



Red Hat OpenShift Data Foundation 4.12

4.12 Release notes

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

Red Hat OpenShift Data Foundation 4.12 4.12 Release notes

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

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

The release notes for Red Hat OpenShift Data Foundation 4.12 summarizes all new features and enhancements, notable technical changes, and any known bugs upon general availability.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	4
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
CHAPTER 1. OVERVIEW	6
1.1. ABOUT THIS RELEASE	6
CHAPTER 2. NEW FEATURES	7
2.1. GENERAL AVAILABILITY OF METROPOLITAN DISASTER RECOVERY (METRO-DR) SOLUTION	7
2.2. GENERAL AVAILABILITY OF LOGICAL VOLUME MANAGER STORAGE FOR SINGLE NODE OPENSIFT CLUSTERS	7
CHAPTER 3. ENHANCEMENTS	8
3.1. SINGLE STACK IPV6 SUPPORT	8
3.2. SUPPORT FOR KMS PROVIDERS USING KMIP	8
3.3. ADJUSTING VERBOSITY LEVELS OF LOGS	8
3.4. ENCRYPTION IN TRANSIT	8
3.5. SUPPORT RESOURCE MODIFICATION FOR MULTICLOUD OBJECT GATEWAY PV POOL PODS	8
3.6. SECURE MODE DEPLOYMENT FOR MULTICLOUD OBJECT GATEWAY	8
3.7. CHANGE IN DEFAULT PERMISSION AND FSGROUPPOLICY	8
CHAPTER 4. TECHNOLOGY PREVIEWS	10
4.1. DISASTER RECOVERY SOLUTIONS FOR OPENSIFT WORKLOADS	10
CHAPTER 5. DEVELOPER PREVIEWS	11
5.1. REPLICAS 1 (NON RESILIENT POOL)	11
5.2. NETWORK FILE SYSTEM NEW CAPABILITIES	11
5.3. ALLOW ROOK-CEPH-OPERATOR-CONFIG ENVIRONMENTAL VARIABLES TO CHANGE DEFAULTS ON UPGRADE	11
5.4. EASY CONFIGURATION OF CEPH TARGET SIZE RATIOS	11
5.5. EPHEMERAL STORAGE FOR PODS	11
5.6. MULTISITE CONFIGURATIONS FOR RGW IN OPENSIFT DATA FOUNDATION	12
5.7. MULTICLOUD OBJECT GATEWAY (MCG) ONLY ON SINGLE NODE CLUSTER	12
5.8. USING TRUSTED CERTIFICATES TO ENSURE TRANSACTIONS ARE SECURE AND PRIVATE	12
CHAPTER 6. BUG FIXES	13
6.1. DISASTER RECOVERY	13
6.2. MULTICLOUD OBJECT GATEWAY	13
6.3. CEPHFS	13
6.4. OPENSIFT DATA FOUNDATION OPERATOR	14
CHAPTER 7. KNOWN ISSUES	15
7.1. DISASTER RECOVERY	15
7.2. CEPHFS	18
7.3. OPENSIFT DATA FOUNDATION CONSOLE	18
CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES	19
8.1. RHBA-2024:3893 OPENSIFT DATA FOUNDATION 4.12.13 BUG FIXES AND SECURITY UPDATES	19
8.2. RHBA-2024:1673 OPENSIFT DATA FOUNDATION 4.12.12 BUG FIXES AND SECURITY UPDATES	19
8.3. RHBA-2024:0630 OPENSIFT DATA FOUNDATION 4.12.11 BUG FIXES AND SECURITY UPDATES	19
8.4. RHSA-2023:7820 OPENSIFT DATA FOUNDATION 4.12.10 BUG FIXES AND SECURITY UPDATES	19
8.5. RHBA-2023:6169 OPENSIFT DATA FOUNDATION 4.12.9 BUG FIXES AND SECURITY UPDATES	19
8.6. RHSA-2023:5377 OPENSIFT DATA FOUNDATION 4.12.8 BUG FIXES AND SECURITY UPDATES	19
8.7. RHBA-2023:4836 OPENSIFT DATA FOUNDATION 4.12.7 BUG FIXES AND SECURITY UPDATES	19

