.

Red Hat OpenShift Data Foundation provides a trusted, enterprise-grade application development environment that simplifies and enhances the user experience across the application lifecycle in a number of ways:

- Provides block storage for databases.
- Shared file storage for continuous integration, messaging, and data aggregation.
- Object storage for cloud-first development, archival, backup, and media storage.
- Scale applications and data exponentially.
- Attach and detach persistent data volumes at an accelerated rate.
- Stretch clusters across multiple data-centers or availability zones.
- Establish a comprehensive application container registry.
- Support the next generation of OpenShift workloads such as Data Analytics, Artificial Intelligence, Machine Learning, Deep Learning, and Internet of Things (IoT).
- Dynamically provision not only application containers, but data service volumes and containers, as well as additional OpenShift Container Platform nodes, Elastic Block Store (EBS) volumes and other infrastructure services.

1.1. ABOUT THIS RELEASE

Red Hat OpenShift Data Foundation 4.9 ([RHSA-2021:5085](#) and [RHSA-2021:5086](#)) is now available. New enhancements, features, and known issues that pertain to OpenShift Data Foundation 4.9 are included in this topic.

Red Hat OpenShift Data Foundation 4.9 is supported on the Red Hat OpenShift Container Platform versions 4.9. For more information, see [Red Hat OpenShift Data Foundation Supportability and Interoperability Guide](#).

With the release of OpenShift Data Foundation 4.9, version 4.5 is now end of life. For more information, see [Red Hat OpenShift Container Platform Life Cycle Policy](#).

CHAPTER 2. NEW FEATURES

This section describes new features introduced in Red Hat OpenShift Data Foundation 4.9.

User interface, product component, and documentation rebranding

OpenShift Container Storage, based on the open source Ceph technology, has expanded its scope and foundational role in a containerized, hybrid cloud environment since its introduction. To better reflect these foundational and infrastructure distinctives, OpenShift Container Storage is now *OpenShift Data Foundation*.

OpenShift Data Foundation 4.9 now includes:

- Improved dashboards for viewing all storage system status and metrics
- Easy to use wizard for creating storage system
- Rebranded user interface and documentation

To view documentation for OpenShift Container Storage version 4.8 and earlier, see [Product Documentation for Red Hat OpenShift Container Storage](#).

To update from OpenShift Container Storage 4.8 to OpenShift Data Foundation 4.9, you must freshly install the OpenShift Data Foundation operator from the OpenShift Container Platform Operator Hub. This fresh operator installation upgrades OpenShift Container Storage version 4.8 and all its components to OpenShift Data Foundation version 4.9.

For more information, see [Upgrading to OpenShift Data Foundation](#).

Multicloud Object Gateway bucket replication

Data replication from one Multicloud Object Gateway (MCG) bucket to another MCG bucket provides higher resiliency and better collaboration options. These buckets can be either data buckets or namespace buckets backed by any supported storage solutions.

For more information, see [Multicloud Object Gateway bucket replication](#).

Ability to view pool compression metrics

In this release, you can view the pool compression metrics, which provide information about the amount of storage space saved, the effectiveness of pool compression when it is enabled, and its impact on capacity consumption. The per-pool metrics available with this release provide information that enables you to reduce cost and consume data more efficiently. Also, you can disable compression if it is not effective.

For more information, see [Pool metrics](#).

Automated scaling of Multicloud Object Gateway endpoint pods

You can use the automated scaling of Multicloud Object Gateway endpoint pods feature to automate the resource adjustments based on increases or decreases to the load. This provides better performance and serviceability to manage your production resources for your S3 load.

For more information, see [Automatic scaling of Multicloud Object Gateway endpoints](#).

Deployment and monitoring layer for pluggable external storage (IBM FlashSystem®)

In this release, you can connect to and monitor IBM FlashSystem® storage using OpenShift Data Foundation. OpenShift Data Foundation extends IBM FlashSystem to file and object storage while providing a single view for both the underlying storage and OpenShift Data Foundation data layer.

