Red Hat OpenShift Data Foundation 4.13

4.13 Release notes

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

Last Updated: 2024-09-09

Red Hat OpenShift Data Foundation 4.13 4.13 Release notes

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

Legal Notice

Abstract

The release notes for Red Hat OpenShift Data Foundation 4.13 summarizes all new features and enhancements, notable technical changes, and any known bugs upon general availability.

Table of Contents

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.

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 4.13 (RHBA-2023:3734 and RHSA-2023:3742) is now available. New enhancements, features, and known issues that pertain to OpenShift Data Foundation 4.13 are included in this topic.

Red Hat OpenShift Data Foundation 4.13 is supported on the Red Hat OpenShift Container Platform version 4.13. 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 4.13.

2.1. GENERAL AVAILABILITY OF DISASTER RECOVERY WITH STRETCH CLUSTERS SOLUTION

With this release, disaster recovery with stretch clusters is generally available. In a high availability stretch cluster solution, a single cluster is stretched across two zones with a third zone as the location for the arbiter. This solution is deployed in the OpenShift Container Platform on-premise data centers. This solution is designed to be deployed where latencies do not exceed 5 ms between zones, with a maximum round-trip time (RTT) of 10 ms between locations of the two zones that are residing in the main on-premise data centers.

For more information, see Disaster recovery with stretch cluster for OpenShift Data Foundation .

2.2. GENERAL AVAILABILITY OF SUPPORT FOR NETWORK FILE SYSTEM

OpenShift Data Foundation supports the Network File System (NFS) service for any internal or external applications running in any operating system (OS) except Mac and Windows OS. 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.

For more information, see Creating exports using NFS .

2.3. SUPPORT FOR ENABLING IN-TRANSIT ENCRYPTION FOR OPENSHIFT DATA FOUNDATION

With this release, OpenShift Data Foundation provides a security enhancement to secure network operations by encrypting all the data moving through the network and systems. The enhanced security is provided using encryption in-transit through Ceph's messenger v2 protocol.

For more information about how to enable in-transit encryption, see the required Deploying OpenShift Data Foundation guide based on the platform.

2.4. SUPPORT FOR AZURE RED HAT OPENSHIFT

With this release, you can use the unmanaged OpenShift Data Foundation on Microsoft Azure on Red Hat OpenShift, which is a managed OpenShift platform on Azure. However, note that OpenShift 4.12, 4.13 versions are not yet available in Azure on Red Hat OpenShift, hence the support here is only for OpenShift Data Foundation 4.10 and 4.11.

For more information, see Deploying OpenShift Data Foundation 4.10 using Microsoft Azure and Azure Red Hat OpenShift and Deploying OpenShift Data Foundation 4.11 using Microsoft Azure and Azure Red Hat OpenShift.

2.5. SUPPORT AGNOSTIC DEPLOYMENT OF OPENSHIFT DATA FOUNDATION ON ANY OPENSHIFT SUPPORTED PLATFORM

This release supports and provides a flexible hosting environment for seamless deployment and upgrade of OpenShift Data Foundation.

For more information, see Deploying OpenShift Data Foundation on any platform.

2.6. SUPPORT INSTALLER PROVISIONED INFRASTRUCTURE DEPLOYMENT OF OPENSHIFT DATA FOUNDATION USING BARE METAL INFRASTRUCTURE

With this release, installer provisioned infrastructure deployment of OpenShift Data Foundation using bare metal infrastructure is fully supported.

For more information, see Deploying OpenShift Data Foundation using bare metal infrastructure and Scaling storage.

2.7. OPENSHIFT DATA FOUNDATION TOPOLOGY IN OPENSHIFT CONSOLE

OpenShift Data Foundation topology provides administrators with rapid observability into important cluster interactions and overall cluster health. This improves the customer experience and their ability to streamline operations to effectively leverage OpenShift Data Foundation to its maximum capabilities.

For more information, see the *View OpenShift Data Foundation Topology* section in any of the Deploying OpenShift Data Foundation guides based on the platform.

