Red Hat OpenShift Data Foundation 4.14

4.14 Release notes

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

Last Updated: 2024-04-01

Red Hat OpenShift Data Foundation 4.14 4.14 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 4.14 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.

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 is designed for FIPS. When running on RHEL or RHEL CoreOS booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries submitted to NIST for FIPS Validation on only the x86_64, ppc64le, and s390X architectures. For more information about the NIST validation program, see Cryptographic Module Validation Program. For the latest NIST status for the individual versions of the RHEL cryptographic libraries submitted for validation, see Compliance Activities and Government Standards.

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 4.14 (RHSA-2023:6832) is now available. New enhancements, features, and known issues that pertain to OpenShift Data Foundation 4.14 are included in this topic.

Red Hat OpenShift Data Foundation 4.14 is supported on the Red Hat OpenShift Container Platform version 4.14. 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 4.14.

2.1. GENERAL AVAILABILITY OF REGIONAL DISASTER RECOVERY

The regional disaster recovery (Regional-DR) is generally available for block and file applications. Improvements made to Regional-DR in this release includes the following fixes among other fixes:

- Ceph improvements that enable Regional-DR to be deployed at scale
- Managing DR with ACM UI for both block and file applications
- Monitoring with new DR metrics

Regional-DR is supported only with OpenShift Data Foundation 4.14 and Red Hat Advanced Cluster Management for Kubernetes 2.9 combinations. Support for existing pre OpenShift Data Foundation 4.14 deployments upgrade to Regional-DR is currently in progress and expected to be available in the near future.

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

2.2. IPV6 AUTO-DETECTION

Red Hat OpenShift Data Foundation version 4.14 introduces IPv6 auto detection and configuration. OpenShift clusters using IPv6 or dual-stack are automatically configured in OpenShift Data Foundation accordingly.

For more information about IPv6, see IPv6 support.

2.3. SUPPORT FOR METRO-DR AND REGIONAL-DR SOLUTION FOR OPENSHIFT DATA FOUNDATION ON IBM Z AND IBM POWER

Red Hat OpenShift Data Foundation version 4.14 now supports the Metro-DR and Regional-DR solution on IBM Z and IBM Power platforms. With the enablement of disaster recovery, business continuity is maintained when any disaster hits a geographical location or a data center. Red Hat Ceph Storage (RHCS) deployment is supported only on x86 architecture on IBM Z and IBM Power.

For more information, see Configuring OpenShift Data Foundation Disaster Recovery for OpenShift Workloads.

2.4. LOG-BASED BUCKET REPLICATION

With this release, Multicloud Object Gateway (MCG) bucket replication is scalable. This helps to replicate data faster on a larger scale. The bucket replication uses the event logs of the Amazon Web Services (AWS) and Microsoft Azure cloud environments to optimize the replication.

For more information, see Enabling log based bucket replication in AWS and Enabling log based bucket replication in Microsoft Azure.

2.5. AUTOSCALING SUPPORT FOR MULTICLOUD OBJECT GATEWAY ENDPOINTS

With this release, two new autoscalers based on HPAV2 (default) and KEDA are available. These autoscalers provide support for MCG endpoint autoscaling using Prometheus metrics.

KEDA is not supported on IBM Power as KEDA images are not available for Power architecture.

2.6. DELETION OF EXPIRED OBJECTS IN THE MULTICLOUD OBJECT GATEWAY LIFECYCLE

Deletion of expired objects is a simplified way that enables handling of unused data. This reduces the storage costs due to accumulated data objects. The unused data is deleted after expiration. Data expiration is a part of Amazon Web Services (AWS) lifecycle management and sets an expiration date for automatic deletion. The minimal time resolution of the lifecycle expiration is one day.

For more information, see Lifecyle bucket configuration in Multicloud Object Gateway .

2.7. SUPPORT FOR STANDALONE MULTICLOUD OBJECT GATEWAY

With this release, you can deploy the Multicloud Object Gateway component using the local storage devices on IBM Z. Deploying only the Multicloud Object Gateway component with the OpenShift Data Foundation provides the flexibility in deployment and helps to reduce the resource consumption.

