



Red Hat OpenShift Data Foundation 4.10

4.10 Release notes

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

Red Hat OpenShift Data Foundation 4.10 4.10 Release notes

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

Legal Notice

Copyright © 2022 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

The release notes for Red Hat OpenShift Data Foundation 4.10 summarizes all new features and enhancements, notable technical changes, and any known bugs upon general availability.

Table of Contents

CHAPTER 1. OVERVIEW	3
1.1. ABOUT THIS RELEASE	3
CHAPTER 2. NEW FEATURES	4
2.1. MULTICLOUD OBJECT GATEWAY SUPPORT FOR NAMESPACE ON TOP OF FILESYSTEM	4
2.2. KUBERNETES NATIVE AUTHENTICATION METHOD TO AUTOMATICALLY AUTHENTICATE AND RENEW SECURITY TOKEN FOR KMS	4
2.3. MINIMUM DEPLOYMENT GENERAL AVAILABILITY SUPPORT	4
2.4. IBM CLOUD HYPER PROTECT CRYPTO SERVICES KEY MANAGEMENT SYSTEM INTEGRATION	4
2.5. STORAGE CLASS SELECTION FOR STANDALONE MULTICLOUD OBJECT GATEWAY FROM THE USER INTERFACE	4
2.6. SUPPORT FOR EXTERNAL MODE ON IBM POWER AND IBM Z INFRASTRUCTURE	4
CHAPTER 3. ENHANCEMENTS	5
3.1. ADDING OPENSIFT DATA FOUNDATION TAINTS FROM THE USER INTERFACE	5
3.2. SUPPORT FOR AWS GP2 AND GP3 CSI DRIVERS	5
3.3. UTILIZATION CARD UPDATE	5
3.4. ROOK	5
3.5. MULTI-CLOUD OBJECT GATEWAY	5
3.6. CSI DRIVER	6
3.7. MANAGEMENT CONSOLE	6
CHAPTER 4. TECHNOLOGY PREVIEWS	7
4.1. DYNAMICALLY PROVISIONED STORAGE FOR SINGLE NODE OPENSIFT CLUSTERS	7
CHAPTER 5. DEVELOPER PREVIEWS	8
5.1. RADOS GATEWAY CONFIGURATION FOR BUCKET NOTIFICATION	8
5.2. REGIONAL-DR WITH ADVANCED CLUSTER MANAGEMENT	8
5.3. METRO-DR MULTIPLE CLUSTERS WITH ACM	8
CHAPTER 6. DEPRECATED FEATURES	9
6.1. THICK-PROVISIONING OF RBD PERSISTENTVOLUMES	9
CHAPTER 7. BUG FIXES	10
7.1. MULTI-CLOUD OBJECT GATEWAY	10
7.2. ROOK	10
7.3. MANAGEMENT CONSOLE	10
7.4. ODF OPERATOR	11
7.5. CSI DRIVER	11
CHAPTER 8. KNOWN ISSUES	12
8.1. ODF-DR	12
8.2. MANAGEMENT CONSOLE	12
8.3. ROOK	12
8.4. ODF OPERATOR	13
8.5. CEPH	13
CHAPTER 9. ASYNCHRONOUS ERRATA UPDATES	14
9.1. RHBA-2022:6675 OPENSIFT DATA FOUNDATION 4.10.6 BUG FIXES AND SECURITY UPDATES	14
9.2. RHBA-2022:5607 OPENSIFT DATA FOUNDATION 4.10.5 BUG FIXES AND SECURITY UPDATES	14
9.3. RHBA-2022:5196 OPENSIFT DATA FOUNDATION 4.10.4 BUG FIXES AND SECURITY UPDATES	14
9.4. RHBA-2022:5023 OPENSIFT DATA FOUNDATION 4.10.3 BUG FIXES AND SECURITY UPDATES	14
9.5. RHBA-2022:4621 OPENSIFT DATA FOUNDATION 4.10.2 BUG FIXES AND SECURITY UPDATES	14
9.6. RHBA-2022:2182 OPENSIFT DATA FOUNDATION 4.10.1 BUG FIXES AND SECURITY UPDATES	15

CHAPTER 1. OVERVIEW

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

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

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

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

1.1. ABOUT THIS RELEASE

Red Hat OpenShift Data Foundation 4.10 ([RHSA-2022:1361](#) and [RHSA-2022:1372](#)) is now available. New enhancements, features, and known issues that pertain to OpenShift Data Foundation 4.10 are included in this topic.