2.8. GENERAL AVAILABILITY OF PERSISTENT VOLUME ENCRYPTION - SERVICE ACCOUNT PER NAMESPACE

OpenShift Data Foundation now provides access to a service account in every OpenShift Container Platform namespace to authenticate with Vault using a Kubernetes service account token. The service account is thus used for KMS authentication for encrypting Persistent Volumes.

For more information, see Data encryption options and Configuring access to KMS using vaulttenantsa.

2.9. SUPPORT OPENSHIFT DUAL STACK WITH ODF USING IPV4

In Openshift Data Foundation single stack, you can either use IPv4 or IPv6. In case OpenShift is configured with dual stack, OpenShift Data Foundation uses IPv4 and this combination is supported.

For more information, see Network requirements.

2.10. SUPPORT FOR BUCKET REPLICATION DELETION

When creating a bucket replication policy, you now have the option to enable deletion so that when data is deleted from the source bucket, the data is deleted from the destination bucket as well. This feature requires logs-based replication, which is currently only supported using AWS.

For more information, see Enabling bucket replication deletion.

2.11. DISASTER RECOVERY MONITORING DASHBOARD

This feature provides reference information to understand the health of disaster recovery (DR) replication relationships such as the following:

- Application level DR health

- Cluster level DR health
- Failover and relocation operation status
- Replication lag status
- Alerts

For more information, see Monitoring disaster recovery health .

CHAPTER 3. ENHANCEMENTS

This section describes the major enhancements introduced in Red Hat OpenShift Data foundation 4.13.

3.1. DISABLE MULTICLOUD OBJECT GATEWAY EXTERNAL SERVICE DURING DEPLOYMENT

With this release, there is an option to deploy OpenShift data Foundation without the Multicloud Object Gateway load balancer service using the command line interface (CLI). You need to use the **disableLoadBalancerService** variable in the **storagecluster** CRD. This provides enhanced security and does not expose services externally to the cluster.

For more information, see the knowledgebase article Install Red Hat OpenShift Data Foundation (previously known as OpenShift Container Storage) 4.X in internal mode using command line interface and Disabling Multicloud Object Gateway external service after deploying OpenShift Data Foundation .

3.2. NETWORK FILE SYSTEM METRICS FOR ENHANCED OBSERVABILITY

Network File System (NFS) metrics dashboard provides observability for NFS mounts such as the following:

- Mount point for any exported NFS shares
- Number of client mounts
- A breakdown statistics of the clients that are connected to help determine internal versus the external client mounts
- Grace period status of the Ganesha server
- Health statuses of the Ganesha server

For more information, see Network File System metrics.

3.3. METRICS TO IMPROVE REPORTING OF UNHEALTHY BLOCKLISTED NODES

With this enhancement, alerts are displayed in OpenShift Web Console to inform about the blocklisted kernel RBD client on a worker node. This helps to reduce any potential operational issue or troubleshooting time.

3.4. ENABLE CEPH EXPORTER WITH LABELED PERFORMANCE COUNTERS IN ROOK

With this enhancement, Ceph exporter is enabled in Rook and provided with labeled performance counters for **rbd-mirror** metrics thereby enhancing scalability for a larger number of images.

3.5. NEW AMAZON WEB SERVICES (AWS) REGIONS FOR MULTICLOUD OBJECT GATEWAY BACKING STORE

With this enhancement, the new regions that were recently added to AWS are included in the list of regions for Multicloud Object Gateway backing store. As a result, it is now possible to deploy default backing store on the new regions.

3.6. ALLOW RBD POOL NAME WITH AN UNDERSCORE OR PERIOD

Previously, creating a storage system in OpenShift Data Foundation using an external Ceph cluster would fail if the RADOS block device (RBD) pool name contained an underscore (_) or a period(.).

With this fix, the Python script (**ceph-external-cluster-details-exporter.py**) is enhanced to contain underscore and period so that an alias for the RBD pool names can be passed in. This alias allows the OpenShift Data Foundation to adopt an external Ceph cluster with RBD pool names containing an underscore(_) or a period(.).