2.8. GENERAL AVAILABILITY OF MULTI-NETWORK PLUG-IN (MULTUS) SUPPORT

With this release, multi-network plug-in, Multus, is generally available. OpenShift Data Foundation supports the ability to use Multus on bare metal infrastructures to improve security and performance by isolating the different types of network traffic. By using Multus, one or more network interfaces on hosts can be reserved for exclusive use of OpenShift Data Foundation.

2.9. GOOGLE CLOUD GENERAL AVAILABILITY SUPPORT

Deployment of OpenShift Data Foundation is now supported on Google Cloud. For more information see, the Deploying OpenShift Data Foundation using Google Cloud guide .

CHAPTER 3. ENHANCEMENTS

This section describes the major enhancements introduced in Red Hat OpenShift Data foundation 4.14.

3.1. SUPPORT FOR HIGHER DISK CAPACITIES AND DISK QUANTITIES

Previously, for local storage deployments, Red Hat recommended 9 devices or less per node and disks size of 4 TiB or less. With this update, the recommended devices per node is now 12 or less, and disks size is 16 TiB or less.

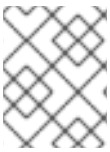

NOTE

Confirm the estimated recovery time using the OpenShift Data Foundation Recovery Calculator. It is recommended that the recovery time for host failure to be under 2 hours.

3.2. FASTER RWO RECOVERY IN CASE OF NODE FAILURES

Previously, it took a long time for ReadWriterOnce (RWO) volumes to recover in case of node failures. With this update, the issue has been fixed.

For the cluster to automatically address node failures and recover RWO volumes faster, manually add one of the following taints to the node:

- **node.kubernetes.io/out-of-service=nodeshutdown:NoExecute**
- **node.kubernetes.io/out-of-service=nodeshutdown:NoSchedule**

3.3. AUTOMATIC SPACE RECLAIMING FOR RBD PERSISTENT VOLUME CLAIMS PVCS

Red Hat OpenShift Data Foundation version 4.14 introduces automatic space reclaiming for RBD persistent volume claims PVCs that are in the namespace that begins with **openshift-**. This means administrators no longer have to manually reclaim space for the RBD PVCs in the namespace that starts with **openshift-** prefix.

3.4. AUTOMATION OF ANNOTATING ENCRYPTED RBD STORAGE CLASSES

Annotation is automatically set when the OpenShift console creates a RADOS block device (RBD) storage class with encryption enabled. This enables customer data integration (CDI) to use host-assisted cloning instead of the default smart cloning.

3.5. LSOS LOCALVOLUMESET AND LOCALVOLUMEDISCOVERY CRS NOW SUPPORT MPATH DEVICE TYPES

With this release, **disk** and **mpath** device types are available for LocalVolumeSet and LocalVolumeDiscovery CRs.

3.6. AUTOMATIC DETECTION OF DEFAULT STORAGECLASS FOR OPENSHIFT VIRTUALIZATION WORKLOADS

OpenShift Data Foundation deployments using OpenShift Virtualization platform will now have a new StorageClass automatically created and it can be set as a default storage class for OpenShift Virtualization. This new StorageClass is optimized for OpenShift virtualization using a specific preset of the underlying storage.

3.7. COLLECT RBD STATUS DETAILS FOR ALL IMAGES

When troubleshooting certain RBD related problems, the status of the RBD-images is an important information. With this release, for the OpenShift Data Foundation internal mode deployment, **odf-must-gather** includes the **rbd status** details, making it faster to troubleshoot RBD related problems.

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 4.14 under Technology Preview support limitations.

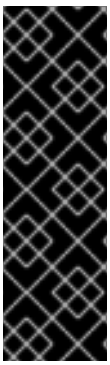

IMPORTANT

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. ALLOW STORAGE CLASSES TO USE NON RESILIENT POOL SINGLE REPLICA

OpenShift Data Foundation allows the option to create a new storage class using non resilient pool with a single replica. This avoids redundant data copies and allows resiliency management on the application level.