Red Hat OpenShift Data Foundation 4.10 is supported on the Red Hat OpenShift Container Platform version 4.10. For more information, see [Red Hat OpenShift Data Foundation Supportability and Interoperability Guide](#).

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 4.10.

2.1. MULTICLOUD OBJECT GATEWAY SUPPORT FOR NAMESPACE ON TOP OF FILESYSTEM

Multicloud Object Gateway (MCG) now has the capability to share data between legacy applications and the cloud native applications, enables easier pipelines in the case of artificial intelligence and machine learning. The object storage capabilities of MCG is expanded to allow access to file systems using Amazon Web Services S3 protocol thereby enabling sharing of data for artificial intelligence and machine learning. For more information, see [Sharing legacy application data with cloud native application using S3 protocol](#).

2.2. KUBERNETES NATIVE AUTHENTICATION METHOD TO AUTOMATICALLY AUTHENTICATE AND RENEW SECURITY TOKEN FOR KMS

For cluster-wide encryption, you can use the Kubernetes authentication method when using HashiCorp Vault Key Management Service for automatic renewal of expired tokens and native integration for a more comprehensive encryption solution. This feature is already available for Persistent Volume encryption. For more information, see [Enabling cluster-wide encryption with KMS using the Kubernetes authentication method](#).

2.3. MINIMUM DEPLOYMENT GENERAL AVAILABILITY SUPPORT

OpenShift Data Foundation (ODF) now supports general availability of deploying the ODF with minimum configuration when the standard deployment resource requirement is not met. For more information, see [minimum deployment resource requirements](#) in the Planning Guide.

2.4. IBM CLOUD HYPER PROTECT CRYPTO SERVICES KEY MANAGEMENT SYSTEM INTEGRATION

OpenShift Data Foundation on IBM Cloud platform now supports Hyper Protect Crypto Services (HPCS) Key Management Services (KMS) as the encryption solution in addition to HashiCorp Vault KMS. HPCS is built on FIPS 140-2 Level 4-certified hardware.

2.5. STORAGE CLASS SELECTION FOR STANDALONE MULTICLOUD OBJECT GATEWAY FROM THE USER INTERFACE

In this release, when you deploy a standalone Multicloud Object Gateway using local storage devices, you have an option to select a storage class from the user interface.

2.6. SUPPORT FOR EXTERNAL MODE ON IBM POWER AND IBM Z INFRASTRUCTURE

In external mode, Red Hat OpenShift Data Foundation exposes the Red Hat Ceph Storage services running outside of the OpenShift Container Platform cluster as storage classes. You can set up Red Hat Ceph Storage on an x86 platform-based OCP environment, and then use it for underlying storage purposes with OpenShift Data Foundation on IBM Power and IBM Z infrastructure.

CHAPTER 3. ENHANCEMENTS

This section describes the major enhancements introduced in Red Hat OpenShift Data foundation 4.10.

3.1. ADDING OPENSIFT DATA FOUNDATION TAINTS FROM THE USER INTERFACE

With this update, you can select the OpenShift Data Foundation(ODF) taint nodes option while creating the ODF cluster and even allows you to add the ODF taint post the cluster creation using the user interface. Adding the ODF taint allows only ODF pods to be run on those worker nodes, thus making it dedicated for ODF. For more information, see the [deployment guide](#) and [managing guide](#) for adding taints from the user interface during deployment and post-deployment respectively.

3.2. SUPPORT FOR AWS GP2 AND GP3 CSI DRIVERS

OpenShift Data Foundation now supports the gp2 CSI and gp3 CSI drivers introduced by AWS. These drivers provide improved storage expansion capabilities (gp2 CSI compared to gp2 in-tree) and a reduced monthly price-point (gp3). For more information, see the [Infrastructure requirements](#) section of the Planning Guide.