8.8. RHBA-2023:4718 OPENSIFT DATA FOUNDATION 4.12.6 BUG FIXES AND SECURITY UPDATES	19
8.9. RHSA-2023:4287 OPENSIFT DATA FOUNDATION 4.12.5 BUG FIXES AND SECURITY UPDATES	20
8.10. RHSA-2023:3609 OPENSIFT DATA FOUNDATION 4.12.4 BUG FIXES AND SECURITY UPDATES	20
8.11. RHSA-2023:3265 OPENSIFT DATA FOUNDATION 4.12.3 BUG FIXES AND SECURITY UPDATES	20
8.12. RHBA-2023:1816 OPENSIFT DATA FOUNDATION 4.12.2 BUG FIXES AND SECURITY UPDATES	20
8.13. RHBA-2023:1170 OPENSIFT DATA FOUNDATION 4.12.1 BUG FIXES AND SECURITY UPDATES	20
8.13.1. New Feature	20
8.13.2. Enhancements	20
8.13.3. Known issues	21

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](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Do let us know how we can make it better. To give feedback:

- For simple comments on specific passages:
 1. Make sure you are viewing the documentation in the *Multi-page HTML* format. In addition, ensure you see the **Feedback** button in the upper right corner of the document.
 2. Use your mouse cursor to highlight the part of text that you want to comment on.
 3. Click the **Add Feedback** pop-up that appears below the highlighted text.
 4. Follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
 1. Go to the [Bugzilla](#) website.
 2. In the **Component** section, choose **documentation**.
 3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
 4. Click **Submit Bug**.

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 technology stack that includes Red Hat Ceph Storage, the Rook.io Operator, and NooBaa's Multicloud Object Gateway technology. OpenShift Data Foundation also supports Logical Volume Manager Storage for single node OpenShift clusters. For more information, see [General availability of logical volume manager storage for single node OpenShift clusters](#).

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 4.12 ([RHBA-2023:0550](#) and [RHBA-2023:0551](#)) is now available. New enhancements, features, and known issues that pertain to OpenShift Data Foundation 4.12 are included in this topic.

Red Hat OpenShift Data Foundation 4.12 is supported on the Red Hat OpenShift Container Platform version 4.12. For more information, see [Red Hat OpenShift Data Foundation Supportability and Interoperability Checker](#).

For Red Hat OpenShift Data Foundation life cycle information, refer to the layered and dependent products life cycle section in [Red Hat OpenShift Container Platform Life Cycle Policy](#) .

CHAPTER 2. NEW FEATURES

This section describes new features introduced in Red Hat OpenShift Data Foundation 4.12.

2.1. GENERAL AVAILABILITY OF METROPOLITAN DISASTER RECOVERY (METRO-DR) SOLUTION

The Metro-DR feature with Red Hat Advanced Cluster Management for Kubernetes 2.7 is now General Available from Red Hat OpenShift Data Foundation version 4.12.1 and higher.

Metro-DR solution ensures protection and business continuity during the unavailability of a data center with no data loss while using multiple clusters synchronous replication. In the public cloud these are similar to protecting from an Availability Zone failure. This solution offers quick recovery of Applications with no data loss.

For more information, see the [planning guide](#) and [Metro-DR solution for OpenShift Data Foundation guide](#).

2.2. GENERAL AVAILABILITY OF LOGICAL VOLUME MANAGER STORAGE FOR SINGLE NODE OPENSIFT CLUSTERS

Logical volume manager storage provides dynamic block storage for the single node OpenShift clusters where resource constraints are more important than feature variety and data resilience. One target application is for Radio Access Networks (RAN) in the Telecommunications market. For more information, see [Installing LVM Storage using RHACM](#).



NOTE

In previous versions, the product was named OpenShift Data Foundation - Logical Volume Manager. With general availability, it has been renamed to logical volume manager storage (LVM Storage or LVMS).