For more information, see the deployment guide for your platform on the customer portal.

4.2. METRO-DR SUPPORT FOR WORKLOADS BASED ON OPENSHIFT VIRTUALIZATION

You can now easily set up Metro-DR to protect your workloads based on OpenShift Virtualization using OpenShift Data Foundation.

For more information, see Knowledgebase article.

CHAPTER 5. DEVELOPER PREVIEWS

This section describes the developer preview features introduced in Red Hat OpenShift Data Foundation 4.14.

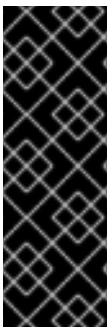

IMPORTANT

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. CUSTOM TIMEOUTS FOR THE RECLAIM SPACE OPERATION

OpenShift Data Foundation now allows you to set a custom timeout value for the reclaim space operation. Previously, depending on RBD volume size and its data pattern, the reclaim space operation might have failed with the error **context deadline exceeded**. Adjusting the timeout value avoids this error.

For more information, see the knowledgebase article Customize timeouts for Reclaim Space Operation .

5.2. EXPANSION OF ENCRYPTED RBD VOLUMES

OpenShift Data Foundation now enables resize capability on encrypted RBD persistent volume claims (PVCs).

For more information, see the knowledgebase article Enabling resize for encrypted RBD PVC.

5.3. IPV6 SUPPORT FOR EXTERNAL MODE

OpenShift Data Foundation now allows users to use the IPv6 Red Hat Ceph Storage external standalone Ceph cluster to connect with the IPV6 OpenShift Container Platform cluster. Users can pass the IPv6 endpoints using the same endpoint flags while running the python script.

5.4. NETWORK FILE SYSTEM SUPPORTS EXPORT SHARING ACROSS NAMESPACES

When OpenShift Data Foundation is used to dynamically create an NFS-export, the **PersistentVolumeClaim** is used to access the NFS-export in a pod. It is not immediately possible to use the same NFS-export for a different application in another OpenShift Namespace. You can now create a second PersistentVolume that can be bound to a second PersistentVolumeClaim in another OpenShift namespace.

For more information, see the knowledgebase article ODF provisioned NFS/PersistentVolume sharing between Namespaces.

5.5. DATA COMPRESSION ON THE WIRE

Data compression on the wire helps in multi-availability zones deployment by lowering latency and network costs. It also helps in cases where the network bandwidth is a bottleneck for performance.

For more information, see the knowledgeabase article In-transit compression.

CHAPTER 6. BUG FIXES

This section describes the notable bug fixes introduced in Red Hat OpenShift Data Foundation 4.14.

6.1. DISASTER RECOVERY

- **Blocklisting no longer leads to pods stuck in an error state**
  Previously, blocklisting due to either network issues or a heavily overloaded or imbalanced cluster with huge tail latency spikes. Because of this, pods got 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**.

  With this fix, blocklisting no longer leads to pods stuck in an error state.

  (BZ#2094320)

- **Ceph now recognizes the global IP assigned by Globalnet**
  Previously, Ceph did not recognize global IP assigned by Globalnet, so disaster recovery solutions could not be configured between clusters with overlapping service CIDR using Globalnet. This issue has been fixed, and now the disaster recovery solution works when service CIDR overlaps.

  (BZ#2104971)

- **PeerReady state is no longer 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**
  Previously, after a disaster recovery (DR) action was initiated, the **PeerReady** condition was initially set to **true** for the duration when the workload was failed over or relocated to the peer cluster. After this it was set to **false** until the cluster from where it was failed over or relocated from was cleaned up for future actions. A user looking at **DRPlacementControl** status conditions for future actions may have recognized this intermediate **PeerReady** state as a peer was ready for action and perform the same. This would result in the operation pending or failing and may have required user intervention to recover from.