3.3. UTILIZATION CARD UPDATE

With this update, you can view an improved graph representation under the Block and File dashboard. For internal mode cluster, the graph indicates an area chart for used capacity and recovery, stack chart for I/O operations and throughput, line chart for latency information for the internal mode cluster. For more information, see [Metrics in the Block and File dashboard](#) .

3.4. ROOK

Update quota limit on Object Bucket Claim

Previously, when the usage exceeded the quota limit, then all the operations on the bucket attached with Object Bucket Claim(OBC) became read only. With this update, you can update the quota limit on OBC based on your requirement.

OSDs are safe when multiple jobs are fired

Previously, when multiple jobs removal were fired in parallel then there was a risk of losing data since it would forcefully remove the OSD.

With this update, if you perform multiple jobs removal then it checks whether the OSD is ok-to-stop first and then proceeds. This implementation waits endlessly and retries every minute thereby keeping the OSD safe from losing data.

3.5. MULTI-CLOUD OBJECT GATEWAY

NooBaa services update

With this update, a new flag is added **disable-load-balancer** that replaces the type of service from LoadBalancer to ClusterIP. This allows you to disable the NooBaa service EXTERNAL-IP.

For instructions, see [knowledgebase article](#) about *Disabling Multicloud Object Gateway external service for private clusters*.

3.6. CSI DRIVER

Automatic reclaim space for RADOS Block Devices

RADOS Block Devices(RBD) PersistentVolumes are thin-provisioned when created, meaning little space from the Ceph cluster is consumed. When data is stored on the PersistentVolume, the consumed storage increases automatically. However, after data is deleted, the consumed storage does not reduce, as the RBD PersistentVolume does not return the free space back to the Ceph cluster. In certain scenarios, it is required that the freed up space is returned to the Ceph cluster so that the other workloads can benefit from it.

With this update, the ReclaimSpace feature allows you to enable automatic reclaiming of freed up space from RBD PersistentVolumes with thin-provisioning. You can add an annotation to your PersistentVolume Claim, create a ReclaimSpaceCronJob for recurring space reclaiming, or run a ReclaimSpaceJob for a one-time operation. For more information, see [Reclaiming space on target volumes](#)


3.7. MANAGEMENT CONSOLE

View the Block and File or Object Service subcomponents on the ODF Dashboard

With this update, you can view the information of the ODF subcomponents, Block and File or Object Service, whenever any of it is down on the OpenShift Data Foundation dashboard.

CHAPTER 4. TECHNOLOGY PREVIEWS

This section describes technology preview features introduced in Red Hat OpenShift Data Foundation 4.10 under Technology Preview support limitations.



IMPORTANT

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

4.1. DYNAMICALLY PROVISIONED STORAGE FOR SINGLE NODE OPENSIFT CLUSTERS

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 [Deploying OpenShift Data Foundation on Single Node Radio Access Network](#).

CHAPTER 5. DEVELOPER PREVIEWS

This section describes developer preview features introduced in Red Hat OpenShift Data Foundation 4.10.



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. RADOS GATEWAY CONFIGURATION FOR BUCKET NOTIFICATION

A simplified way with new custom resources to configure the RGW bucket notification is now available. For more information, see [Bucket Notification in Rook](#).

5.2. REGIONAL-DR WITH ADVANCED CLUSTER MANAGEMENT

Regional-DR is composed of Red Hat Advanced Cluster Management for Kubernetes (RHACM) and OpenShift Data Foundation components to provide application and data mobility across OpenShift Container Platform clusters. It is built on Asynchronous data replication and hence could have a potential data loss but provides the protection against a broad set of failures. For more information, see [Planning guide](#) and [Configuring Regional-DR with Advanced Cluster Management](#).

5.3. METRO-DR MULTIPLE CLUSTERS WITH ACM

Metro-DR capability provides volume persistent data and metadata replication across sites that are geographically dispersed. In the public cloud these would be similar to protecting from an Availability Zone failure. Metro-DR ensures business continuity during the unavailability of a data center with no data loss. For more information, see [Planning guide](#) and [Configuring Metro-DR with Advanced Cluster Management](#).

CHAPTER 6. DEPRECATED FEATURES

This section describes the deprecated features introduced in Red Hat OpenShift Data foundation 4.10.