Starting with this release, in addition to dynamic storage, logical volume manager storage provides the following new features:

- Provides the ability to control or restrict the volume group to your preferred disks by enabling you to manually select the local paths of the disks by path or by name. For more information, see [Installing the OpenShift Data Foundation Logical Volume Manager Operator using RHACM](#).
- Provides the ability to install and use logical volume manager storage on single node OpenShift clusters with additional worker nodes. This helps you to use logical volume manager storage on your desired single node OpenShift architecture. For more information, see [Installing the OpenShift Data Foundation Logical Volume Manager Operator using RHACM](#) and [Scaling storage of Single Node OpenShift cluster](#).

CHAPTER 3. ENHANCEMENTS

This section describes the major enhancements introduced in Red Hat OpenShift Data foundation 4.12.

3.1. SINGLE STACK IPV6 SUPPORT

Single Stack IPv6 is now supported in Red Hat OpenShift Data Foundation. For more information, see [Single Stack IPv6 support](#).

3.2. SUPPORT FOR KMS PROVIDERS USING KMIP

This release introduces support for Key Management System (KMS) providers using Key Management Interoperability Protocol (KMIP) which uses client certificate for authentication. Thales CipherTrust Manager works well with OpenShift Data Foundation 4.12. For more information, see [CipherTrust Manager](#).

3.3. ADJUSTING VERBOSITY LEVELS OF LOGS

The amount of space consumed by debugging logs can become a significant issue. With this update, it is possible to adjust and therefore control the amount of storage that can be consumed by debugging logs as the space consumed by the debugging logs can be a significant issue at times. For more information, see [Adjusting verbosity level of logs](#).

3.4. ENCRYPTION IN TRANSIT

With this enhancement, the IPsec framework provides *Encryption in transit* for a virtualized network that is used for pods and services. The virtualized network is provided by the Open Virtual Network (OVN)-Kubernetes Container Network Interface (CNI) plug-in. For more information, see [Encryption in transit](#).

3.5. SUPPORT RESOURCE MODIFICATION FOR MULTICLOUD OBJECT GATEWAY PV POOL PODS

This enhancement enables you to fine-tune the performance of backingstores that are based on Multicloud Object Gateway (MCG) persistent volume (PV) pools. It provides the ability to modify the CPU and memory resource and limit for PV pool based backingstores to improve MCG's performance for their workloads.

For more information, see [Creating a local Persistent Volume-backed backingstore](#) .

3.6. SECURE MODE DEPLOYMENT FOR MULTICLOUD OBJECT GATEWAY

With this enhancement, it is possible to deploy Multicloud Object Gateway (MCG) in a secure mode and restricts any external access. This provides fine grained control over subnets that have access to MCG deployment. For more information, see [Enabling secure mode deployment for Multicloud Object Gateway](#).

3.7. CHANGE IN DEFAULT PERMISSION AND FSGROUPPOLICY

Permissions of newly created volumes now defaults to a more secure 755 instead of 777. FSGroupPolicy is now set to File (instead of ReadWriteOnceWithFSType in ODF 4.11) to allow application access to

volumes based on FSGroup. This involves Kubernetes using fsGroup to change permissions and ownership of the volume to match user requested fsGroup in the pod's SecurityPolicy.

**NOTE**

Existing volumes with a huge number of files may take a long time to mount since changing permissions and ownership takes a lot of time.

For more information, see this [knowledgebase solution](#).

CHAPTER 4. TECHNOLOGY PREVIEWS

This section describes the technology preview features introduced in Red Hat OpenShift Data Foundation 4.12 under Technology Preview support limitations.



IMPORTANT

Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and 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.

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

4.1. DISASTER RECOVERY SOLUTIONS FOR OPENSIFT WORKLOADS

The OpenShift Data Foundation disaster recovery (DR) capability enables DR across multiple OpenShift Container Platform clusters, and is categorized as follows:

- **Regional disaster recovery (Regional-DR)**
Regional-DR solution provides automated protection for block volumes, asynchronous replication, and protects business functionalities when a disaster strikes at a geographical location. In the public cloud this is similar to protecting from a region failure. For more information, see the [planning guide](#) and [Regional-DR solution for OpenShift Data Foundation](#) guide.
- **Multicluster monitoring in Red Hat Advanced Cluster Management console**
Multicluster monitoring is a single simplified view of storage health and capacity spread across multiple clusters. This multicluster monitoring enables you to manage the storage capacity and monitor the OpenShift Data Foundation clusters from the Red Hat Advanced Cluster Management (RHACM) user interface. This monitoring capability applies to both DR and non-DR clusters. For more information, see [Monitoring multicluster storage health](#).
- **Availability of Regional-DR asynchronously for CephFS volumes**
Regional-DR solution now expands customer DR workload capabilities by adding Regional-DR tasks such as orchestration, failover, and relocate on CephFS volumes using the OpenShift console that is similar to the OpenShift Data Foundation experience with Regional-DR on Ceph RBD volumes. For more information, see the [planning guide](#) and [Regional-DR solution](#) guide.

CHAPTER 5. DEVELOPER PREVIEWS

This section describes the developer preview features introduced in Red Hat OpenShift Data Foundation 4.12.



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.

5.1. REPLICAS 1 (NON RESILIENT POOL)

Applications that manage resiliency at the application level can now use storage class with single replica without data resiliency and high availability.

5.2. NETWORK FILE SYSTEM NEW CAPABILITIES

With this release, OpenShift Data Foundation provides Network File System (NFS) v4.1 and v4.2 service for any internal or external applications. The NFS service helps to migrate data from any environment to the OpenShift environment, for example, data migration from Red Hat Gluster Storage file system to OpenShift environment. NFS features also include volume expansion, snapshot creation and deletion, and volume cloning.

For more information, see [Resource requirements for using Network File system](#) and [Creating exports using NFS](#).

5.3. ALLOW ROOK-CEPH-OPERATOR-CONFIG ENVIRONMENTAL VARIABLES TO CHANGE DEFAULTS ON UPGRADE

This update allows the **rook-ceph-operator-config** environmental variables to change the defaults when OpenShift Data Foundation is upgraded from version 4.5 to another version. This was not possible in the earlier versions.

5.4. EASY CONFIGURATION OF CEPH TARGET SIZE RATIOS

With this update, it is possible to change the target size ratio for any pool. In the previous versions, the pools deployed by rook in the Ceph cluster were assigned a **target_ratio** of **0.49** for both RBD and CephFS data and this could cause an under-allocation of PGs for the RBD pool and an over-allocation of PGs for the CephFS metadata pool. For more information, see [Configuration of pool target size ratios](#).

5.5. EPHEMERAL STORAGE FOR PODS

Ephemeral volume support enable a user to specify ephemeral volumes in its pod specification and tie the lifecycle of the PVC with the pod.

5.6. MULTISITE CONFIGURATIONS FOR RGW IN OPENSIFT DATA FOUNDATION

This feature supports multisite configurations such as Zone, ZoneGroup, or Realm for internal or external OpenShift Data Foundation clusters. This setup helps to replicate data into different sites and recover the data in case of failure.

5.7. MULTICLOUD OBJECT GATEWAY (MCG) ONLY ON SINGLE NODE CLUSTER

In this release, a lightweight object storage solution is provided for single node OpenShift (SNO) clusters using MCG with backingstore layered on top of local storage. Previously, deployments running on SNO could only use block storage.

5.8. USING TRUSTED CERTIFICATES TO ENSURE TRANSACTIONS ARE SECURE AND PRIVATE

This feature provides *in-transit* encryption for Object Storage between OpenShift Data Foundation and Red Hat Ceph Storage when using external mode. It enables all data to be encrypted in transit and at rest. For more information, see [knowledgebase article](#) on *how to use trusted certificates*.

CHAPTER 6. BUG FIXES

This section describes the notable bug fixes introduced in Red Hat OpenShift Data Foundation 4.12.

6.1. DISASTER RECOVERY

- **async replication can no longer be set to 0**
Previously, you could enter any value for **Sync schedule**. This meant you could set **async** replication to **0**, which caused an error. With this update, a number input has been introduced that does not allow a value lower than 1. **async** replication now works correctly.