  With this fix, **PeerReady** state is no longer set to **true** on failed or relocated workloads until the cluster is cleaned up, so there is no longer any confusion for the user.

  (BZ#2138855)

- **The application no longer stays in Cleaningup state when the ACM hub is recovered after a disaster**
  Previously, when the ACM hub was lost during a disaster and was recovered using the backups, VRG ManifestWork and DRPC status were not restored. This caused the application to stay in Cleaningup state.

  With this fix, Ramen now ensures that VRG ManifestWork is part of the ACM backup and rebuilds the DRPC status if it is empty after a hub recovery, and the application successfully migrates to the failover cluster.

  (BZ#2162469)

- **STS based applications can now be relocated as expected**

Previously, relocating STS based applications would fail due to an underlying bug. This bug has been fixed, and relocating STS based applications now works as expected.

(BZ#2224325)

- **Ramen reconciles as expected after hub restore**
  Previously, while working with an active/passive Hub Metro-DR setup, you may have 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 was impacted. In such an event, all disaster recovery orchestration activities related to that workload stopped until the Ramen pod was restarted.

  This issue has been fixed, and Ramen reconciles as expected after hub restore.

  (BZ#2175201)

- **Managed resources are no longer deleted during hub recovery**
  Previously, during hub recovery, OpenShift Data Foundation encountered 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 have been unintentionally deleted.

  This issue has been fixed, and no managed resources are deleted during hub recovery.

  (BZ#2211643)

6.1.1. DR upgrade

This section describes bug fixes related to upgrading Red Hat OpenShift Data Foundation from version 4.13 to 4.14 in disaster recovery environment.

- **Failover or relocate is no longer blocked for workloads that existed prior to upgrade**
  Previously, a failover or a relocate was blocked for workloads that existed prior to the upgrade. This was because OpenShift Data Foundation Disaster Recovery solution protects persistent volume claim (PVC) data in addition to the persistent volume (PV) data, and the workload did not have the PVC data backed up.

  With this fix, failover or relocate is no longer blocked for workloads that existed prior to upgrade.

  (BZ#2229568)

- **DRPC no longer has incorrect values cached**
  Previously, when OpenShift Data Foundation was upgraded, the disaster recovery placement control (DRPC) may have had an incorrect value cached in **status.preferredDecision.ClusterNamespace**. This issue has been fixed, and the incorrect value is no longer cached.

  (BZ#2229567)

6.2. MULTICLOUD OBJECT GATEWAY

- **Virtual-host style is now available on NooBaa buckets**
  Previously, Virtual-host style did not work on NooBaa buckets because the NooBaa endpoints and core were not aware of the port of the DNS configuration.

With this update, the NooBaa operator passes the port of the DNS to the core and endpoints, making Virtual-host style available.

(BZ#2183092)

- **Dummy credentials are no longer printed to the logs**
  Previously, dummy credentials were printed to the logs which could lead to confusion. This bug has been fixed, and the credentials are no longer printed.

  (BZ#2189866)

- **NooBaa now falls back to using a backing store with type pv-pool when credentials are not provided in the limited time**
  When the cloud credentials operator cannot or fails to provide a secret after the cloud credential request was created, for example, before installing NooBaa the cloud credential operator mode was set to manual mode and no additional necessary actions were done. The provided secret includes the credentials needed for creating the target bucket for the default backing store. This means, the default backing store was not created and Noobaa was stuck in phase Configuring.

  With this fix, if the cloud credential request was sent and we could not get the secret in the limited time that was defined (10 minutes), then NooBaa would fall back to using a backing store with type pv-pool. This means the system should be in status Ready and the default backing store should be with type pv-pool.

  (BZ#2242854)

- **Postgresql DB password no longer displayed in clear text in core and endpoint logs**
  Previously, the internal Postgresql client in noobaa-core printed a connections parameters object to the log, and this object contained the password to connect to Postgresql DB.

  With this fix, the password information is omitted from the connection object that is printed to the log, and the messages to the logs contain only the nonsensitive connection details.