6.1. THICK-PROVISIONING OF RBD PERSISTENTVOLUMES

The feature RADOS Block Devices (RBD) thick provisioning with new storage class capability is now deprecated. Administrators used this feature to configure Storage Quotas for their tenants.

CHAPTER 7. BUG FIXES

This section describes notable bug fixes introduced in Red Hat OpenShift Data Foundation 4.10.

7.1. MULTI-CLOUD OBJECT GATEWAY

Install the Multicloud Object Gateway with a secure transfer

Previously, when the Microsoft Azure resource group was configured with a policy to enforce secure transfer for storage accounts, the installation of Multicloud Object Gateway (MCG) was stuck on the creation of the default backing store. This was because the MCG failed to create a storage account for the default backing store. With this update, a flag is added to allow HTTPS traffic only when you create a storage account. Now, you can install the MCG only in an environment that enforces secure transfer. ([BZ#1970123](#))

Object expiration lifecycle update

Previously, the object lifecycle expiration was set by days and mistakenly counted in minutes. With this update, object expiration lifecycle is counted in days not in minutes. ([BZ#2034661](#))

Noobaa handles the upload request headers as expected

With this update, Noobaa saves the correct content_encoding type that is sent in the upload request headers and returns the proper HEAD/GET operations. ([BZ#2054074](#))

7.2. ROOK

CephObjectStore is reaching to Ready state

Previously, there was an error with CephObjectStore where it would not attain the Ready state if it was deleted and then a new CephObjectStore was created with the same name.

With this update, any new CephObjectStore can be able to reach the Ready state even if a previously deleted CephObjectStore has the same name. ([BZ#1974344](#))

Adding an upgrade flag to grant new permissions

With this update, you can upgrade the **cephCSIKeyrings**, for example, client.csi-cephfs-provision with new permissions caps. To upgrade all the **cephCSIKeyrings** run **python3 /etc/ceph/create-external-cluster-resources.py --upgrade**. The upgrade flag is required when you already have an ODF deployment with RHCS(external Ceph storage system) and now you are either upgrading or adding a new ODF deployment(multi-tenant) to the RHCS cluster. The upgrade flag is not required when you are freshly creating an ODF deployment with RHCS cluster. ([BZ#2044983](#))

7.3. MANAGEMENT CONSOLE

OpenShift Data Foundation user interface available in camel case

Previously, OpenShift Data Foundation user interface used upper case to store the vault Key Management System (KMS) configs in the csi-kms-connection-details config map. However, Ceph Container Storage Interface (CSI) supports the upper case for a user interface at limited places. Ceph CSI recommends using camel cases in most places. As a result, the csi-kms-connection-details config map is mixed with both upper and lower cases which caused confusion. With this update, the user interface is moving to the camel case while supporting the upper case for backward compatibility. ([BZ#2005801](#))

7.4. ODF OPERATOR

Define OverprovisionControl for custom and built-in storage-classes

Previously, it was not possible to define OverProvisionControl user-defined custom storage-classes, because it was rejected as invalid for the entire StorageCluster CRD. This was because the original solution was restricted only to built-in OpenShift Container Storage (OCS) storage-classes.

With this update, you can define OverprovisionControl for both default as well as user-defined storage-classes. ([BZ#2024545](#))

Automated the creation of cephobjectstoreuser for object bucket claim metrics collector

With this update, the cephobjectstoreuser known as **prometheus-user** to collect data from the RGW server is automatically created. ([BZ#1999952](#))

7.5. CSI DRIVER

Revised permissions on the staging path

Previously, while performing Node mount the permissions on the staging path were explicitly set by the Ceph Container Storage Interface (CSI) driver.

With this update, this issue has been fixed, and, in some scenarios this can help to avoid extra overhead of pod startup delay. ([BZ#2024870](#))

CHAPTER 8. KNOWN ISSUES

This section describes known issues in Red Hat OpenShift Data Foundation 4.10.