For more information, see [Deploy OpenShift Data Foundation using IBM FlashSystem](#) .

CHAPTER 3. ENHANCEMENTS

This section describes the major enhancements introduced in Red Hat OpenShift Data foundation 4.9.

Deletion of data is allowed when the storage cluster is full

Previously, when the storage cluster was full, the Ceph Manager hung on checking pool permissions while reading the configuration file. The Ceph Metadata Server (MDS) did not allow write operations to occur when the Ceph OSD was full, resulting in an **ENOSPACE** error. When the storage cluster hit full ratio, users could not delete data to free space using the Ceph Manager volume plugin.

With this release, the new FULL capability is introduced. With the FULL capability, the Ceph Manager bypasses the Ceph OSD full check. The **client_check_pool_permission** option is disabled by default whereas, in previous releases, it was enabled. With the Ceph Manager having FULL capabilities, the MDS no longer blocks Ceph Manager calls. This results in allowing the Ceph Manager to free up space by deleting subvolumes and snapshots when a storage cluster is full.

Standalone Multicloud Object Gateway component deployment

With this release, you can deploy OpenShift Data Foundation with only the Multicloud Object Gateway component in a standalone mode. In this mode, there is no CephCluster accompanying the StorageCluster, and hence Multicloud Object Gateway is not using a Ceph-based storage volume.

Movement of Core and DB pods is enabled when a node fails

OpenShift Container Platform does not mark the node as disconnected unless it is deleted. As a result, Core and DB pods, which are the statefulsets are not automatically evicted on such failed nodes. With this update, when a node fails, the DB and Core pods are evicted and moved to a new node.

Volume snapshot restore to a different pool

With this update, you can restore a volume snapshot of persistent volume claim (PVC) into a different pool than the parent volume. Previously, a volume snapshot could only be restored into the same pool.

Multiple file systems are not created with existing pools

With this update, after you create the **filesystem.yaml**, multiple file systems with the existing pool are not created even if you delete or recreate the **filesystem.yaml**. This avoids data loss.

Auto-detection of Vault's Secret Key/Value store version

With this enhancement, Vault's Secret Key/Value store version is auto-detected.

Configuring VAULT_BACKEND parameter for HashiCorp Vault is now allowed


With this update, you can configure the **VAULT_BACKEND** parameter for selecting the type of backend used by HashiCorp Vault. The autodetection of the backend used by HashiCorp Vault does not always work correctly. In case of a non-common configuration, the automatically detected configuration parameter might be set incorrectly. By allowing you to configure the **VAULT_BACKEND** parameter, non-common configurations can be forced to use a particular type of backend.

Human-readable format for output of time in the Multicloud Object Gateway CLI

With this release, the output of time in the Multicloud Object Gateway (MCG) CLI shows human-readable format (days-hours-minutes-seconds) instead of minutes and seconds.

CHAPTER 4. TECHNOLOGY PREVIEWS

This section describes technology preview features introduced in Red Hat OpenShift Data Foundation 4.9 under Technology Preview support limitations.



IMPORTANT

Technology Preview features are provided with a limited support scope, as detailed on the Customer Portal: [Technology Preview Features Support Scope](#).

PV encryption - service account per namespace

As of OpenShift Data Foundation 4.9, you can use service accounts to authenticate a tenant with Vault as a technology preview. For more information, see [Persistent volume encryption](#).

Alerts to control overprovision

With this release, you can get alerts for the overprovision. This enables you to define a quota on the amount of persistent volume claims (PVCs) consumed from a storage cluster based on the specific application namespace. For more information, see [Overprovision level policy control](#).

CHAPTER 5. DEVELOPER PREVIEWS

This section describes developer preview features introduced in Red Hat OpenShift Data Foundation 4.9.