  (BZ#2240778)

6.3. CEPH CONTAINER STORAGE INTERFACE (CSI)

- **CSI CephFS and RBD holder Pods no longer use the old cephcsi image after upgrade**
  Previously, after upgrade CSI CephFS and RBD holder Pods were not getting updated because they were using the old **cephcsi** image.

  With this fix, the daemonset object for CSI CephFS and RBD holder is also upgraded, and the CSI holders pods use the latest **cephcsi** image.

  (BZ#2219536)

- **More reliable and controlled resynchronization process**
  Previously, the **resync** command was not triggered effectively leading to sync issues and inability to disable image mirroring. It was because CephCSI had a dependency on the image mirror state to issue **resync** commands which was unreliable due to the unpredictable changes in the state.

  With this fix, when a volume is being demoted, CephCSI saves the timestamp of the image creation. When the **resync** command is issued, CephCSI compares the saved timestamp with the current creation timestamp and the **resync** proceeds only if the timestamps match. Also,

CephCSI examines the state of the images and the last snapshot timestamps to determine whether resync is required or if an error message needs to be displayed. This results in a more reliable and controlled resynchronization process.

(BZ#2165941)

6.4. OPENSHIFT DATA FOUNDATION OPERATOR

- **There is no longer unnecessary network latency because of S3 clients not able to talk to RGW in same zone**
  Previously, when using the Ceph object store, and while requesting transfer to another zone, there was unnecessary network latency because the S3 clients were unable to talk to RGW in the same zone.

  With this fix, the annotation, "service.kubernetes.io/topology-mode" is added to the RGW service so that the request is routed to the RGW server in the same zone. As a result, pods are routed to the nearest RGW service.

  (BZ#2209098)

6.5. OPENSHIFT DATA FOUNDATION CONSOLE

- **Volume type dropdown is removed from the user interface**
  Previously, for the internal OpenShift Data Foundation installations, the user interface showed HDD, SSD, or both in the Volume type drop down for the existing clusters even though the internal installations should have assumed the disks to be SSD.

  With this fix, Volume type dropdown is removed from the user interface and always assume it to be SSD.

  (BZ#2239622)

- **OpenShift Data Foundation Topology rook-ceph-operator deployment now shows correct resources**
  Previously, the owner references for CSI pods and other pods were set to rook-ceph-operator that caused the mapping to show these pods as part of the deployment too.

  With this fix, the mapping pods approach is changed to top down instead of bottom up, which ensures that only the pods that are related to the deployment are shown.

  (BZ#2209251)

- **CSS properties are set to dynamically adjust the height of the resource list to changes in window size**
  Previously, the sidebar of the topology view resources list was not adjusting to its length based on the window size because the CSS properties were not applied properly to the sidebar.

  With this fix, the CSS properties are set to dynamically adjust the height of the resource list to the changes in window size both in full screen and normal screen mode.

  (BZ#2209258)

- **Add capacity operation no longer fails when moving from LSO to default storage classes**
  Previously, the add capacity operation used to fail when moving from LSO to default storage classes because the persistent volumes (PVs) for expansion were not created correctly.

With this fix, the add capacity operation using a non-LSO storage class is not allowed when a storage cluster is initially created using a LSO based storage class.

(BZ#2213183)

- **Resource utilization of OpenShift Data Foundation topology now matches the metrics**
  Previously, the resource utilization of OpenShift Data Foundation topology did not match the metrics because the metrics query used in the sidebar for resources list of nodes and deployment were different.

  With this fix, the metric queries are made the same and as a result the values are same in both the places.

  (BZ#2214033)

- **Topology view for external mode is now disabled**
  Previously, topology view showed a blank screen for external mode as external mode is not supported in topology view.

  With this fix, external mode is disabled and a message is appears instead of the blank screen.

  (BZ#2216707)

- **Topology no longer shows rook-ceph-operator on every node**
  Previously, the topology view showed the Rook-Ceph operator deployment in all the nodes as Rook-Ceph operator deployment is an owner of multiple pods that are actually not related to it.