([BZ#2114501](#))

- **Deletion of Application now deletes pods and PVCs correctly**
Previously, when deleting an application from the RHACM console, DRPC did not get deleted. Not deleting DRPC leads to not deleting the VRG as well as the VR. If the VRG/VR is not deleted, the PVC finalizer list will not be cleaned up, causing the PVC to stay in a **Terminating** state.

With this update, deleting an application from the RHACM console deletes the required dependent DRPC and related resources on the managed clusters, freeing up the PVCs as well for required garbage collection.

([BZ#2108716](#))

- **Deleting the internal VolumeReplicaionGroup resource from where a workload failed over or relocated from no longer causes errors**
Due to a bug in the disaster recovery (DR) reconciler, during deletion of the internal **VolumeReplicaionGroup** resource on a managed cluster, from where a workload failed over or relocated from, a persistent volume claim (PVC) was attempted to be protected. The resulting cleanup operation did not complete and would report the **PeerReady** condition on the **DRPlacementControl** for the application to be **False**. This meant the application that was failed over or relocated, could not be relocated or failed over again because the **DRPlacementControl** resource was reporting its **PeerReady** condition as **False**.

With this update, during deletion of the internal **VolumeReplicationGroup** resource, a PVC is not attempted to be protected again, thereby avoiding the issue of a stalled cleanup. This results in **DRPlacementControl** reporting **PeerReady** as **True** post auto completion of the cleanup.

([BZ#2116605](#))

6.2. MULTICLOUD OBJECT GATEWAY

- **StorageCluster no longer goes into Error state while waiting for StorageClass creation**
When an Red Hat OpenShift Data Foundation **StorageCluster** is created, it waits for the underlying pools to be created before the **StorageClass** is created. During this time, the cluster returns an error for the reconcile request until the pools are ready. Because of this error, the **Phase** of the **StorageCluster** is set to **Error**. With this update, this error is caught during pool creation, and the **Phase** of the **StorageCluster** is **Progressing**.

([BZ#2004027](#))

6.3. CEPHFS

- **There is no longer an issue with bucket metadata when updating from RHCS 5.1 to a later version**

RADOS Gateway (RGW) as shipped with Red Hat Ceph Storage (RHCS) version 5.1 inadvertently contained logic related to not-yet-GA support for dynamic bucket-index resharding in multisite replication setups. This logic was intentionally removed from RHCS 5.2. A side effect of this history is that sites which have upgraded to RHCS 5.1 cannot upgrade to RHCS 5.2, since version 5.2's bucket metadata handling is not compatible with that of RHCS 5.1. This situation is now resolved with the upgrade to RHCS 5.3. As a result, RHCS 5.3 is able to operate on buckets created in all prior versions, including 5.1.

([BZ#2115645](#))

6.4. OPENSIFT DATA FOUNDATION OPERATOR

- **There is no longer a Pod Security Violation Alert when the ODF operator is installed**

OpenShift Data Foundation version 4.11 introduced new POD Security Admission standards which give warnings on running of privileged pods. The ODF operator deployment uses a few pods which needed privileged access. Because of this, after the ODF operator was deployed, a Pod Security Violation alert started firing.

With this release, OLM now automatically labels the Namespace, which is prefixed by **openshift-**, for relevant Pod security Admission standards.

([BZ#2110628](#))

CHAPTER 7. KNOWN ISSUES

This section describes the known issues in Red Hat OpenShift Data Foundation 4.12.

7.1. DISASTER RECOVERY

- **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 might 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](#))

- **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](#))

- **Creating an application namespace for the managed clusters**

Application namespace needs to exist on RHACM managed clusters for disaster recovery (DR) related pre-deployment actions and hence is pre-created when an application is deployed at the RHACM hub cluster. However, if an application is deleted at the hub cluster and its corresponding namespace is deleted on the managed clusters, they reappear on the managed cluster.