IMPORTANT

Developer preview feature is subject to Developer preview support limitations. Developer preview releases are not intended to be run in production environments. The clusters deployed with the developer preview features are considered to be development clusters and are not supported through the Red Hat Customer Portal case management system. If you need assistance with developer preview features, reach out to the ocs-devpreview@redhat.com mailing list and a member of the Red Hat Development Team will assist you as quickly as possible based on availability and work schedules.

Regional-DR with Advanced Cluster Management

Regional-DR solution provides an automated "one-click" recovery in the event of a regional disaster. The protected applications are automatically redeployed to a designated OpenShift Container Platform with OpenShift Data Foundation cluster that is available in another region.

For more information, see [Configuring Regional-DR with Advanced Cluster Management](#).

Quota support for object data

You can now set quota options for object bucket claims (OBC) to avoid resource starvation and increase the usage of the product. You set the quota during the OBC creation using the options **maxObjects** and **maxSize** in the custom resource definitions (CRD). You can also update these options after the OBC creation.

For more information, see <https://access.redhat.com/articles/6541861>.

IPv6 support

With this release, IPv6 single-stack and dual-stack can be used with OpenShift Data Foundation.

CHAPTER 6. BUG FIXES

This section describes notable bug fixes introduced in Red Hat OpenShift Data Foundation 4.9.

Multicloud Object Gateway storage class deleted during uninstall

Previously, Multicloud Object Gateway (MCG) storage class which was deployed as a part of OpenShift Data Foundation deployment was not deleted during uninstall.

With this update, the Multicloud Object Gateway (MCG) storage class gets removed while uninstalling OpenShift Data Foundation.

([BZ#1892709](#))

OpenShift Container Platform alert when OpenShift Container Storage quorum lost

Previously, **CephMonQuorumAtRisk** alert was fired when **mon** quorum was about to be lost, but there was no alert triggered after losing the quorum. This resulted in no notification being sent when the **mon** quorum was completely lost.

With this release, a new alert, **CephMonQuorumLost** is introduced. This alert is triggered when you have only one node left and a single **mon** is running on it. However, at this point the cluster will be in unrecoverable state and the alert serves as a notification of the issue.

([BZ#1944513](#))

Reduce the **mon_data_avail_warn** from 30 % to 15%

Previously, the **mon_data_avail_warn** alert was triggered when the **mon** store was less than 30% and it did not match the threshold value of OpenShift Container Platform's garbage collector for images, which is 15%. With this release, you will see the alert when the available storage at the **mon** store location is less than 15% and not less than 30%.

([BZ#1964055](#))

OSD pods do not log anything if the initial deployment is OpenShift Container Storage 4.4

Previously, object storage daemon (OSD) logs were not generated when OpenShift Container Storage 4.4 was deployed. With this update, the OSD logs are generated correctly.

([BZ#1974343](#))

Multicloud Object Gateway was not able to initialize in a fresh deployment

Previously, after the internal database change from MongoDB to PostgreSQL, duplicate entities that should be unique could be added to the database (MongoDB prevented duplicate entities earlier) due to which Multicloud Object Gateway (MCG) was not working. With this release, duplicate entities are prevented.

([BZ#1975645](#))

PVC is restored when using two different backend paths for the encrypted parent

Previously, when restoring a persistent volume claim (PVC) from a volume snapshot into a different storage class with a different encryption **KMSID**, the restored PVC went into the **Bound** state and the restored PVC failed to get attached to a Pod. This was because the encryption passphrase was being copied with the parent PVC's storage class encryption **KMSID config**. With this release, the restored

PVC's encryption passphrase is copied with the correct encryption **KMSID config** from the destination storage class. Hence, the PVC is successfully restored into a storage class with a different encryption **KMSID** than its parent PVC.

(BZ#1975730)

Deletion of data is allowed when the storage cluster is full

Previously, when the storage cluster was full, the Ceph Manager hung on checking pool permissions while reading the configuration file. The Ceph Metadata Server (MDS) did not allow write operations to occur when the Ceph OSD was full, resulting in an **ENOSPACE** error. When the storage cluster hit full ratio, users could not delete data to free space using the Ceph Manager and **ceph-volume** plugin.

