# OpenShift Dedicated 4

## Policies and service definition

Policies and service definition for OpenShift Dedicated

Last Updated: 2021-09-28

# OpenShift Dedicated 4 Policies and service definition

Policies and service definition for OpenShift Dedicated

## Legal Notice



## Abstract

Policies and service definition for your OpenShift Dedicated cluster

# Table of Contents

# CHAPTER 1. OPENSHIFT DEDICATED SERVICE DEFINITION

## 1.1. ACCOUNT MANAGEMENT

### 1.1.1. Billing

Each OpenShift Dedicated cluster requires a minimum annual base cluster purchase and there are two billing options available for each cluster: Standard and Customer Cloud Subscription (CCS).

Standard OpenShift Dedicated clusters are deployed in to their own cloud infrastructure accounts, each owned by Red Hat. Red Hat is responsible for this account, and cloud infrastructure costs are paid directly by Red Hat. The customer only pays the Red Hat subscription costs.

In the CCS model, the customer pays the cloud infrastructure provider directly for cloud costs and the cloud infrastructure account is part of a customer's Organization, with specific access granted to Red Hat. The customer will have restricted access to this account, but will be able to view billing and usage information. In this model, the customer pays Red Hat for the CCS subscription and pays the cloud provider for the cloud costs. It is the customer's responsibility to pre-purchase or provide Reserved Instance (RI) compute instances to ensure lower cloud infrastructure costs.

Additional resources can be purchased for an OpenShift Dedicated Cluster, including:

- Additional nodes (can be different types and sizes through the use of machine pools)
- Middleware (JBoss EAP, JBoss Fuse, and so on) - additional pricing based on specific middleware component
- Additional storage in increments of 500 GB (standard only; 100 GB included)
- Additional 12 TiB Network I/O (standard only; 12 TB included)
- Load Balancers for Services are available in bundles of 4; enables non-HTTP/SNI traffic or non-standard ports (standard only)

### 1.1.2. Cluster self-service

Customers can create, scale, and delete their clusters from OpenShift Cluster Manager (OCM), provided that they have pre-purchased the necessary subscriptions.

Actions available in OpenShift Cluster Manager (OCM) must not be directly performed from within the cluster as this might cause adverse affects, including having all actions automatically reverted.

### 1.1.3. Cloud providers

OpenShift Dedicated offers OpenShift Container Platform clusters as a managed service on the following cloud providers:

- Amazon Web Services (AWS)
- Google Cloud Platform (GCP)

### 1.1.4. Compute

Single availability zone clusters require a minimum of 2 worker nodes for Customer Cloud Subscription (CCS) clusters deployed to a single availability zone. A minimum of 4 worker nodes is required for standard clusters. These 4 worker nodes are included in the base subscription.

Multiple availability zone clusters require a minimum of 3 worker nodes for Customer Cloud Subscription (CCS) clusters, 1 deployed to each of 3 availability zones. A minimum of 9 worker nodes are required for standard clusters. These 9 worker nodes are included in the base subscription, and additional nodes must be purchased in multiples of 3 to maintain proper node distribution.

Worker nodes must all be the same type and size within a single OpenShift Dedicated cluster.

**NOTE**

The default machine pool node type and size cannot be changed after the cluster has been created.

Control and infrastructure nodes are also provided by Red Hat. There are at least 3 control planenodes that handle etcd and API-related workloads. There are at least 2 infrastructure nodes that handle metrics, routing, the web console, and other workloads. Control and infrastructure nodes are strictly for Red Hat workloads to operate the service, and customer workloads are not permitted to be deployed on these nodes.

**NOTE**

Approximately 1 vCPU core and 1 GiB of memory are reserved on each worker node and removed from allocatable resources. This is necessary to run processes required by the underlying platform. This includes system daemons such as udev, kubelet, container runtime, and so on, and also accounts for kernel reservations. OpenShift Container Platform core systems such as audit log aggregation, metrics collection, DNS, image registry, SDN, and so on might consume additional allocatable resources to maintain the stability and maintainability of the cluster. The additional resources consumed might vary based on usage.

