OpenShift Data Foundation

Subscription Guide

(2024/08)

An overview of how subscriptions for OpenShift Data Foundation work

Last Updated: 2024-08-19

Legal Notice

Copyright © 2022 Red Hat, Inc.

The text 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

This document provides an overview on how OpenShift Data Foundations subscriptions work.

What is Red Hat OpenShift Data Foundation ? 5
Where can Red Hat OpenShift Data Foundation be deployed ? 5
How can Red Hat OpenShift Data Foundation be deployed ? 6
Internal mode OpenShift Data Foundation 6
Internal mode deployed on OpenShift infrastructure nodes 6
External mode OpenShift Data Foundation 7
External mode OpenShift Data Foundation with multiple OpenShift clusters 8
Red Hat OpenShift Data Foundation Editions 9
Red Hat OpenShift Data Foundation Essentials Edition 9
Red Hat OpenShift Data Foundation Advanced Edition 9
Mixed workload definition 9
Feature comparison Red Hat OpenShift Data Foundation Essentials vs Advanced Edition 10
Red Hat OpenShift Platform Plus 10
Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced 11
How are OpenShift Data Foundation subscriptions counted ? 12
Bare-metal deployments 12
Socket-based subscriptions 12
Example 1 - Internal Mode 12
Example 2 - External mode with 1 OpenShift cluster and 1 OpenShift Data Foundation cluster 12
Example 3 - External mode with 3 OpenShift clusters and 1 OpenShift Data Foundation cluster 13
Virtualized deployments 13
Cores versus vCPUs and Hyper-Threading Technology 13
Example 1 - Internal Mode 13
Example 2 - Internal Mode with infrastructure nodes 14
OpenShift Data Foundation deployments without worker-nodes 14
OpenShift Data Foundation Capacity boundaries 14
Example - OpenShift clusters and 1 External Mode OpenShift Data Foundation cluster with 12PB raw capacity 15
Support Level of Subscriptions 16
Disaster Recovery deployments and subscription coverage 16
Hot DR 16
Warm DR 16
Cold DR 17

Migration and swing upgrades **17**

What is Red Hat OpenShift Data Foundation ?

Red Hat OpenShift Data Foundation (OpenShift Data Foundation) is a software-defined persistent storage solution for containerized applications running on OpenShift and for OpenShift infrastructure components and services which require storing data. OpenShift Data Foundation provides persistent file, block or object storage services. Storage services can be provisioned by administrators, by users or API-driven. As a software-defined storage solution, OpenShift Data Foundation provides the software components to deploy, virtualize, provision, manage and monitor all storage related services. OpenShift Data Foundation is typically deployed as software containers running on physical or virtual servers. These servers must provide physical or virtual disks to be consumed and virtualized by the storage software.

Where can Red Hat OpenShift Data Foundation be deployed ?

Red Hat OpenShift Data Foundation can be deployed on most platforms that Red Hat OpenShift Container Platform can be deployed on, but it can also be deployed as a dedicated standalone storage cluster. There are fundamentally two different deployment flavors:

- Red Hat OpenShift Data Foundation deployed on physical servers
  => Bare-metal nodes
    => Socket-based subscriptions will be required

- Red Hat OpenShift Data Foundation deployed on virtual servers or in the public cloud
  => Virtual nodes
    => Core-based subscriptions will be required

Depending on the deployment platform, different units of measure will be used to count subscriptions for Red Hat OpenShift Data Foundation.

How can Red Hat OpenShift Data Foundation be deployed ?

Red Hat OpenShift Data Foundation can be deployed in two different deployment modes: *Internal mode* and *External mode*

Internal mode OpenShift Data Foundation

Internal mode means that all OpenShift Data Foundation software components are deployed as containers directly on the OpenShift cluster nodes together with other containerized applications. In internal mode, OpenShift Data Foundation is typically deployed on OpenShift worker nodes or infrastructure nodes. These nodes need to provide physical or virtual storage devices and capacity which can be consumed by Red Hat OpenShift Data Foundation. In internal mode the storage services which OpenShift Data Foundation provides can only be consumed by the same OpenShift cluster where OpenShift Data Foundation is deployed.
Internal mode is currently the only deployment model that is supported for OpenShift deployments in the public cloud or on VMware.