3.7. OSD REPLICAS ARE SET TO MATCH THE NUMBER OF FAILURE DOMAINS

Previously, an unbalanced situation used to occur when the number of repliacas did not match the number of failure domains.

With this enhancement, OSD replicas are set to match the number of failure domains thereby avoiding the imbalance. For example, when a cluster is deployed on 4 zone cluster with 4 nodes, 4 OSD replicas are created.

3.8. 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 4.13 under Technology Preview support limitations.

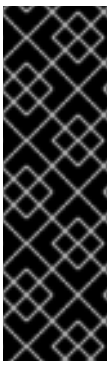

IMPORTANT

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. REGIONAL DISASTER RECOVERY FOR RADOS BLOCK DEVICE

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 regional failure.

For more information, see Regional-DR section in the Planning guide and Regional-DR solution for OpenShift Data Foundation.

CHAPTER 5. DEVELOPER PREVIEWS

This section describes the developer preview features introduced in Red Hat OpenShift Data Foundation 4.13.

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. ALLOW OVERRIDE OF THE DEFAULT NOOBAA BACKING STORE

You can override the default backing store and remove it if you do not want to use the default configuration of backing store by using the **manualDefaultBackingStore** flag. This provides flexibility to customize your backingstore configuration and tailor it to your specific needs. By leveraging this feature, you can further optimize your system and enhance its performance.

For more information, see the knowledgebase article Allow override of default NooBaa backing store .

5.2. ABILITY TO CLONE AND RESTORE VOLUMES ACROSS DIFFERENT CEPH CSI STORAGE CLASSES

You can now clone and restore persistent volume claims (PVCs) across different Ceph CSI managed storage classes. You can take a snapshot of a PVC created with one Ceph CSI storage class and clone it into a PVC created using another Ceph CSI storage class. This ability helps in the following scenarios:

- To reduce the storage footprint
- To move from a storage class with replica 3 to replica 2 or replica 1 where you have your own high availability or underlying storage providing resiliency.
- To clone from a slower storage class to a faster storage class.
- To take a snapshot of an image on a storage class that has lower failure-domain tolerance and clone it to a storage class that has higher failure-domain tolerance

For more information, see the knowledgebase article Ability to clone and restore volumes across different Ceph CSI storage classes.

5.3. ENABLE EXTERNAL MODE USING SSL WITH RADOS GATEWAY

OpenShift Data Foundation now provides in-transit encryption for object storage between OpenShift Data Foundation and Red Hat Ceph Storage when using external mode. This enables Multicloud Object Gateway to use SSL when in external mode.

For more information, see the knowledgebase article Setting up TLS-enabled RGW in external Red Hat Ceph Storage for OpenShift Data Foundation.

5.4. ALLOW OCS-OPERATOR TO DEPLOY ACTIVE AND STANDBY MGR PODS

OpenShift Data Foundation now allows ocs-operator to deploy two MGR pods, one of them is active and the other is a standby. This helps to improve the reliability of MGR pods.

For more information, see the knowledgebase article Allowing ocs-operator to deploy two MGR PODs active and standby.

5.5. NETWORK FILE SYSTEM SUPPORT FOR ACTIVE DIRECTORY AND FREEIPA

In addition to LDAP, OpenShift Data Foundation provides Network File System (NFS) support for Active Directory and FreeIPA. This helps with better control of access management and leveraging of existing tools.

For more information, see the knowledgebase article Setting up Ceph NFS-Ganesha and using a Windows AD for user ID lookups.

CHAPTER 6. BUG FIXES

This section describes the notable bug fixes introduced in Red Hat OpenShift Data Foundation 4.13.

6.1. MULTICLOUD OBJECT GATEWAY

- **Reconcile of disableLoadBalancerService field is ignored in OpenShift Data Foundation operator**
  Previously, any change to the **disableLoadBalancerService** field for Multicloud Object Gateway (MCG) was overridden due to the OpenShift Data Foundation operator reconciliation.