With this release, the new FULL feature is introduced. This feature gives the Ceph Manager FULL capability, and bypasses the Ceph OSD full check. Additionally, the **client_check_pool_permission** option can be disabled. With the Ceph Manager having FULL capabilities, the MDS no longer blocks Ceph Manager calls. This allows the Ceph Manager to free up space by deleting subvolumes and snapshots when a storage cluster is full.

(BZ#1978769)

Keys are completely destroyed in Vault after deleting encrypted persistent volume claims (PVCs) while using the kv-v2 secret engine

HashiCorp Vault added a feature for the key-value store v2 where deletion of the stored keys makes it possible to recover the contents in case the metadata of the deleted key is not removed in a separate step. When using key-value v2 storage for secrets in HashiCorp Vault, deletion of volumes did not remove the metadata of the encryption passphrase from the KMS.

With this update, the keys in HashiCorp Vault is completely destroyed by default when a PVC is deleted. You can set the new configuration option **VAULT_DESTROY_KEYS** to **false** to enable the previous behavior. In that case, the metadata of the keys will be kept in HashiCorp Vault so that recovery of the encryption passphrase of the removed PVC is possible.

(BZ#1979244)

Multicloud Object Gateway object bucket creation is going to Pending Phase

Previously, after the internal database change from MongoDB to PostgreSQL, duplicate entries that should be unique could be added to the database (MongoDB prevented duplicate entries earlier). As a result, creation of new resources such as buckets, backing stores, and so on failed. With this release, duplicate entries are prevented.

(BZ#1980299)

Deletion of CephBlockPool gets stuck and blocks the creation of new pools

Previously, in a Multus enabled cluster, the Rook Operator did not have access to the object storage daemon (OSD) network as it did not have the network annotations. As a result, the **rbd** type commands during a pool cleanup would hang because the OSDs could not be contacted.

With this release, the operator proxies the **rbd** command through a sidecar container in the **mgr** pod and runs successfully during the pool cleanup.

(BZ#1983756)

Standalone Multicloud Object Gateway failing to connect

Previously, the Multicloud Object Gateway (MCG) CR was not updated properly because of the change in the internal DB from MongoDB to PostgreSQL. This caused issues in certain flows. As a result, MCG components were not able to communicate with one another and MCG failures occurred on upgrade.

With this release, MCG CR issue is fixed.

([BZ#1984284](#))

Monitoring spec is getting reset in CephCluster resource in external mode

Previously, when OpenShift Container Storage was upgraded, the monitoring endpoints would get reset in external CephCluster's monitoring spec. This was not an expected behavior and was due to the way monitoring endpoints were passed to the CephCluster. With this update, the way endpoints are passed is changed. Before the CephCluster is created, the endpoints are accessed directly from the JSON secret, **rook-ceph-external-cluster-details** and the CephCluster spec is updated. As a result, the monitoring endpoint specs in the CephCluster is updated properly with appropriate values even after the OpenShift Container Storage upgrade.

([BZ#1984735](#))

CrashLoopBackOff state of noobaa-db-pg-0 pod when enabling hugepages

Previously, enabling **hugepages** on OpenShift Container Platform cluster caused the Multicloud Object Gateway (MCG) database pod to go into a **CrashLoopBackOff** state. This was due to wrong initialization of PostgreSQL. With this release, MCG database pod's initialization of PostgreSQL is fixed.

([BZ#1995271](#))

Multicloud Object Gateway unable to create new object bucket claims

Previously, performance degradation when working against the Multicloud Object Gateway (MCG) DB caused back pressure on all the MCG components which resulted in failure to execute flows within the system such as configuration flows and I/O flows.

With this update, the most time consuming queries are fixed, the DB is cleared quickly, and no back pressure is created.