### 1.1.5. AWS compute types

OpenShift Dedicated offers the following worker node types and sizes on AWS:

**General purpose**

- M5.xlarge (4 vCPU, 16 GiB)
- M5.2xlarge (8 vCPU, 32 GiB)
- M5.4xlarge (16 vCPU, 64 GiB)
- M5.8xlarge (32 vCPU, 128 GiB)
- M5.12xlarge (48 vCPU, 192 GiB)
- M5.16xlarge (64 vCPU, 256 GiB)
- M5.24xlarge (96 vCPU, 384 GiB)

**Memory-optimized**

- R5.xlarge (4 vCPU, 32 GiB)
- R5.2xlarge (8 vCPU, 64 GiB)
- R5.4xlarge (16 vCPU, 128 GiB)
- R5.8xlarge (32 vCPU, 256 GiB)
- R5.12xlarge (48 vCPU, 384 GiB)
- R5.16xlarge (64 vCPU, 512 GiB)
- R5.24xlarge (96 vCPU, 768 GiB)

**Compute-optimized**

- C5.2xlarge (8 vCPU, 16 GiB)
- C5.4xlarge (16 vCPU, 32 GiB)
- C5.9xlarge (36 vCPU, 72 GiB)
- C5.12xlarge (48 vCPU, 96 GiB)
- C5.18xlarge (72 vCPU, 144 GiB)
- C5.24xlarge (96 vCPU, 192 GiB)

### 1.1.6. Google Cloud compute types

OpenShift Dedicated offers the following worker node types and sizes on Google Cloud that are chosen to have a common CPU and memory capacity that are the same as other cloud instance types:

**General purpose**

- custom-4-16384 (4 vCPU, 16 GiB)
- custom-8-32768 (8 vCPU, 32 GiB)
- custom-16-65536 (16 vCPU, 64 GiB)
- custom-32-131072 (32 vCPU, 128 GiB)
- custom-48-196608 (48 vCPU, 192 GiB)
- custom-64-262144 (64 vCPU, 256 GiB)
- custom-96-393216 (96 vCPU, 384 GiB)

**Memory-optimized**

- custom-4-32768-ext (4 vCPU, 32 GiB)
- custom-8-65536-ext (8 vCPU, 64 GiB)
- custom-16-131072-ext (16 vCPU, 128 GiB)

- custom-32-262144 (32 vCPU, 256 GiB)
- custom-48-393216 (48 vCPU, 384 GiB)
- custom-64-524288 (64 vCPU, 512 GiB)
- custom-96-786432 (96 vCPU, 768 GiB)

**Compute-optimized**

- custom-8-16384 (8 vCPU, 16 GiB)
- custom-16-32768 (16 vCPU, 32 GiB)
- custom-36-73728 (36 vCPU, 72 GiB)
- custom-48-98304 (48 vCPU, 96 GiB)
- custom-72-147456 (72 vCPU, 144 GiB)
- custom-96-196608 (96 vCPU, 192 GiB)

### 1.1.7. Regions and availability zones

The following AWS regions are supported by OpenShift Container Platform 4 and are supported for OpenShift Dedicated:

- af-south-1 (Cape Town, AWS opt-in required)
- ap-east-1 (Hong Kong, AWS opt-in required)
- ap-northeast-1 (Tokyo)
- ap-northeast-2 (Seoul)
- ap-south-1 (Mumbai)
- ap-southeast-1 (Singapore)
- ap-southeast-2 (Sydney)
- ca-central-1 (Central Canada)
- eu-central-1 (Frankfurt)
- eu-north-1 (Stockholm)
- eu-south-1 (Milan, AWS opt-in required)
- eu-west-1 (Ireland)
- eu-west-2 (London)
- eu-west-3 (Paris)
- me-south-1 (Bahrain, AWS opt-in required)

- sa-east-1 (São Paulo)
- us-east-1 (N. Virginia)
- us-east-2 (Ohio)
- us-west-1 (N. California)
- us-west-2 (Oregon)

The following Google Cloud regions are currently supported:

- asia-east1, Changhua County, Taiwan
- asia-east2, Hong Kong
- asia-northeast1, Tokyo, Japan
- asia-northeast2, Osaka, Japan
- asia-northeast3, Seoul, Korea
- asia-south1, Mumbai, India
- asia-southeast1, Jurong West, Singapore
- asia-southeast2, Jakarta, Indonesia
- europe-north1, Hamina, Finland
- europe-west1, St. Ghislain, Belgium
- europe-west2, London, England, UK
- europe-west3, Frankfurt, Germany
- europe-west4, Eemshaven, Netherlands
- europe-west6, Zürich, Switzerland
- northamerica-northeast1, Montréal, Québec, Canada
- southamerica-east1, Osasco (São Paulo), Brazil
- us-central1, Council Bluffs, Iowa, USA
- us-east1, Moncks Corner, South Carolina, USA
- us-east4, Ashburn, Northern Virginia, USA
- us-west1, The Dalles, Oregon, USA
- us-west2, Los Angeles, California, USA
- us-west3, Salt Lake City, Utah, USA
- us-west4, Las Vegas, Nevada, USA

Multi-AZ clusters can only be deployed in regions with at least 3 availability zones (see AWS and Google Cloud).

Each new OpenShift Dedicated cluster is installed within a dedicated Virtual Private Cloud (VPC) in a single Region, with the option to deploy into a single Availability Zone (Single-AZ) or across multiple Availability Zones (Multi-AZ). This provides cluster-level network and resource isolation, and enables cloud-provider VPC settings, such as VPN connections and VPC Peering. Persistent volumes are backed by cloud block storage and are specific to the availability zone in which they are provisioned. Persistent volumes do not bind to a volume until the associated pod resource is assigned into a specific availability zone in order to prevent unschedulable pods. Availability zone-specific resources are only usable by resources in the same availability zone.

**WARNING**

The region and the choice of single or multi availability zone cannot be changed once a cluster has been deployed.

### 1.1.8. Service level agreement (SLA)

Any SLAs for the service itself are defined in Appendix 4 of the Red Hat Enterprise Agreement Appendix 4 (Online Subscription Services).

### 1.1.9. Limited support status

You must not remove or replace any native OpenShift Dedicated components or any other component installed and managed by Red Hat. If using cluster administration rights, Red Hat is not responsible for any actions taken by you or any of your authorized users, including actions that might affect infrastructure services, service availability, and data loss.

If any actions that affect infrastructure services, service availability, or data loss are detected, Red Hat will notify the customer of such and request either that the action be reverted or to create a support case to work with Red Hat to remedy any issues.

### 1.1.10. Support

OpenShift Dedicated includes Red Hat Premium Support, which can be accessed by using the Red Hat Customer Portal.

See the Scope of Coverage Page for more details on what is covered with included support for OpenShift Dedicated.

See OpenShift Dedicated SLAs for support response times.

## 1.2. LOGGING

OpenShift Dedicated provides optional integrated log forwarding to Amazon CloudWatch.

### 1.2.1. Cluster audit logging

Cluster audit logs are available through Amazon CloudWatch, if the integration is enabled. If the

integration is not enabled, you can request the audit logs by opening a support case. Audit log requests must specify a date and time range not to exceed 21 days. When requesting audit logs, customers should be aware that audit logs are many GB per day in size.

### 1.2.2. Application logging

Application logs sent to **STDOUT** are collected by Fluentd and forwarded to Amazon CloudWatch through the cluster logging stack, if it is installed.

## 1.3. MONITORING

### 1.3.1. Cluster metrics

OpenShift Dedicated clusters come with an integrated Prometheus/Grafana stack for cluster monitoring including CPU, memory, and network-based metrics. This is accessible through the web console and can also be used to view cluster-level status and capacity/usage through a Grafana dashboard. These metrics also allow for horizontal pod autoscaling based on CPU or memory metrics provided by an OpenShift Dedicated user.

### 1.3.2. Cluster status notification

Red Hat communicates the health and status of OpenShift Dedicated clusters through a combination of a cluster dashboard available in OpenShift Cluster Manager (OCM), and email notifications sent to the email address of the contact that originally deployed the cluster.

## 1.4. NETWORKING

### 1.4.1. Custom domains for applications

To use a custom hostname for a route, you must update your DNS provider by creating a canonical name (CNAME) record. Your CNAME record should map the OpenShift canonical router hostname to your custom domain. The OpenShift canonical router hostname is shown on the **Route Details** page after a Route is created. Alternatively, a wildcard CNAME record can be created once to route all subdomains for a given hostname to the cluster's router.

### 1.4.2. Custom domains for cluster services

Custom domains and subdomains are not available for the platform service routes, for example, the API or web console routes, or for the default application routes.