Workaround: **openshift-dr** maintains a namespace **manifestwork** resource in the managed cluster namespace at the RHACM hub. These resources need to be deleted after the application deletion. For example, as a cluster administrator, execute the following command on the hub cluster: **oc delete manifestwork -n <managedCluster namespace> <drPlacementControl name>-<namespace>-ns-mw.**

([BZ#2059669](#))

- **RBD mirror scheduling is getting stopped for some images**

The Ceph manager daemon gets blocklisted due to different reasons, which causes the scheduled RBD mirror snapshot from being triggered on the cluster where the image(s) are primary. All RBD images that are mirror enabled (hence DR protected) do not list a schedule when examined using **rbd mirror snapshot schedule status -p ocs-storagecluster-cephblockpool**, and hence are not actively mirrored to the peer site.

Workaround: Restart the Ceph manager deployment, on the managed cluster where the images are primary, to overcome the blacklist against the currently running instance, this can be done by scaling down and then later scaling up the ceph manager deployment as follows:

```
$ oc -n openshift-storage scale deployments/rook-ceph-mgr-a --replicas=0
$ oc -n openshift-storage scale deployments/rook-ceph-mgr-a --replicas=1
```

Result: Images that are DR enabled and denoted as primary on a managed cluster start reporting mirroring schedules when examined using **rbd mirror snapshot schedule status -p ocs-storagecluster-cephblockpool**

[\(BZ#2067095\)](#)

- **ceph df reports an invalid MAX AVAIL value when the cluster is in stretch mode**

When a crush rule for a Red Hat Ceph Storage cluster has multiple "take" steps, the **ceph df** report shows the wrong maximum available size for the map. The issue will be fixed in an upcoming release.

[\(BZ#2100920\)](#)

- **Ceph does not recognize the global IP assigned by Globalnet**

Ceph does not recognize global IP assigned by Globalnet, so disaster recovery solution cannot be configured between clusters with overlapping service CIDR using Globalnet. Due to this disaster recovery solution does not work when service **CIDR** overlaps.

[\(BZ#2102397\)](#)

- **Both the DRPCs protect all the persistent volume claims created on the same namespace**

The namespaces that host multiple disaster recovery (DR) protected workloads, protect all the persistent volume claims (PVCs) within the namespace for each DRPlacementControl resource in the same namespace on the hub cluster that does not specify and isolate PVCs based on the workload using its **spec.pvcSelector** field.

This results in PVCs, that match the DRPlacementControl **spec.pvcSelector** across multiple workloads. Or, if the selector is missing across all workloads, replication management to potentially manage each PVC multiple times and cause data corruption or invalid operations based on individual DRPlacementControl actions.

Workaround: Label PVCs that belong to a workload uniquely, and use the selected label as the DRPlacementControl **spec.pvcSelector** to disambiguate which DRPlacementControl protects and manages which subset of PVCs within a namespace. It is not possible to specify the **spec.pvcSelector** field for the DRPlacementControl using the user interface, hence the DRPlacementControl for such applications must be deleted and created using the command line.

Result: PVCs are no longer managed by multiple DRPlacementControl resources and do not cause any operation and data inconsistencies.

[\(BZ#2111163\)](#)

- **MongoDB pod is in CrashLoopBackoff because of permission errors reading data in ceph rbd volume**

The OpenShift projects across different managed clusters have different security context constraints (SCC), which specifically differ in the specified UID range and/or **FSGroups**. This leads to certain workload pods and containers failing to start post failover or relocate operations within these projects, due to filesystem access errors in their logs.

Workaround: Ensure workload projects are created on all managed clusters with the same project-level SCC labels, allowing them to use the same filesystem context when failed over or relocated. Pods will no longer fail post-DR actions on filesystem-related access errors.

[\(BZ#2114573\)](#)

- **Application is stuck in Relocating state during relocate**

Multicloud Object Gateway allowed multiple persistent volume (PV) objects of the same name or namespace to be added to the S3 store on the same path. Due to this, Ramen does not restore the PV because it detected multiple versions pointing to the same **claimRef**.

Workaround: Use S3 CLI or equivalent to clean up the duplicate PV objects from the S3 store. Keep only the one that has a timestamp closer to the failover or relocate time.

Result: The restore operation will proceed to completion and the failover or relocate operation proceeds to the next step.

([BZ#2120201](#))

- **Application is stuck in a FailingOver state when a zone is down**

At the time of a failover or relocate, if none of the s3 stores are reachable then the failover or relocate process hangs. If the DR logs indicate that the S3 store is not reachable, then troubleshooting and getting the s3 store operational will allow the DR to proceed with the failover or relocate operation.