([BZ#1998680](#))

Buckets fail during creation because of an issue with checking attached resources

Previously, because of a problem in checking resources attached to a bucket during its creation, the bucket would fail to be created. The conditions in the resource validation during bucket creation have been fixed, and the buckets are created as expected.

([BZ#2000588](#))

NooBaa Operator still checks for noobaa-db service after upgrading

Previously, when OpenShift Container Storage was updated from version 4.6, there was a need to retain the old and the new **noobaa-db** StatefulSets for migration purposes. The code still supports both the names of sets. A failure message was generated on the old **noobaa-db** StatefulSet due to a small issue in the code which caused the operator to check the status of the old **noobaa-db** StatefulSet even though it was no longer relevant.

With this update, the operator stops checking the status of the old **noobaa-db** StatefulSet.

([BZ#2008821](#))

Changes to the config maps of the Multicloud Object Gateway (MCG) DB pod does not get reconciled after upgrade

Previously, changes to the config maps of the MCG DB pod did not apply after an upgrade. The flow has been fixed to properly take the variables from the config maps for the DB pod.

(BZ#2012930)

CHAPTER 7. KNOWN ISSUES

This section describes known issues in Red Hat OpenShift Data Foundation 4.9.

odf-operator is missing when OpenShift Container Storage is upgraded from version 4.8 to 4.9

Currently, while upgrading the **ocs-operator**, if you change the channel in the OpenShift Container Storage subscription without installing the **odf-operator**, the cluster will only have the OpenShift Data Foundation and Multicloud Object Gateway (MCG) installed, and the 'odf-operator' will be missing from the cluster.

Workaround: Install the **odf-operator** from the graphical user interface (GUI) or backend. Ensure that the subscription name is **odf-operator** if you create it via the backend.

([BZ#2050251](#))

Multicloud Object Gateway insecure storage account does not support TLS 1.2

Multicloud Object Gateway (MCG) does not support Microsoft Azure storage account configured with Transport Layer Security (TLS) 1.2. As a result, you cannot create the default backing store or any new backing store on a storage account which is with 1.2 only policy.

([BZ#1970123](#))

Critical alert notification is sent after installation of arbiter storage cluster, when Ceph object user for cephobjectstore fails to be created during storage cluster reinstallation

In a storage cluster containing a CephCluster and one or more **CephObjectStores**, if the **CephCluster** resource is deleted before all of the **CephObjectStore** resources are fully deleted, the Rook Operator can still keep connection details about the **CephObjectStores** in memory. If the same **CephCluster** and **CephObjectStores** are re-created, the **CephObjectStores** might enter **Failed** state.

To avoid this issue, you can delete the **CephObjectStores** completely before removing the CephCluster. If you do not want to wait for the CephObjectStores to be deleted, restart the Rook Operator (by deleting the Operator Pod) to avoid the issue if done after uninstall. If you are actively experiencing this issue, restart the Rook Operator to resolve it by clearing the Operator's memory of old CephObjectStore connection details.

([BZ#1974344](#))

Poor performance of stretch clusters on CephFS

Workloads with many small metadata operations might exhibit poor performance because of the arbitrary placement of metadata server (MDS) on multi-site OpenShift Data Foundation clusters.

([BZ#1982116](#))

rook-ceph-operator-config ConfigMap is not updated when OpenShift Container Storage is upgraded from version 4.5 to other version

ocs-operator uses the **rook-ceph-operator-config ConfigMap** to configure **rook-ceph-operator** behaviors, however it only creates it once and then does not reconcile it. This raises the problem that it will not update the default values for the product as they evolve.

Workaround: Administrators can manually change the **rook-ceph-operator-config** values.

([BZ#1986016](#))

Automate the creation of `cephobjectstoreuser` for object bucket claim metrics collector

Currently, the object bucket claim (OBC) metrics collection fails because the `ocs-metrics-exporter` expects the Ceph object store user named `prometheus-user`.

Workaround: Manually, create `prometheus-user` and provide appropriate permissions after the storage cluster creation. Refer to the Prerequisites section of the Knowledge Base article <https://access.redhat.com/articles/6541861> for more information.