8.1. ODF-DR

Creating application namespace for managed clusters

Application namespace needs to exist on managed clusters for Disaster Recovery(DR) related pre-deployment actions and hence is pre-created when an application is deployed at the ACM hub cluster. However, if an application is deleted at the ACM 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 ACM hub, these resources need to be deleted post the application deletion. For example, as cluster administrator execute the following command on the ACM hub cluster, "oc delete manifestwork -n <managedCluster namespace> <drPlacementControl name>-<namespace>-ns-mw".

([BZ#2059669](#))

8.2. MANAGEMENT CONSOLE

OpenShift console disables plugin and all its extension when the network connection is lost

When the user is accessing the Data Foundation dashboard for the first time and in between if the network connectivity is lost then the plugin and extensions of OpenShift Container Platform console are also deactivated for that instance. This happens because an error occurs due to network disruption between the browser and the cluster while resolving any of the required modules.

Workaround: Ensure to have stable network connectivity between the browser and the cluster, refresh the page and make sure that everything is working smoothly.

([BZ#2072965](#))

Standalone Multicloud Object Gateway deployment with external Key Management Service fails

The standalone Multicloud Object Gateway(MCG) deployment using an external Key Management Service(KMS) fails due to a crash in the user interface.

Workaround: There is currently no workaround for this issue, and a fix is expected in one of the upcoming releases.

([BZ#2074810](#))

8.3. ROOK

IBM FlashSystem is not supported with ODF 4.10 due to a failure of Rook-Ceph to run OSDs

Rook-Ceph prepares job fails that result in no OSDs running due to the presence of the environment variable starting with "IBM_".

Workaround: Currently, there is no workaround for this issue, and a fix is expected in one of the upcoming releases of Red Hat OpenShift Data Foundation.

[\(BZ#2073920\)](#)

8.4. ODF OPERATOR

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

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

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

[\(BZ#2004027\)](#)

8.5. CEPH

Poor performance of stretch clusters on CephFS

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

[\(BZ#1982116\)](#)

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

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

[\(BZ#2007376\)](#)

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.

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

CHAPTER 9. ASYNCHRONOUS ERRATA UPDATES

9.1. RHBA-2022:6675 OPENSIFT DATA FOUNDATION 4.10.6 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.6 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:6675](#) advisory.

9.2. RHBA-2022:5607 OPENSIFT DATA FOUNDATION 4.10.5 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.5 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:5607](#) advisory.

Documentation updates

- Added a new section on how to access the legacy application data from the **openshift-storage** namespace. In almost all cases, the data you need to access is not in the **openshift-storage** namespace, but in the namespace that the legacy application uses.
In order to access data stored in another namespace, you need to create a Persistent Volume Claim (PVC) in the **openshift-storage** namespace that points to the same CephFS volume that the legacy application uses. For more information, see *Accessing legacy application data from the openshift-storage namespace* in the [Managing hybrid and multicloud resources guide](#).
- Added a new section on how to change resources in the OpenShift Data Foundation components. When you install OpenShift Data Foundation, it comes with pre-defined resources that the OpenShift Data Foundation pods can consume.
In some situations with higher I/O load, it might be required to increase these limits. For more information, see *Changing resources for the OpenShift Data Foundation components* in the [Troubleshooting guide](#).

9.3. RHBA-2022:5196 OPENSIFT DATA FOUNDATION 4.10.4 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.4 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:5196](#) advisory.

9.4. RHBA-2022:5023 OPENSIFT DATA FOUNDATION 4.10.3 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.3 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:5023](#) advisory.

9.5. RHBA-2022:4621 OPENSIFT DATA FOUNDATION 4.10.2 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.2 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:4621](#) advisory.

Documentation updates

Added a section on how to remove the default bucket created by the Multicloud Object Gateway (MCG). The MCG creates a default bucket in the cloud. You need to remove this default bucket. For more information, see *Removing the default bucket created by the Multicloud Object Gateway* in the Red Hat Knowledgebase solution [Uninstalling OpenShift Data Foundation in Internal mode](#).

9.6. RHBA-2022:2182 OPENSIFT DATA FOUNDATION 4.10.1 BUG FIXES AND SECURITY UPDATES

OpenShift Data Foundation release 4.10.1 is now available. The bug fixes that are included in the update are listed in the [RHBA-2022:2182](#) advisory.