([BZ#2121680](#))

- **PeerReady state is set to true when a workload is failed over or relocated to the peer cluster until the cluster from where it was failed over or relocated from is cleaned up**

After a disaster recovery (DR) action is initiated, the **PeerReady** condition is initially set to **true** for the duration when the workload is failed over or relocated to the peer cluster. After this it is set to **false** until the cluster from where it was failed over or relocated from is cleaned up for future actions. A user looking at **DRPlacementControl** status conditions for future actions may recognize this intermediate **PeerReady** state as a peer is ready for action and perform the same. This will result in the operation pending or failing and may require user intervention to recover from.

Workaround: Examine both **Available** and **PeerReady** states before performing any actions. Both should be **true** for a healthy DR state for the workload. Actions performed when both states are true will result in the requested operation progressing

([BZ#2138855](#))

- **Disaster recovery workloads remain stuck when deleted**

When deleting a workload from a cluster, the corresponding pods might not terminate with events such as **FailedKillPod**. This might cause delay or failure in garbage collecting dependent DR resources such as the **PVC**, **VolumeReplication**, and **VolumeReplicationGroup**. It would also prevent a future deployment of the same workload to the cluster as the stale resources are not yet garbage collected.

Workaround: Reboot the worker node on which the pod is currently running and stuck in a terminating state. This results in successful pod termination and subsequently related DR API resources are also garbage collected.

([BZ#2159791](#))

- **Blocklisting can lead to Pods stuck in an error state**

Blocklisting due to either network issues or a heavily overloaded or imbalanced cluster with huge tail latency spikes. Because of this, Pods get stuck in **CreateContainerError** with the message **Error: relabel failed /var/lib/kubelet/pods/cb27938e-f66f-401d-85f0-9eb5cf565ace/volumes/kubernetes.io~csi/pvc-86e7da91-29f9-4418-80a7-4ae7610bb613/mount: lsetxattr /var/lib/kubelet/pods/cb27938e-f66f-401d-85f0-9eb5cf565ace/volumes/kubernetes.io~csi/pvc-86e7da91-29f9-4418-80a7-4ae7610bb613/mount/#ib_16384_0.dblwr: read-only file system.**

Workaround: Reboot the node to which these pods are scheduled and failing by following these steps:

1. Cordon and then drain the node having the issue
2. Reboot the node having the issue
3. Uncordon the node having the issue

([BZ#2094320](#))

7.2. CEPHFS

- **Poor performance of the 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 Data Foundation clusters.

([BZ#1982116](#))

- **SELinux relabelling issue with a very high number of files**

When attaching volumes to pods in Red Hat OpenShift Container Platform, the pods sometimes do not start or take an excessive amount of time to start. This behavior is generic and it is tied to how SELinux relabelling is handled by the Kubelet. This issue is observed with any filesystem based volumes having a very high file counts. In OpenShift Data Foundation, the issue is seen when using CephFS based volumes with a very high number of files. There are different ways to workaround this issue. Depending on your business needs you can choose one of the workarounds from the knowledgebase solution <https://access.redhat.com/solutions/6221251>.

([Jira#3327](#))

7.3. OPENSIFT DATA FOUNDATION CONSOLE

- **OpenShift Data Foundation dashboard crashes after upgrade**

When OpenShift Container Platform and OpenShift Data Foundation are upgraded, the Data Foundation dashboard under the Storage section crashes with a "404: Page not found" error when dashboard link is clicked. This is because the pop-up that refreshes the console does not appear.

Workaround: Perform a hard refresh of the console. This brings back the dashboard and it will no longer crash.

([BZ#2157876](#))

CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES

8.1. RHBA-2024:3893 OPENSIFT DATA FOUNDATION 4.12.13 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.13 is now available. The bug fixes that are included in the update are listed in the [RHBA-2024:3893](#) advisory.

8.2. RHBA-2024:1673 OPENSIFT DATA FOUNDATION 4.12.12 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.12 is now available. The bug fixes that are included in the update are listed in the [RHBA-2024:1673](#) advisory.