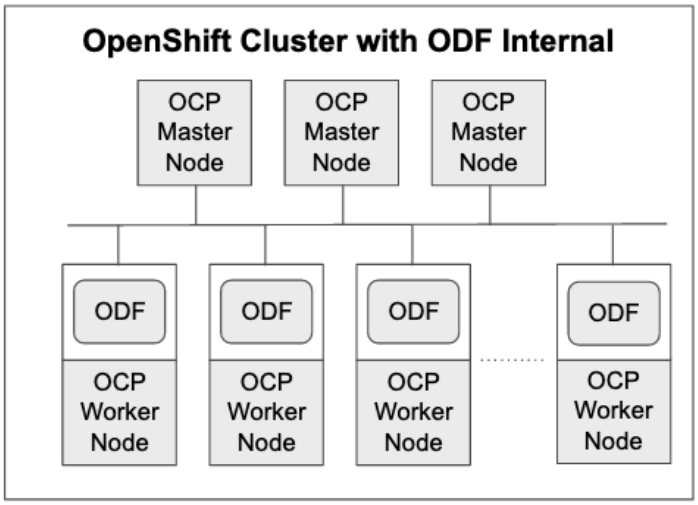

Figure 1: OpenShift with OpenShift Data Foundation in Internal Mode

Internal mode deployed on OpenShift infrastructure nodes

Note: OpenShift Data Foundation can also be deployed on OpenShift infrastructure nodes but for subscriptions only the worker node cores or sockets are counted. By deploying OpenShift Data Foundation on infrastructure nodes, customers can potentially reduce the amount of cores needed on the worker nodes which can help to decrease the subscription cost as worker node cores / sockets is the main measure for counting subscriptions in OpenShift.

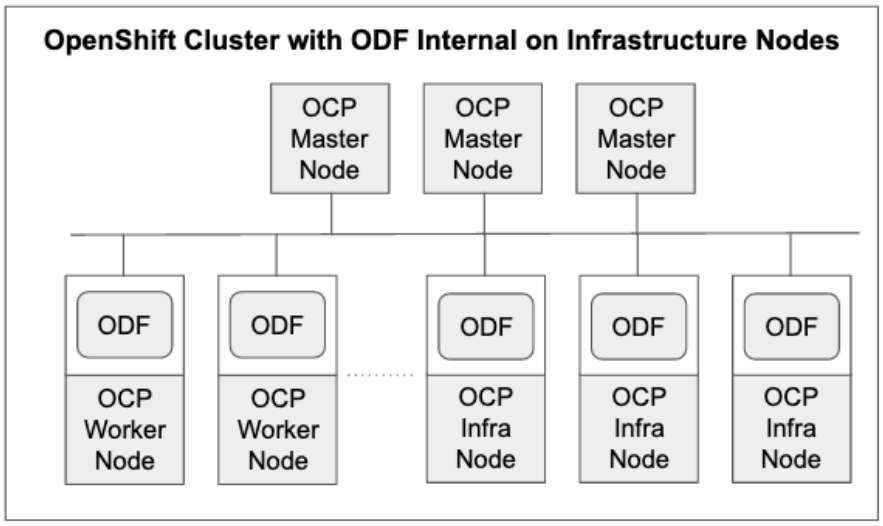

Figure 2: OpenShift with OpenShift Data Foundation in Internal Mode deployed in OpenShift Infrastructure Nodes

External mode OpenShift Data Foundation

External mode means that OpenShift Data Foundation is deployed on a dedicated storage cluster, typically a separate Ceph cluster, which provides its storage services to one or multiple OpenShift clusters. External mode is only available for on-premise and bare-metal deployments. The integration and management of these storage services is provided by the OpenShift Data Foundation Operator which needs to be deployed on each OpenShift cluster which requires persistent storage. Attaching a Red Hat Ceph Storage cluster directly to OpenShift without these integration components is not supported.

Please note: Technically an ODF External Mode deployment consists of a dedicated Red Hat Ceph Storage cluster and the ODF integration bits running on OpenShift. However, here we view this from a subscription perspective only. ODF Advanced (part of OpenShift Platform Plus with ODF Advanced) contains all subscriptions required to deploy a dedicated Ceph storage cluster including RHEL.