### 1.4.3. Domain validated certificates

OpenShift Dedicated includes TLS security certificates needed for both internal and external services on the cluster. For external routes, there are two, separate TLS wildcard certificates that are provided and installed on each cluster, one for the web console and route default hostnames and the second for the API endpoint. *Let's Encrypt* is the certificate authority used for certificates. Routes within the cluster, for example, the internal API endpoint, use TLS certificates signed by the cluster's built-in certificate authority and require the CA bundle available in every pod for trusting the TLS certificate.

### 1.4.4. Custom certificate authorities for builds

OpenShift Dedicated supports the use of custom certificate authorities to be trusted by builds when pulling images from an image registry.

### 1.4.5. Load balancers

OpenShift Dedicated uses up to 5 different load balancers:

- Internal control plane load balancer that is internal to the cluster and used to balance traffic for internal cluster communications.
- External control plane load balancer that is used for accessing the OpenShift Container Platform and Kubernetes APIs. This load balancer can be disabled in OpenShift Cluster Manager (OCM). If this load balancer is disabled, Red Hat reconfigures the API DNS to point to the internal control load balancer.
- External control plane load balancer for Red Hat that is reserved for cluster management by Red Hat. Access is strictly controlled, and communication is only possible from allowlisted bastion hosts.
- Default router/ingress load balancer that is the default application load balancer, denoted by **apps** in the URL. The default load balancer can be configured in OpenShift Cluster Manager (OCM) to be either publicly accessible over the internet, or only privately accessible over a pre-existing private connection. All application routes on the cluster are exposed on this default router load balancer, including cluster services such as the logging UI, metrics API, and registry.
- Optional: Secondary router/ingress load balancer that is a secondary application load balancer, denoted by **apps2** in the URL. The secondary load balancer can be configured in OpenShift Cluster Manager (OCM) to be either publicly accessible over the internet, or only privately accessible over a pre-existing private connection. If a 'Label match' is configured for this router load balancer, then only application routes matching this label will be exposed on this router load balancer, otherwise all application routes are also exposed on this router load balancer.
- Optional: Load balancers for services that can can be mapped to a service running on OpenShift Dedicated to enable advanced ingress features, such as non-HTTP/SNI traffic or the use of non-standard ports. These can be purchased in groups of 4 for standard clusters, or they can be provisioned without charge in Customer Cloud Subscription (CCS) clusters; however, each AWS account has a quota that limits the number of Classic Load Balancers that can be used within each cluster.

### 1.4.6. Network usage

For standard OpenShift Dedicated clusters, network usage is measured based on data transfer between inbound, VPC peering, VPN, and AZ traffic. On a standard OpenShift Dedicated base cluster, 12 TB of network I/O is provided. Additional network I/O can be purchased in 12 TB increments. For CCS OpenShift Dedicated clusters, network usage is not monitored, and is billed directly by the cloud provider.

### 1.4.7. Cluster ingress

Project administrators can add route annotations for many different purposes, including ingress control through IP allowlisting.

Ingress policies can also be changed by using **NetworkPolicy** objects, which leverage the **ovs-networkpolicy** plugin. This allows for full control over the ingress network policy down to the pod level, including between pods on the same cluster and even in the same namespace.

All cluster ingress traffic goes through the defined load balancers. Direct access to all nodes is blocked by cloud configuration.

### 1.4.8. Cluster egress

Pod egress traffic control through **EgressNetworkPolicy** objects can be used to prevent or limit outbound traffic in OpenShift Dedicated.

Public outbound traffic from the control plane and infrastructure nodes is required and necessary to maintain cluster image security and cluster monitoring. This requires the **0.0.0.0/0** route to belong only to the internet gateway; it is not possible to route this range over private connections.

OpenShift Dedicated clusters use NAT Gateways to present a public, static IP for any public outbound traffic leaving the cluster. Each subnet a cluster is deployed into receives a distinct NAT Gateway. For clusters deployed on AWS with multiple availability zones, up to 3 unique static IP addresses can exist for cluster egress traffic. For clusters deployed on Google Cloud, regardless of availability zone topology, there will by 1 static IP address for worker node egress traffic. Any traffic that remains inside the cluster or does not go out to the public internet will not pass through the NAT Gateway and will have a source IP address belonging to the node that the traffic originated from. Node IP addresses are dynamic, and therefore a customer should not rely on allowlisting individual IP address when accessing private resources.