  With this fix, reconcile of **disableLoadBalancerService** field is ignored in OpenShift Data Foundation operator and, as a result, any value set for this field in NooBaa CR is retained and not overridden.

  (BZ#2186171)

- **Performance improvement for non optimized database related flows on deletions**
  Previously, non optimized database related flows on deletions caused Multicloud Object Gateway to spike in CPU usage and perform slowly on mass delete scenarios. For example, reclaiming a deleted object bucket claim (OBC).

  With this fix, indexes for the bucket reclaimer process are optimized, a new index is added to the database to speed up the database cleaner flows, and bucket reclaimer changes are introduced to work on batches of objects.

  (BZ#2181535)

- **OpenShift generated certificates used for MCG internal flows to avoid errors**
  Previously, there were errors in some of the Multicloud Object Gateway (MCG) internal flows due to self-signed certificate that resulted in failed client operations. This was due to the use of self-signed certification in internal communication between MCG components.

  With this fix, OpenShift Container Platform generated certificate is used for internal communications between MCG components, thereby avoiding the errors in the internal flows.

  (BZ#2168867)

- **Metric for number of bytes used by Multicloud Object Gateway bucket**
  Previously, there was no metric to show the number of bytes used by Multicloud Object Gateway bucket.

  With this fix, a new metric **NooBaa_bucket_used_bytes** is added, which shows the number of bytes used by the Multicloud Object Gateway bucket.

  (BZ#2168010)

- **Public access disabled for Microsoft Azure blob storage**
  Previously, the default container created in Microsoft Azure was with public access enabled and caused security concerns.

  With this fix, the default container created will not have the public access enabled by default which means **AllowBlobPublicAccess** is set to false.

  (BZ#2167821)

- **Multicloud Object Gateway bucket buckets are deleted even when the replication rules are set**
  Previously, if replication rules were set for a Multicloud Object Gateway bucket, the bucket was not considered to be eligible for deletion, thereby the buckets would stay without getting deleted. With this fix, the replication rules on a specific bucket are updated when the bucket is being deleted and as a result the bucket is deleted.

  (BZ#2168788)

- **Database init container ownership replaced with Kubernetes FSGroup**
  Previously, Multicloud Object Gateway (MCG) failed to come up and serve when **init** container for MCG database (DB) pod failed to change ownership.

  With this fix, DB **init** container ownership is replaced with Kubernetes FSGroup. (BZ#2115616)

6.2. CEPHFS

- **cephfs-top is able to display more than 100 clients**
  Previously, when you tried to load more than 100 clients to **cephfs-top**, in a few instances, it showed a blank screen and went into hung state as **cephfs-top** could not accommodate the clients in the display due to less or no space. Because the clients were displayed based on **x_coord_map** calculations, **cephfs-top** could not accommodate more clients in the display.

  This issue is fixed as a part of another BZ in Ceph when **ncurses** scrolling and a new way of displaying clients were introduced in **cephfs-top**. The **x_coord_map** calculation was also dropped. So, **cephfs-top** now displays 200 or more clients.

  (BZ#2067168)

6.3. CEPH CONTAINER STORAGE INTERFACE (CSI)

- **RBD Filesystem PVC expands even when the StagingTargetPath is missing**
  Previously, the RADOS block device (RBD) Filesystem persistent volume claim (PVC) expansion was not successful when the **StagingTartgetPath** was missing in the **NodeExpandVolume** remote procedure call (RPC) and Ceph CSI was not able to get the device details to expand.

  With this fix, Ceph CSI goes through all the mount references to identify the **StageingTargetPath** where the RBD image is mounted. As a result, RBD Filesystem PVC expands successfully even when the **StagingTargetPath** is missing.

  (BZ#2164617)

- **Default memory and CPU resource limit increased**
  Previously, **odf-csi-addons-operator** had low memory resource limit and as a result the **odf-csi-addons-operator** pod was **OOMKilled** (out of memory).