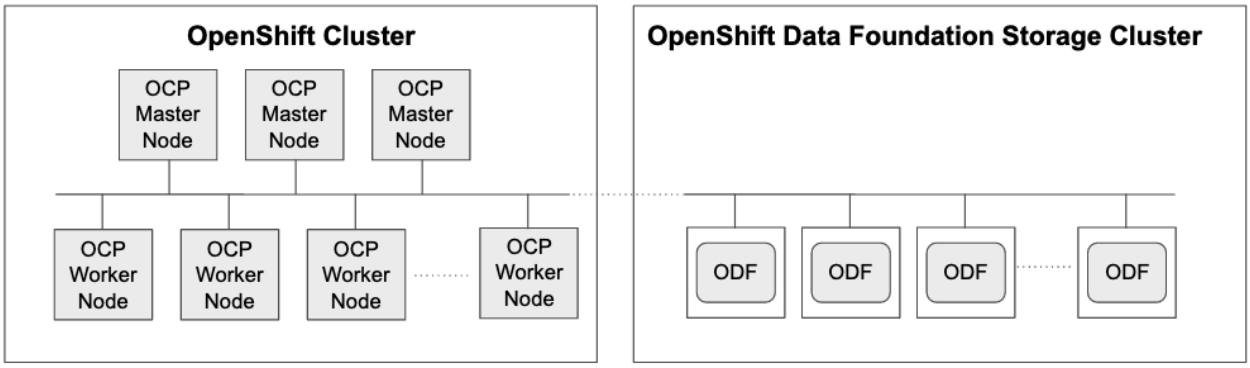

Figure 2: OpenShift with OpenShift Data Foundation in External Mode

External mode OpenShift Data Foundation with multiple OpenShift clusters

In External Mode, the OpenShift Data Foundation storage services can also be consumed by multiple OpenShift clusters and also workloads which are not deployed on OpenShift.

Note: A proper scoping of workload requirements needs to be conducted before multiple OpenShift clusters are attached to one single OpenShift Data Foundation Storage cluster as performance, throughput and latencies are not infinite. We recommend getting Red Hat Global Professional Services or a certified Red Hat OpenShift Partner involved to conduct the scoping.

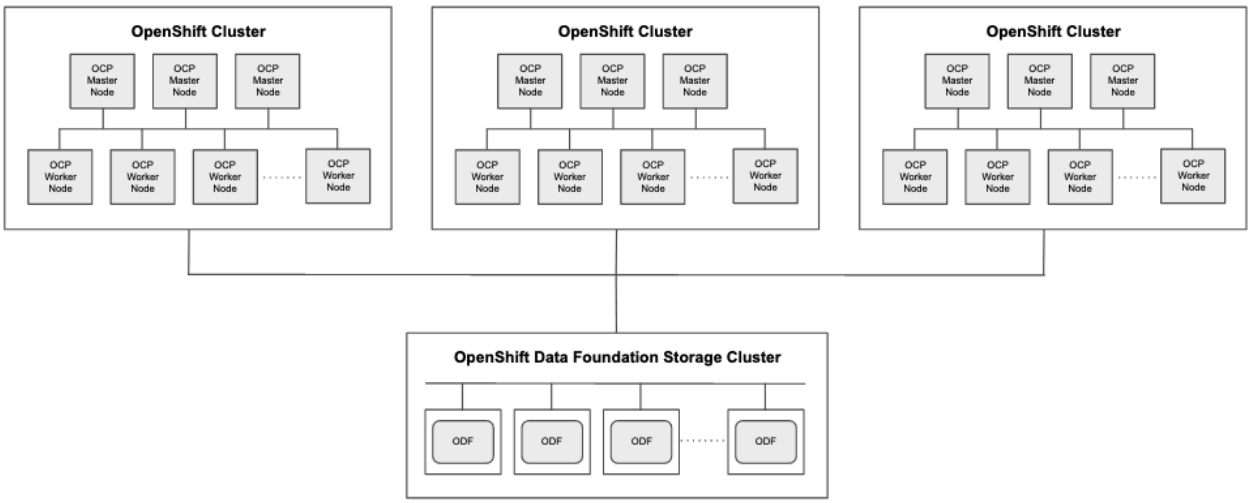

Figure 3: Multiple OpenShift clusters with one OpenShift Data Foundation Storage cluster

Red Hat OpenShift Data Foundation Editions

There are two different flavors of OpenShift Data Foundation available, **OpenShift Data Foundation Essentials** and **OpenShift Data Foundation Advanced.**

Red Hat OpenShift Data Foundation Essentials Edition

provides persistent and highly available storage for applications running on OpenShift and OpenShift internal services like logging, metrics and registry, etc. The recommended way to get Red Hat OpenShift Data Foundation Essentials is through getting *Red Hat OpenShift Platform Plus* which includes OpenShift Data Foundation Essentials.