  With this fix, the mapping mechanism of deployment to node in the topology view is changed and as a result, Rook-Ceph operator deployment is shown only in one node.

  (BZ#2233027)

- **The console user interface no longer shows SDN instead of OVN**
  Previously, the console user interface showed SDN instead of OVN even though OpenShift Container Platform has moved from SDN to OVN.

  With this fix, the text has been changed from SDN to OVN and as a result, the text for managing network shows OVN.

  (BZ#2233731)

- **Resource names must follow the rule, "starts and ends with a lowercase letter or number", or regex returns an error**
  Previously, due to invalid regex validation for input name of object bucket claim (OBC), backing store, blocking pool, namespace store, and bucket class, the rule "starts and ends with a lowercase letter or number" was violated when symbols or capital letter was entered in the beginning of the name.

  With this release, the issue is fixed and if the resource name does not follow the rule, "starts and ends with a lowercase letter or number", regex returns an error.

  (BZ#2193109)

6.6. ROOK

- **ODF monitoring is no longer missing any metric values**

Previously, there was a missing port for the service monitor of ceph-exporter. This meant that Ceph daemons performance metrics were missing.

With this fix, the port for ceph-exporter service monitor has been added, and Ceph daemons performance metrics are visible in prometheus.

(BZ#2221488)

- **OSD pods no longer continue flapping if there is a network issue**
  Previously, if OSD pods started flapping because of a network issue, they would continue flapping. This would adversely impact the system.

  With this fix, flapping OSD pods are marked as down after a certain amount of time, and no longer impact the system.

  (BZ#2223959)

- **MDS are no longer unnecessarily restarted**
  Previously, MDS pods were unnecessarily restarted because the liveness probe restarted the MDS without checking the **ceph fs dump**.

  With this fix, the liveness probe monitors the MDS in **ceph fs dump** and restarts the MDS only if MDS is missed in the dump output.

  (BZ#2234377)

CHAPTER 7. KNOWN ISSUES

This section describes the known issues in Red Hat OpenShift Data Foundation 4.14.

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)

- **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)

- **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)

- **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)

- **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)

- **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 4.14 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 4.14 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 multiple 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)

- **Failover of apps from c1 to c2 cluster hang in FailingOver**
  The failover action is not disabled by Ramen when data is not uploaded to the s3 store due to s3 store misconfiguration.This means the cluster data is not available on the failover cluster during the failover. Therefore, failover cannot be completed.

  Workaround: Inspect the ramen logs after initial deployment to insure there are no s3 configuration errors reported.

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

  (BZ#2248723)

- **Potential risk of data loss after hub recovery**
  A potential data loss risk exists following hub recovery due to an eviction routine designed to clean up orphaned resources. This routine identifies and marks **AppliedManifestWorks** instances lacking corresponding **ManifestWorks** for collection. A hardcoded grace period of one hour is provided. After this period elapses, any resources associated with the **AppliedManifestWork** become subject to garbage collection.

  If the hub cluster fails to regenerate corresponding **ManifestWorks** within the initial one hour window, data loss could occur. This highlights the importance of promptly addressing any issues that might prevent the recreation of **ManifestWorks** post-hub recovery to minimize the risk of data loss.

  (BZ-2252933)

7.1.1. DR upgrade

This section describes the issues and workarounds related to upgrading Red Hat OpenShift Data Foundation from version 4.13 to 4.14 in disaster recovery environment.

- **Incorrect value cached status.preferredDecision.ClusterNamespace**
  When OpenShift Data Foundation is upgraded from version 4.13 to 4.14, 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 is inaccessible after crash or shutdown tests are run**
In a stretch cluster, when a monitor is revived and is in the probing stage for other monitors to receive the latest information such as **MonitorMap** or **OSDMap**, it is unable to enter **stretch_mode** at the time it is in the probing stage. This prevents it from correctly setting the elector's **disallowed_leaders** list.