Customers can determine their public static IP addresses by running a pod on the cluster and then querying an external service. For example:

```
$ oc run ip-lookup --image=busybox -i -t --restart=Never --rm -- /bin/sh -c "/bin/nslookup -type=a
myip.opendns.com resolver1.opendns.com | grep -E 'Address: [0-9.]+'"
```

### 1.4.9. Cloud network configuration

OpenShift Dedicated allows for the configuration of a private network connection through several cloud provider managed technologies:

- VPN connections
- AWS VPC peering
- AWS Transit Gateway
- AWS Direct Connect
- Google Cloud VPC Network peering
- Google Cloud Classic VPN
- Google Cloud HA VPN

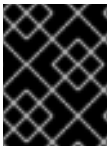

**IMPORTANT**

Red Hat SREs do not monitor private network connections. Monitoring these connections is the responsibility of the customer.

### 1.4.10. DNS forwarding

For OpenShift Dedicated clusters that have a private cloud network configuration, a customer can specify internal DNS servers available on that private connection that should be queried for explicitly provided domains.

## 1.5. STORAGE

### 1.5.1. Encrypted-at-rest OS/node storage

Control plane nodes use encrypted-at-rest-EBS storage.

### 1.5.2. Encrypted-at-rest PV

EBS volumes used for persistent volumes (PVs) are encrypted-at-rest by default.

### 1.5.3. Block storage (RWO)

Persistent volumes (PVs) are backed by AWS EBS and Google Cloud persistent disk block storage, which uses the ReadWriteOnce (RWO) access mode. On a standard OpenShift Dedicated base cluster, 100 GB of block storage is provided for PVs, which is dynamically provisioned and recycled based on application requests. Additional persistent storage can be purchased in 500 GB increments.

PVs can only be attached to a single node at a time and are specific to the availability zone in which they were provisioned, but they can be attached to any node in the availability zone.

Each cloud provider has its own limits for how many PVs can be attached to a single node. See AWS instance type limits or Google Cloud Platform custom machine types for details.

### 1.5.4. Shared storage (RWX)

The AWS CSI Driver can be used to provide RWX support for OpenShift Dedicated on AWS. A community Operator is provided to simplify setup.

## 1.6. PLATFORM

### 1.6.1. Cluster backup policy

**IMPORTANT**

It is critical that customers have a backup plan for their applications and application data.

Application and application data backups are not a part of the OpenShift Dedicated service. All Kubernetes objects in each OpenShift Dedicated cluster are backed up to facilitate a prompt recovery in the unlikely event that a cluster becomes irreparably inoperable.

The backups are stored in a secure object storage (Multi-AZ) bucket in the same account as the cluster. Node root volumes are not backed up because Red Hat Enterprise Linux CoreOS is fully managed by the OpenShift Container Platform cluster and no stateful data should be stored on the root volume of a node.

The following table shows the frequency of backups:

| Component | Snapshot Frequency | Retention | Notes |
|---|---|---|---|
| Full object store backup | Daily at 0100 UTC | 7 days | This is a full backup of all Kubernetes objects. No persistent volumes (PVs) are backed up in this backup schedule. |
| Full object store backup | Weekly on Mondays at 0200 UTC | 30 days | This is a full backup of all Kubernetes objects. No PVs are backed up in this backup schedule. |
| Full object store backup | Hourly at 17 minutes past the hour | 24 hours | This is a full backup of all Kubernetes objects. No PVs are backed up in this backup schedule. |

### 1.6.2. Autoscaling

Node autoscaling is not available on OpenShift Dedicated at this time.

### 1.6.3. Daemon sets

Customers may create and run DaemonSets on OpenShift Dedicated. In order to restrict DaemonSets to only running on worker nodes, use the following nodeSelector:

```
...
spec:
  nodeSelector:
    role: worker
...
```

### 1.6.4. Multiple availability zone

In a multiple availability zone cluster, control nodes are distributed across availability zones and at least three worker nodes are required in each availability zone.

### 1.6.5. Node labels

Custom node labels are created by Red Hat during node creation and cannot be changed on OpenShift Dedicated clusters at this time.

### 1.6.6. OpenShift version

OpenShift Dedicated is run as a service and is kept up to date with the latest OpenShift Container Platform version.

### 1.6.7. Upgrades

Refer to OpenShift Dedicated Life Cycle for more information on the upgrade policy and procedures.