Red Hat OpenShift Data Foundation Advanced Edition

comes with additional data security, disaster recovery features and more deployment options than OpenShift Data Foundation Essentials. The recommended way to implement OpenShift Data Foundation Advanced Edition is through getting *Red Hat OpenShift Platform Plus with OpenShift Data Foundation Advanced.* This product bundle contains all the necessary technologies required to implement a container platform with enhanced resilience, security and manageability. ODF Advanced Edition also supports **mixed workloads**.

Mixed workload definition

Mixed workloads for ODF Advanced are defined as ODF can provide S3 object storage for applications deployed on OpenShift and outside of OpenShift e.g. applications running on RHEL. The use of Ceph RBD or CephFS directly for external applications or platforms beyond the use as persistent storage for OpenShift through ODF Advanced is not supported and not covered by the subscription. Mixed workloads are only supported for ODF Advanced deployed in external mode.

Feature comparison Red Hat OpenShift Data Foundation Essentials vs Advanced Edition

Product Component	OPP	OPP w ODF Adv
Red Hat OpenShift Container Platform	☑	☑
Red Hat Advanced Cluster Management for Kubernetes	☑	☑
Red Hat Advanced Container Security for Kubernetes	☑	☑
Red Hat Quay	☑	☑
OpenShift Data Foundation		
Block & File	☑	☑
Object	☑	☑
Multi Cloud Gateway	☑	☑
Node and Disk resiliency	☑	☑
Storage Operators based automation	☑	☑
Compression	☑	☑
Deduplication for Multi Cloud Object Gateway	☑	☑
Local Snapshots & Clones	☑	☑
Basic Cluster-wide Encryption	☑	☑
Cross-availability zone HA (3 zones)	☑	☑
Internal Deployment Mode	☑	☑
Stretch Cluster (OCP and ODF stretched over 2 zones + arbiter)	☐	☑
Advanced Encryption with KMS support	☐	☑
Metro DR (sync replication) (1)	☐	☑
Regional DR (async replication) (2)	☐	☑
External Deployment Mode	☐	☑
Multi-Cluster Support	☐	☑
Mixed-Workload Support (3)	☐	☑
Hosted Control Plane Support (4)	☐	☑
Raw Capacity included (5)	256TB	

(1) Requires Stretched Storage Cluster and RH ACM

(2) Requires RH ACM

(3) Mixed workloads definition: applications or clients that require S3 object storage deployed on OpenShift or externally. This does not cover usage of Ceph RBD or CephFS services beyond OpenShift and through ODF Advanced.

(4) Subscription count is based on the hosted workers. OPP+Advanced should be used for this case instead of OPP.

(5) Expandable through Capacity Expansion Packs

Red Hat OpenShift Platform Plus

Red Hat OpenShift Platform Plus contains all components required in a modern Hybrid Cloud Platform for stateful applications. It also includes persistent storage and high availability features and can be deployed virtualized or on bare-metal servers. Red Hat OpenShift Data Foundation Essentials is part of this solution. By definition, Red Hat OpenShift Data Foundation Essentials always has to be deployed in Internal Mode (directly on the OpenShift nodes).

Red Hat OpenShift Platform Plus contains

- Red Hat OpenShift Container Platform
- **Red Hat OpenShift Data Foundation Essentials**
- Red Hat Advanced Cluster Management
- Red Hat Advanced Container Security
- Red Hat Quay

Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced

In addition to all components and features of Red Hat OpenShift Platform Plus, this solution also contains Red Hat OpenShift Data Foundation Advanced which by definition can also be deployed as a separate and dedicated Red Hat OpenShift Data Foundation storage cluster, which is called External Mode.

Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced contains

- Red Hat OpenShift Container Platform
- **Red Hat OpenShift Data Foundation Advanced**
- Red Hat Advanced Cluster Management
- Red Hat Advanced Container Security
- Red Hat Quay

More detailed information on Red Hat OpenShift Platform Plus sizing and subscriptions can be found here:

https://www.redhat.com/en/resources/self-managed-openshift-sizing-subscription-guide

How are OpenShift Data Foundation subscriptions counted ?