  With this fix, the default memory and the CPU resource limit has been increased and **odf-csi-addons-operator OOMKills** are not observed.

  (BZ#2172365)

6.4. OPENSHIFT DATA FOUNDATION OPERATOR

- **Two separate routes for secure and insecure ports**

Previously, http request failures occured as route ended up using the secure port because the port in RGW service for its **openshiftroute** was not defined.

With this fix, insecure port for the existing OpenShift for RGW are defined properly and a new route with secure port is created, thereby avoiding the http request failures. Now, two routes are available for RGW, the existing route uses the insecure port and the new separate route uses the secure port.

(BZ#2104148)

- **Reflects correct state of the Ceph cluster in external mode**
  Previously, when OpenShift Data Foundation is deployed in external mode with a Ceph cluster, the negative conditions such as storagecluster **ExternalClusterStateConnected** were not cleared from the storage cluster even when the associated Ceph cluster was in a good state.

  With this fix, the negative conditions are removed from the storage cluster when the Ceph cluster is in a positive state, thereby reflecting the correct state of the Ceph cluster.

  (BZ#2172189)

- **nginx configurations are added through the ConfigMap**
  Previously, when IPv6 was disabled at node's kernel level, **IPv6 listen** directive of **nginx** configuration for the **odf-console** pod gave an error. As a result, OpenShift Data Foundation was stuck with **odf-console** not available and **odf-console** is in **CrashLoopBackOff** errors.

  With this fix, all the **nginx** configurations are added through the ConfigMap created by the **odf-operator**.

  (BZ#2173161)

6.5. OPENSHIFT DATA FOUNDATION CONSOLE

- **User interface correctly passes the PVC name to the CR**
  Previously, while creating NamespaceStore in the user interface (UI) using file system, the UI would pass the entire persistent volume claim (PVC) object to the CR instead of just the PVC name that is required to be passed to the CR's **spec.nsfs.pvcName** field. As a result, an error was seen on the UI.

  With this fix, only the PVC name is passed to the CR instead of the entire PVC object.

  (BZ#2158922)

- **Refresh popup is shown when OpenShfit Data Foundation is upgraded**
  Previously, when OpenShfit Data Foundation was upgraded, OpenShift Container Platform did not show the **Refresh** button due to lack of awareness about the changes. OpenShift used to not perform checks to know the changes in the **version** field of the **plugin-manifest.json** file present in the **odf-console** pod.

  With this fix, OpenShift Container Platform and OpenShift Data Foundation are configured to poll the manifest for OpenShift Data Foundation user interface. Based on the change in version a **Refresh** popup is shown.

  (BZ#2157876)

6.6. ROOK

- **StorageClasses are created even if the RGW endpoint is not reachable** Previously, in OpenShift Data Foundation external mode deployment, if the RADOS gateway (RGW) endpoints were not reachable and Rook fails to configure the CephObjectStore, creation of RADOS block device (RBD) and CephFS also would fail as these were tightly coupled in the python script,**create-external-cluster-resources.py**.
  With this fix, the issues in the python script was fixed to make separate calls instead of failing or showing errors and the StorageClasses are created.

  (BZ#2139451)

CHAPTER 7. KNOWN ISSUES

This section describes the known issues in Red Hat OpenShift Data Foundation 4.13.

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 blocklist 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 cephrbd 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)

- **PeerReady state is set totrue 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)

- **Application failover hangs in FailingOver state when the managed clusters are on different versions of OpenShift Container Platform and OpenShift Data Foundation**
  Disaster Recovery solution with OpenShift Data Foundation 4.13 protects and restores persistent volume claim (PVC) data in addition to the persistent volume (PV) data. If the primary cluster is on an older OpenShift Data Foundation version and the target cluster is updated to 4.13 then the failover will be stuck as the S3 store will not have the PVC data.

  Workaround: When upgrading the Disaster Recovery clusters, the primary cluster must be upgraded first and then the post-upgrade steps must be run.