### 1.6.8. Windows containers

Windows containers are not available on OpenShift Dedicated at this time.

### 1.6.9. Container engine

OpenShift Dedicated runs on OpenShift 4 and uses CRI-O as the only available container engine.

### 1.6.10. Operating system

OpenShift Dedicated runs on OpenShift 4 and uses Red Hat Enterprise Linux CoreOS as the operating system for all control plane and worker nodes.

### 1.6.11. Kubernetes Operator support

All Operators listed in the OperatorHub marketplace should be available for installation. Operators installed from OperatorHub, including Red Hat Operators, are not SRE managed as part of the OpenShift Dedicated service. Refer to the Red Hat Customer Portal for more information on the supportability of a given Operator.

## 1.7. SECURITY

### 1.7.1. Authentication provider

Authentication for the cluster is configured as part of the OpenShift Cluster Manager (OCM) cluster creation process. OpenShift is not an identity provider, and all access to the cluster must be managed by the customer as part of their integrated solution. Provisioning multiple identity providers provisioned at the same time is supported. The following identity providers are supported:

- GitHub or GitHub Enterprise OAuth
- GitLab OAuth
- Google OAuth
- LDAP
- OpenID connect

### 1.7.2. Privileged containers

Privileged containers are not available by default on OpenShift Dedicated. The **anyuid** and **nonroot** Security Context Constraints are available for members of the **dedicated-admins** group, and should address many use cases. Privileged containers are only available for **cluster-admin** users.

### 1.7.3. Customer administrator user

In addition to normal users, OpenShift Dedicated provides access to an OpenShift Dedicated-specific group called **dedicated-admin**. Any users on the cluster that are members of the **dedicated-admin** group:

- Have administrator access to all customer-created projects on the cluster.
- Can manage resource quotas and limits on the cluster.

- Can add and manage **NetworkPolicy** objects.
- Are able to view information about specific nodes and PVs in the cluster, including scheduler information.
- Can access the reserved **dedicated-admin** project on the cluster, which allows for the creation of service accounts with elevated privileges and also gives the ability to update default limits and quotas for projects on the cluster.

### 1.7.4. Cluster administration role

As an administrator of OpenShift Dedicated with Customer Cloud Subscriptions (CCS), you have access to the **cluster-admin** role. While logged in to an account with the **cluster-admin** role, users have mostly unrestricted access to control and configure the cluster. There are some configurations that are blocked with webhooks to prevent destablizing the cluster, or because they are managed in OpenShift Cluster Manager (OCM) and any in-cluster changes would be overwritten.

### 1.7.5. Project self-service

All users, by default, have the ability to create, update, and delete their projects. This can be restricted if a member of the **dedicated-admin** group removes the self-provisioner role from authenticated users:

```
$ oc adm policy remove-cluster-role-from-group self-provisioner system:authenticated:oauth
```

Restrictions can be reverted by applying:

```
$ oc adm policy add-cluster-role-to-group self-provisioner system:authenticated:oauth
```

### 1.7.6. Regulatory compliance

See OpenShift Dedicated Process and Security Overview for the latest compliance information.

### 1.7.7. Network security

With OpenShift Dedicated on AWS, AWS provides a standard DDoS protection on all Load Balancers, called AWS Shield. This provides 95% protection against most commonly used level 3 and 4 attacks on all the public facing Load Balancers used for OpenShift Dedicated. A 10-second timeout is added for HTTP requests coming to the haproxy router to receive a response or the connection is closed to provide additional protection.

# CHAPTER 2. RESPONSIBILITY ASSIGNMENT MATRIX

Understanding the Red Hat, cloud provider, and customer responsibilities for the OpenShift Dedicated managed service.

## 2.1. OVERVIEW OF RESPONSIBILITIES FOR OPENSHIFT DEDICATED

While Red Hat manages the OpenShift Dedicated service, the customer shares responsibility with respect to certain aspects. The OpenShift Dedicated services are accessed remotely, hosted on public cloud resources, created in either Red Hat or customer-owned cloud service provider accounts, and have underlying platform and data security that is owned by Red Hat.