([BZ#1999952](#))

`StorageCluster` and `StorageSystem` `ocs-storagecluster` are in error state for a few minutes when installing `StorageSystem`

During `StorageCluster` creation, there is a small window of time where it will appear in an error state before moving on to a successful/ready state. This is an intermittent but expected behavior, and will usually resolve itself.

Workaround: Wait and watch status messages or logs for more information.

([BZ#2004027](#))

Tenant config does not override `backendpath` if the key is specified in upper case

Key Management Service (KMS) provider options set in a Tenants namespace is more advanced than the key/value settings that the OpenShift Container Storage user interface supports. As a result, the configuration options for KMS providers set in the Tenants namespace need to be formatted as camel case, instead of upper case. It might be confusing for the users that have access to the KMS provider configuration in the `openshift-storage` namespace, and the configuration in a Tenants namespace as options in the `openshift-storage` namespace are in upper case, whereas the options in the Tenants namespace are in camel case.

Workaround: Use camel case formatting for the KMS provider options.

([BZ#2005801](#))

Deleting a protected application that has been failed over and later relocated does not delete the RADOS block device image on the secondary or failover site

Deleting a disaster recovery (DR) protected workload may leak RADOS block device (RBD) images on the secondary DR cluster. The deleted images would then occupy space on the secondary cluster. To resolve this issue, use a toolbox pod to detect and clean up the images on the secondary cluster that are no longer in use for DR protection. This workaround ensures space reclamation on the secondary cluster.

([BZ#2005919](#))

Failover action reports RADOS block device image mount failed on the pod with RPC error still in use

Failing over a disaster recovery (DR) protected workload may result in pods using the volume on the failover cluster to be stuck in reporting RADOS block device (RBD) image is still in use. This prevents the pods from starting up for a long duration (upto several hours).

([BZ#2007376](#))

Relocate action results in PVC's in Termination state and workload is not moved to a preferred cluster

While relocating a disaster recovery (DR) protected workload, results in workload not stopping on the

current primary cluster and PVCs of the workload remaining in the terminating state. This prevents pods and PVCs from being relocated to the preferred cluster. To recover the issue perform a failover action to move the workload to the preferred cluster. The workload would be recovered on the preferred cluster but may include a loss of data as the action is a failover.

[\(BZ#2019931\)](#)

Failover action reports RADOS block device image mount failed on the pod with RPC error fsck

Failing over a disaster recovery (DR) protected workload may result in pods not starting with volume mount errors that state the volume has file system consistency check (fsck) errors. This prevents the workload from failing over to the failover cluster.

[\(BZ#2021460\)](#)

Overprovision Level Policy Control does not support custom storage class

OpenShift Data Foundation limits the allowed storage classes in **overprovision-control** to Ceph subtypes only. As a result, if a user-defined storage class is used in the **overprovision-control**, the **StorageCluster** CRD is defined as invalid and that storage class cannot have the **overprovision-control**.

[\(BZ#2024545\)](#)

CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES

8.1. RHBA-2022:0865 OPENSIFT DATA FOUNDATION 4.9.4 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.9.4 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:0865](#) advisory.

8.2. RHBA-2022:0684 OPENSIFT DATA FOUNDATION 4.9.3 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.9.3 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:0684](#) advisory.

8.3. RHBA-2022:0346 OPENSIFT DATA FOUNDATION 4.9.2 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.9.2 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:0346](#) advisory.

Documentation updates

Added a chapter on how to enable console plugin if it is disabled. The console plugin provides a custom interface that is included in the Web Console. You can enable the console plugin option either from the graphical user interface (GUI) or command-line interface. For more information, see [Enabling the Red Hat OpenShift Data Foundation console plugin](#).

8.4. RHSA-2022:0032 OPENSIFT DATA FOUNDATION 4.9.1 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.9.1 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:0032](#) advisory.