In general Red Hat OpenShift Data Foundation subscriptions are counted the same way as Red Hat OpenShift are counted: **every OpenShift worker node needs an active Red Hat subscription**. Subscriptions will differ depending on where OpenShift / OpenShift Data Foundation is being deployed.

Bare-metal deployments

Whenever OpenShift / OpenShift Data Foundation is deployed on bare-metal servers, subscriptions are based on CPU sockets.

Socket-based subscriptions

One OpenShift Data Foundation socket-based subscriptions can be used for physical servers with up to 2 CPU sockets or up to 64 cores (unlimited cores for NFV cases) in total and cannot be used for virtualized OpenShift deployments.

Example 1 - Internal Mode

OpenShift is deployed on 3 master nodes and **6 worker nodes**. OpenShift Data Foundation Is deployed in the same environment running on the 6 worker nodes. Each of these physical nodes have either 1 or 2 CPU sockets. You would require:

> 6 x Red Hat OpenShift Platform Plus (includes OpenShift Data Foundation Essentials), 1-2 socket subscription

Example 2 - External mode with 1 OpenShift cluster and 1 OpenShift Data Foundation cluster

OpenShift is deployed on 3 masters and **2 worker nodes**. OpenShift Data Foundation is deployed in External mode on 10 storage nodes. You would require:

> 2 x Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced, 1-2 socket subscription

Example 3 - External mode with 3 OpenShift clusters and 1 OpenShift Data Foundation cluster

OpenShift is deployed on 3 clusters with 3 masters and **4 worker nodes,** each. OpenShift Data Foundation is deployed in External mode on 10 storage nodes. You would require:

> 3 x 4 = 12 x Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced, 1-2 socket subscription

Virtualized deployments

Whenever OpenShift / OpenShift Data Foundation is deployed on virtualized servers like a Hypervisor platform, a private cloud or public cloud infrastructure, subscriptions will be based on CPU cores.

Cores versus vCPUs and Hyper-Threading Technology

Making a determination about whether or not a particular system uses one or more cores is currently dependent on whether or not that system has hyperthreading available. Note that Hyper-Threading is a feature only available on Intel CPUs but e.g. AMD CPUs have a similar feature called SMT. For simplicity we will refer to Hyper-Threading generically in this document. To determine whether a particular system supports Hyper-Threading, please visit https://access.redhat.com/solutions/7714.

For systems where Hyper-Threading is enabled and where one Hyper-Thread equates to one visible system core, a calculation of cores at a ratio of 2 cores = 4 vCPUs is used.

In other words, a 2-core subscription covers 4 vCPUs in a Hyper-Threaded system. A large VM might have 8 vCPUs, equating to 4 subscription cores. As subscriptions come in units of 2-cores, you would need two 2-core subscriptions to cover these 4 cores or 8 vCPUs.

Where Hyper-Threading is not enabled and where each visible system core correlates directly to an underlying physical core, a calculation of 2 cores = 2 vCPUs is used.

Example 1 - Internal Mode

OpenShift is deployed on 3 master nodes and **6 worker nodes**. OpenShift Data Foundation Is deployed in the same environment running on the 6 worker nodes. Each of the worker nodes is provisioned with **20 cores** / 40 vCPUs. You would require:

6 x 20 / 2 = 60 x Red Hat OpenShift Platform Plus, 2-core subscription

Example 2 - Internal Mode with infrastructure nodes

OpenShift is deployed in the public cloud on 3 master nodes, **2 worker nodes** and 5 infrastructure nodes. OpenShift Data Foundation Is deployed in the same environment running on the 5 infrastructure nodes. Each of the worker nodes is provisioned with **20 cores** / 40 vCPUs. The infrastructure nodes are provisioned with 40 cores / 80 vCPUs. You would require:

2 x 20 / 2 = 20 x Red Hat OpenShift Platform Plus, 2-core subscription

OpenShift Data Foundation deployments without worker-nodes

Hosted Control Plane

OpenShift Data Foundation aligns with OpenShift subscription guidance and will be using 1:1 Ration with OpenShift on the managed hosted clusters, meaning OpenShift Platform Plus + Advanced is required here.

ACM Hub Cluster

If OpenShift Data Foundation is used to provide storage for OpenShift internal infrastructure services such as ACM, ACS or Quay deployed on a dedicated OpenShift cluster without any worker nodes, 20 x OpenShift Data Foundation Advanced subscriptions (20 x 2-core) need to be acquired for that infrastructure. However if ODF is subscribed in the managed clusters already, then ODF can also be deployed for ACM without any additional cost.
This rule also applies to OpenShift Platform Plus deployments where no worker nodes are subscribed but OPP is used to deploy ACM, ACS or Quay with ACM deployment mode of active/backup.