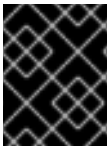

### IMPORTANT

If the **cluster-admin** role is enabled on a cluster, see the responsibilities and exclusion notes in the Red Hat Enterprise Agreement Appendix 4 (Online Subscription Services) .

| Resource | Incident and operations management | Change management | Identity and access management | Security and regulation compliance | Disaster recovery |
|---|---|---|---|---|---|
| Customer data | Customer | Customer | Customer | Customer | Customer |
| Customer applications | Customer | Customer | Customer | Customer | Customer |
| Developer services | Customer | Customer | Customer | Customer | Customer |
| Platform monitoring | Red Hat | Red Hat | Red Hat | Red Hat | Red Hat |
| Logging | Red Hat | Shared | Shared | Shared | Red Hat |
| Application networking | Shared | Shared | Shared | Red Hat | Red Hat |
| Cluster networking | Red Hat | Shared | Shared | Red Hat | Red Hat |
| Virtual networking | Shared | Shared | Shared | Shared | Shared |
| Master and infrastructure nodes | Red Hat | Red Hat | Red Hat | Red Hat | Red Hat |
| Worker nodes | Red Hat | Red Hat | Red Hat | Red Hat | Red Hat |
| Cluster version | Red Hat | Shared | Red Hat | Red Hat | Red Hat |
| Capacity management | Red Hat | Shared | Red Hat | Red Hat | Red Hat |

| Resource | Incident and operations management | Change management | Identity and access management | Security and regulation compliance | Disaster recovery |
|---|---|---|---|---|---|
| Virtual storage | Red Hat and cloud provider | Red Hat and cloud provider | Red Hat and cloud provider | Red Hat and cloud provider | Red Hat and cloud provider |
| Physical infrastructure and security | Cloud provider | Cloud provider | Cloud provider | Cloud provider | Cloud provider |

## 2.2. SHARED RESPONSIBILITY MATRIX

The customer and Red Hat share responsibility for the monitoring and maintenance of an OpenShift Dedicated cluster. This documentation illustrates the delineation of responsibilities by area and task.

### 2.2.1. Incident and operations management

The customer is responsible for incident and operations management of customer application data and any custom networking the customer might have configured for the cluster network or virtual network.

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Application networking | Monitor cloud load balancers and native OpenShift router service, and respond to alerts. | • Monitor health of service load balancer endpoints<br>• Monitor health of application routes, and the endpoints behind them.<br>• Report outages to Red Hat. |
| Virtual networking | Monitor cloud load balancers, subnets, and public cloud components necessary for default platform networking, and respond to alerts. | Monitor network traffic that is optionally configured through VPC to VPC connection, VPN connection, or Direct connection for potential issues or security threats. |

### 2.2.2. Change management

Red Hat is responsible for enabling changes to the cluster infrastructure and services that the customer will control, as well as maintaining versions for the master nodes, infrastructure nodes and services, and worker nodes. The customer is responsible for initiating infrastructure change requests and installing and maintaining optional services and networking configurations on the cluster, as well as all changes to customer data and customer applications.

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Logging | • Centrally aggregate and monitor platform audit logs.<br>• Provide and maintain a logging operator to enable the customer to deploy a logging stack for default application logging.<br>• Provide audit logs upon customer request. | • Install the optional default application logging operator on the cluster.<br>• Install, configure, and maintain any optional app logging solutions, such as logging sidecar containers or third-party logging applications.<br>• Tune size and frequency of application logs being produced by customer applications if they are affecting the stability of the logging stack or the cluster.<br>• Request platform audit logs through a support case for researching specific incidents. |
| Application networking | • Set up public cloud load balancers. Provide the ability to set up private load balancers and up to one additional load balancer when required.<br>• Set up native OpenShift router service. Provide the ability to set the router as private and add up to one additional router shard.<br>• Install, configure, and maintain OpenShift SDN components for default internal pod traffic.<br>• Provide the ability for the customer to manage **NetworkPolicy** and **EgressNetworkPolicy** (firewall) objects. | • Configure non-default pod network permissions for project and pod networks, pod ingress, and pod egress using **NetworkPolicy** objects.<br>• Use OpenShift Cluster Manager to request a private load balancer for default application routes.<br>• Use OpenShift Cluster Manager to configure up to one additional public or private router shard and corresponding load balancer.<br>• Request and configure any additional service load balancers for specific services.<br>• Configure any necessary DNS forwarding rules. |