8.3. RHBA-2024:0630 OPENSIFT DATA FOUNDATION 4.12.11 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.11 is now available. The bug fixes that are included in the update are listed in the [RHBA-2024:0630](#) advisory.

8.4. RHSA-2023:7820 OPENSIFT DATA FOUNDATION 4.12.10 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.10 is now available. The bug fixes that are included in the update are listed in the [RHSA-2023:7820](#) advisory.

8.5. RHBA-2023:6169 OPENSIFT DATA FOUNDATION 4.12.9 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.9 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:6169](#) advisory.

8.6. RHSA-2023:5377 OPENSIFT DATA FOUNDATION 4.12.8 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.8 is now available. The bug fixes that are included in the update are listed in the [RHSA-2023:5377](#) advisory.

8.7. RHBA-2023:4836 OPENSIFT DATA FOUNDATION 4.12.7 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.7 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:4836](#) advisory.

8.8. RHBA-2023:4718 OPENSIFT DATA FOUNDATION 4.12.6 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.6 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:4718](#) advisory.

8.9. RHSA-2023:4287 OPENSIFT DATA FOUNDATION 4.12.5 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.5 is now available. The bug fixes that are included in the update are listed in the [RHSA-2023:4287](#) advisory.

8.10. RHSA-2023:3609 OPENSIFT DATA FOUNDATION 4.12.4 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.4 is now available. The bug fixes that are included in the update are listed in the [RHSA-2023:3609](#) advisory.

8.11. RHSA-2023:3265 OPENSIFT DATA FOUNDATION 4.12.3 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.3 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:3265](#) advisory.

8.12. RHBA-2023:1816 OPENSIFT DATA FOUNDATION 4.12.2 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.2 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:1816](#) advisory.

8.13. RHBA-2023:1170 OPENSIFT DATA FOUNDATION 4.12.1 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.12.1 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:1170](#) advisory.

8.13.1. New Feature

General availability of Metropolitan disaster recovery (Metro-DR) solution

Red Hat OpenShift Data Foundation Metro-DR feature with Red Hat Advanced Cluster Management for Kubernetes 2.7 is General Available now.



NOTE

The Regional-DR solution for both Blocks and Files is offered as Technology Preview and is subject to Technology Preview support limitations.

For more information, see the [planning guide](#) and [Metro-DR solution for OpenShift Data Foundation guide](#).

8.13.2. Enhancements

Fixed read performance issues as found by COS

The read operations performance of Multicloud Object Gateway database is improved with this enhancement. To achieve this, certain regular expressions that are used by some of the queries that run against the database to serve the required data are pre-compiled. This saves time when running in real-time. ([BZ#2149861](#))

Added missing annotation to CSV for disconnected environment support and RelatedImages field

The **multicluster-orchestrator** operator is listed under operators supporting disconnected mode installations with this enhancement. To display this operator, the disconnected mode support annotation is added to CSV as the user interface (UI) uses this annotation. ([BZ#2166223](#))

8.13.3. Known issues

Cannot initiate failover of application from hub console

While working with an active/passive Hub Metro-DR setup, you might come across a rare scenario where the Ramen reconciler stops running after exceeding its allowed rate-limiting parameters. As reconciliation is specific to each workload, only that workload is impacted. In such an event, all disaster recovery orchestration activities related to that workload stop until the Ramen pod is restarted.

Workaround: Restart the Ramen pod on the Hub cluster.

```
$ oc delete pods <ramen-pod-name> -n openshift-operators
```

([BZ#2175201](#))

Cannot failover applications from console after the repeated Active hub zone failure

During multiple hub recoveries, in the event of a double failure such as when both the hub and managed clusters are going down, you may not be able to initiate a failover from the RHACM console if the last action was relocate.

Workaround: Use the CLI to set the **DRPC.spec.action** field to Failover.

```
$ oc edit drpc -n app-1 app-1-placement-1-drpc
```

```
spec
  action: Failover
```

Result: Failover of the workload will be initiated to the failover cluster.

([BZ#2176028](#))