OpenShift Data Foundation Capacity boundaries

Each Red Hat OpenShift Container Platform Cluster with Red Hat OpenShift Data Foundation is entitled up to 256 TB of raw storage capacity. Additional storage capacity can be purchased through Red Hat OpenShift Data Foundation Capacity Expansion Packs. Capacity Expansion Packs cannot be shared across multiple ODF clusters but the storage capacities of each OpenShift cluster (256TB) adds to the total capacity of the central Ceph storage cluster in case of an external mode deployment. Also for mixed workloads like e.g. an external application which requires S3 object storage, the additional capacity required for this application has to be added to the total capacity provided for OpenShift. The following OpenShift Data Foundation Capacity Expansion Packs are available:

256TB, 512TB, 1PB, 2PB, 3PB, 4PB, 5PB and 10PB

Example - OpenShift clusters and 1 External Mode OpenShift Data Foundation cluster with 12PB raw capacity

OpenShift is deployed on 3 bare-metal clusters with 3 masters and **4 worker nodes,** each. OpenShift Data Foundation is deployed in External mode on 10 storage nodes with a total raw capacity of 12PB. You would require:

3 x 4 = 12 x Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced, 1-2 socket subscription
1x Red Hat OpenShift Data Foundation Capacity Expansion Pack, 10PB subscription
1x Red Hat OpenShift Data Foundation Capacity Expansion Pack, 2PB subscription

Support Level of Subscriptions

Red Hat OpenShift Data Foundation is available in two different support levels:

1. Premium Support: Coverage is 24 x 7
2. Standard Support: Coverage is Standard business hours

More detailed information can be found here:
https://access.redhat.com/support/offerings/production/sla
For any serious production deployment it's strongly advised to default to Premium Support.

Disaster Recovery deployments and subscription coverage

Red Hat defines three types of disaster recovery (DR) environments – hot, warm, and cold. Paid OpenShift and OpenShift Data Foundation Advanced subscriptions for the DR-site are needed for hot DR only. For Metro DR or Regional DR solutions, Red Hat OpenShift Data Foundation Advanced and Red Hat Advanced Cluster Management for Kubernetes subscriptions are required for the primary site and also for the DR-site. In case of an OpenShift cluster with Openshift Data Foundation stretched across two sites with data volumes being synchronously replicated across the two sites, it requires an OpenShift Data Foundation Advanced subscription. The best way to meet all technical requirements for DR is to involve Red Hat Consulting and use the Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced solution bundle.

Hot DR

Hot DR systems are defined as fully functional and running concurrently to the production systems. When Data Volumes are actively being replicated either synchronously or asynchronously between clusters they are considered to be Hot DR systems. Hot DR systems are ready to immediately activate their Applications, receive client IO traffic and take over in the event of a disaster within the primary environment.

Warm DR

Warm DR systems are defined as already stocked with hardware representing a reasonable facsimile of that found in the primary site, but containing no customer data. Warm DR systems should not be participating in active data volume replication either synchronously or

asynchronously between clusters. Warm DR recovery requires customer's data having to be restored onto the existing cluster hardware from a last backup located outside of the cluster.

Cold DR

Cold DR systems are defined as having the infrastructure in place, but not the full technology (hardware, software) needed to restore service.

For both warm DR and cold DR, the Red Hat OpenShift subscriptions can be transferred from the primary environment to the DR environment when the disaster occurs to restore service and maintain compliance with Red Hat's subscription terms.

Migration and swing upgrades

Red Hat OpenShift 4 provides in-place upgrades between minor versions. However, for customers who are upgrading from Red Hat OpenShift 3 or need to perform a swing upgrade due to maintenance windows or other considerations, your Red Hat OpenShift subscription will cover both the original and destination infrastructure of a one-way migration until such migration is complete. This also applies to Red Hat OpenShift Data Foundation.
During the migration, Red Hat's subscription management tools will show your environment as being out-of-compliance based on the number of OpenShift or OpenShift Data Foundation subscriptions you purchased. Red Hat allows this for major version upgrades and will not require the purchase of additional subscriptions to get back into compliance during the migration.