  (BZ#2214306)

- **When DRPolicy is applied to mutiple applications under same namespace, volume replication group is not created**
  When a DRPlacementControl (DRPC) is created for applications that are co-located with other applications in the namespace, the DRPC has no label selector set for the applications. If any subsequent changes are made to the label selector, the validating admission webhook in the OpenShift Data Foundation Hub controller rejects the changes.

  Workaround: Until the admission webhook is changed to allow such changes, the DRPC **validatingwebhookconfigurations** can be patched to remove the webhook:

  ```
  $ oc patch validatingwebhookconfigurations vdrplacementcontrol.kb.io-lq2kz --type=json --patch='[{"op": "remove", "path": "/webhooks"}]'
  ```

  (BZ#2210762)

- **Deletion of certain managed resources associated with subscription-based workload**
  During hub recovery, OpenShift Data Foudation encounters a known issue with Red Hat Advanced Cluster Management version 2.7.4 (or higher) where certain managed resources associated with the subscription-based workload might be unintentionally deleted. Currently, this issue does not have any known workaround.

  (BZ#2211643)

- **Peer ready status shown as Unknown after hub recovery**
  After the zone failure and hub recovery, occasionally, the peer ready status of the subscription and application set applications in their disaster recovery placement control (DRPC) is shown as **Unknown**. This is a cosmetic issue and does not impact the regular functionality of Ramen and is limited to the visual appearance of the DRPC output when viewed using the **oc** command.

  Workaround: Use the YAML output to know the correct status:

  ```
  $ oc get drpc -o yaml
  ```

  (BZ#2211883)

7.1.1. DR upgrade

This section describes the issues and workarounds related to upgrading Red Hat OpenShift Data Foundation from version 4.12 to 4.13 in disaster recovery environment.

- **Failover or relocate is blocked for workloads that existed prior to upgrade**

OpenShift Data Foundation Disaster Recovery solution protects persistent volume claim (PVC) data in addition to the persistent volume (PV) data. A failover or a relocate is blocked for workloads that existed prior to the upgrade as they do not have the PVC data backed up.

Workaround:

1. Perform the following steps on the primary cluster for each workload to ensure that the OpenShift Data Foundation Disaster Recovery solution backs up the PVC data in addition to PV data:

   a. Edit the volume replication group (VRG) status:

   ```
   $ oc edit --subresource=status vrg -n <namespace>  <name>
   ```

   b. Set the ClusterDataProtected status to **False** for each **vrg.Status.ProtectedPVCs** conditions. This populates the S3 store with applications PVCs.

   c. Restart the Ramen pod by scaling it down to 0 and back to 1.

   d. Ensure that the S3 store is populated with the application PVCs by waiting for the **ClusterDataProtected** status of all the **vrg.Status.ProtectedPVCs** conditions to turn to **TRUE** again.

2. When you want to initiate the failover or relocate operation, perform the following step:

   a. Clear the **status.preferredDecision.ClusterNamespace** field by editing the DRPC status subresource before initiating a failover or a relocate operation (if not already edited):

   ```
   $ oc edit --subresource=status drpc -n <namespace>  <name>
   ```

   (BZ#2215462)

- **Incorrect value cached status.preferredDecision.ClusterNamespace**
  When OpenShift Data Foundation is upgraded from version 4.12 to 4.13, the disaster recovery placement control (DRPC) might have incorrect value cached in **status.preferredDecision.ClusterNamespace**. As a result, the DRPC incorrectly enters the **WaitForFencing** PROGRESSION instead of detecting that the failover is already complete. The workload on the managed clusters is not affected by this issue.

  Workaround:

  1. To identify the affected DRPCs, check for any DRPC that is in the state **FailedOver** as CURRENTSTATE and are stuck in the **WaitForFencing** PROGRESSION.