  Assuming that the revived monitor actually has the best score, it will think that it is best fit to be a leader in the current election round and will cause the election phase of the monitors to get stuck because it will keep proposing itself and will keep getting rejected by the surviving monitors because of the **disallowed_leaders** list. This leads to the monitors getting stuck in election, and Ceph eventually becomes unresponsive.

  To workaround this issue, when stuck in election and Ceph becomes unresponsive, reset the Connectivity Scores of each monitor by using the command:

```
`ceph daemon mon.{name} connection scores reset`
```

  If this doesn't work, restart the monitors one by one. Election will then be unstuck, monitors will be able to elect a leader, form a quorum, and Ceph will become responsive again.

  (BZ#2241937)

- **Ceph reports no active mgr after workload deployment**
After workload deployment, Ceph manager loses connectivity to MONs or is unable to respond to its liveness probe.

This causes the ODF cluster status to report that there is "no active mgr". This causes multiple operations that use the Ceph manager for request processing to fail. For example, volume provisioning, creating CephFS snapshots, and others.

To check the status of the ODF cluster, use the command **oc get cephcluster -n openshift-storage**. In the status output, the **status.ceph.details.MGR_DOWN** field will have the message "no active mgr" if your cluster has this issue.

To workaround this issue, restart the Ceph manager pods using the following commands:

```
# oc scale deployment -n openshift-storage rook-ceph-mgr-a --replicas=0
```

```
# oc scale deployment -n openshift-storage rook-ceph-mgr-a --replicas=1
```

After running these commands, the ODF cluster status reports a healthy cluster, with no warnings or errors regarding **MGR_DOWN**.

(BZ#2244873)

- **CephBlockPool creation fails when custom deviceClass is used in StorageCluster**
  Due to a known issue, CephBlockPool creation fails when custom deviceClass is used in StorageCluster.

  (BZ#2248487)

7.3. OPENSHIFT DATA FOUNDATION CONSOLE

- **Missing NodeStageVolume RPC call blocks new pods from going into Running state**
  NodeStageVolume RPC call is not being issued blocking some pods from going into **Running** state. The new pods are stuck in **Pending** forever.

  To workaround this issue, scale down all the affected pods at once or do a node reboot. After applying the workaround, all pods should go into Running state.

  (BZ#2244353)

- **Backups are failing to transfer data**
  In some situations, backups fail to transfer data, and snapshot PVC is stuck in Pending state.

  (BZ#2248117)

CHAPTER 8. DEPRECATED FEATURES

This section describes the deprecated features introduced in Red Hat OpenShift Data foundation 4.14.

8.1. RED HAT VIRTUALIZATION PLATFORM

Starting Red Hat OpenShift Data Foundation 4.14, OpenShift Data Foundation deployed on Installer-provisioned infrastructure (IPI) deployment of OpenShift on Red Hat Virtualization Platform (RHV) is no longer supported.

For more information, see OpenShift Container Platform 4.14 release notes.

CHAPTER 9. ASYNCHRONOUS ERRATA UPDATES

9.1. RHBA-2024:1579 OPENSHIFT DATA FOUNDATION 4.14.6 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.6 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:1579 advisory.

9.2. RHBA-2024:1043 OPENSHIFT DATA FOUNDATION 4.14.5 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.5 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:1043 advisory.

9.3. RHBA-2024:0315 OPENSHIFT DATA FOUNDATION 4.14.4 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.4 is now available. The bug fixes that are included in the update are listed in the RHBA-2024:0315 advisory.

9.4. RHBA-2023:7869 OPENSHIFT DATA FOUNDATION 4.14.3 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.3 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:7869 advisory.

9.5. RHBA-2023:7776 OPENSHIFT DATA FOUNDATION 4.14.2 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.2 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:7776 advisory.

9.6. RHBA-2023:7696 OPENSHIFT DATA FOUNDATION 4.14.1 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.14.1 is now available. The bug fixes that are included in the update are listed in the RHBA-2023:7696 advisory.