  2. To clear the incorrect value edit the DRPC subresource and delete the line, **status.PreferredCluster.ClusterNamespace**:

  ```
  $ oc edit --subresource=status drpc -n <namespace>  <name>
  ```

  3. To verify the DRPC status, check if the PROGRESSION is in **COMPLETED** state and **FailedOver** as CURRENTSTATE.
  (BZ#2215442)

7.2. CEPH

- **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 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)

- **Ceph becomes unresponsive during net split between data zones**
  In a two-site stretch cluster with arbiter, during netsplit between data zones, the Rook operator throws an error stating **failed to get ceph status**. This happens because when netsplit is induced to zone1 and zone2, zone1 and arbiter, the zone1 is cut off from the cluster for about 20 minutes. However, it is found that the error thrown in the rook-operator log is not relevant and OpenShift Data Foundation operates normally in netsplit situation.

  (BZ#2185180)

- **CephOSDCriticallyFull and CephOSDNearFull alerts not firing as expected**
  The alerts **CephOSDCriticallyFull** and **CephOSDNearFull** do not fire as expected because **ceph_daemon** value format has changed in Ceph provided metrics and these alerts rely on the old value format.

  (BZ#2215239)

7.3. CSI DRIVER

- **Automatic flattening of snapshots does not work**
  When there is a single common parent RBD PVC, if volume snapshot, restore, and delete snapshot are performed in a sequence more than 450 times, it is further not possible to take volume snapshot or clone of the common parent RBD PVC.

  To workaround this issue, instead of performing volume snapshot, restore, and delete snapshot in a sequence, you can use PVC to PVC clone to completely avoid this issue.

  If you hit this issue, contact customer support to perform manual flattening of the final restore PVCs to continue to take volume snapshot or clone of the common parent PVC again.

  (BZ#2232163)

7.4. OPENSHIFT DATA FOUNDATION CONSOLE

- **Unable to initiate failover of applicationSet based applications after hub recovery**
  When a primary managed cluster is down after hub recovery, it is not possible to use the user interface (UI) to trigger failover of **applicationSet** based applications if its last action was **Relocate**.

Workaround: To trigger failover from the command-line interface(CLI), set the **DRPC.spec.action** field to **Failover** as follows:

```
$ oc edit drpc -n openshift-gitops app-placement-drpc
```

```
[...]
spec
  action: Failover
[...]
```

(BZ#2209288)

- **Topology view does not work for external mode**
  In external mode, the nodes are not labelled with OCS label. As a result, the topology view cannot show nodes at the first level.

  (BZ#2213739)

CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES

8.1. RHBA-2024:6399 OPENSHIFT DATA FOUNDATION 4.13.11 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.11 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:6399 advisory.

8.2. RHBA-2024:4358 OPENSHIFT DATA FOUNDATION 4.13.10 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.10 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:4358 advisory.

8.3. RHBA-2024:3865 OPENSHIFT DATA FOUNDATION 4.13.9 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.9 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:3865 advisory.

8.4. RHBA-2024:1657 OPENSHIFT DATA FOUNDATION 4.13.8 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.8 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:1657 advisory.

8.5. RHBA-2024:0540 OPENSHIFT DATA FOUNDATION 4.13.7 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.7 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:0540 advisory.

8.6. RHBA-2023:7870 OPENSHIFT DATA FOUNDATION 4.13.6 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.6 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:7870 advisory.

8.7. RHBA-2023:7775 OPENSHIFT DATA FOUNDATION 4.13.5 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.5 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:7775 advisory.

8.8. RHBA-2023:6146 OPENSHIFT DATA FOUNDATION 4.13.4 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.4 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:6146 advisory.

8.9. RHSA-2023:5376 OPENSHIFT DATA FOUNDATION 4.13.3 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.3 is now available. The bug fixes that are included in the update are listed in the RHSA-2023:5376 advisory.

8.10. RHBA-2023:4716 OPENSHIFT DATA FOUNDATION 4.13.2 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.2 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:4716 advisory.

8.11. RHSA-2023:4437 OPENSHIFT DATA FOUNDATION 4.13.1 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.13.1 is now available. The bug fixes that are included in the update are listed in the RHSA-2023:4437 advisory.