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Cluster networking | • Set up cluster management components, such as public or private service endpoints and necessary integration with virtual networking components.<br>• Set up internal networking components required for internal cluster communication between worker, infrastructure, and master nodes. | • Provide optional non-default IP address ranges for machine CIDR, service CIDR, and pod CIDR if needed through OpenShift Cluster Manager when the cluster is provisioned.<br>• Request that the API service endpoint be made public or private on cluster creation or after cluster creation through OpenShift Cluster Manager. |
| Virtual networking | • Set up and configure virtual networking components required to provision the cluster, including virtual private cloud, subnets, load balancers, internet gateways, NAT gateways, etc.<br>• Provide the ability for the customer to manage VPN connectivity with on-premises resources, VPC to VPC connectivity, and Direct connectivity as required through OpenShift Cluster Manager.<br>• Enable customers to create and deploy public cloud load balancers for use with service load balancers. | • Set up and maintain optional public cloud networking components, such as VPC to VPC connection, VPN connection, or Direct connection.<br>• Request and configure any additional service load balancers for specific services. |
| Cluster version | • Communicate schedule and status of upgrades for minor and maintenance versions.<br>• Publish changelogs and release notes for minor and maintenance upgrades. | • Work with Red Hat to establish maintenance start times for upgrades.<br>• Test customer applications on minor and maintenance versions to ensure compatibility. |

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Capacity management | • Monitor utilization of control plane (master nodes and infrastructure nodes).<br>• Scale or resize control plane nodes to maintain quality of service.<br>• Monitor utilization of customer resources including Network, Storage and Compute capacity. Where autoscaling features are not enabled alert customer for any changes required to cluster resources (for example, new compute nodes to scale, additional storage, etc). | • Use the provided OpenShift Cluster Manager controls to add or remove additional worker nodes as required.<br>• Respond to Red Hat notifications regarding cluster resource requirements. |

### 2.2.3. Identity and access management

The Identity and Access Management matrix includes responsibilities for managing authorized access to clusters, applications, and infrastructure resources. This includes tasks such as providing access control mechanisms, authentication, authorization, and managing access to resources.

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Logging | • Adhere to an industry standards-based tiered internal access process for platform audit logs.<br>• Provide native OpenShift RBAC capabilities. | • Configure OpenShift RBAC to control access to projects and by extension a project's application logs.<br>• For third-party or custom application logging solutions, the customer is responsible for access management. |
| Application networking | Provide native OpenShift RBAC and **dedicated-admin** capabilities. | • Configure OpenShift dedicated-admins and RBAC to control access to route configuration as required.<br>• Manage Org Admins for Red Hat organization to grant access to OpenShift Cluster Manager. OCM is used to configure router options and provide service load balancer quota. |

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Cluster networking | • Provide customer access controls through OpenShift Cluster Manager.<br>• Provide native OpenShift RBAC and **dedicated-admin** capabilities. | • Manage Red Hat organization membership of Red Hat accounts.<br>• Manage Org Admins for Red Hat organization to grant access to OpenShift Cluster Manager.<br>• Configure OpenShift dedicated-admins and RBAC to control access to route configuration as required. |
| Virtual networking | Provide customer access controls through OpenShift Cluster Manager. | Manage optional user access to public cloud components through OpenShift Cluster Manager. |

### 2.2.4. Security and regulation compliance

The following are the responsibilities and controls related to compliance:

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Logging | Send cluster audit logs to a Red Hat SIEM to analyze for security events. Retain audit logs for a defined period of time to support forensic analysis. | Analyze application logs for security events. Send application logs to an external endpoint through logging sidecar containers or third-party logging applications if longer retention is required than is offered by the default logging stack. |
| Virtual networking | • Monitor virtual networking components for potential issues and security threats.<br>• Leverage additional public cloud provider tools for additional monitoring and protection. | • Monitor optionally-configured virtual networking components for potential issues and security threats.<br>• Configure any necessary firewall rules or data center protections as required. |

### 2.2.5. Disaster recovery

Disaster recovery includes data and configuration backup, replicating data and configuration to the disaster recovery environment, and failover on disaster events.

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Virtual networking | Restore or recreate affected virtual network components that are necessary for the platform to function. | • Configure virtual networking connections with more than one tunnel where possible for protection against outages as recommended by the public cloud provider.<br>• Maintain failover DNS and load balancing if using a global load balancer with multiple clusters. |

## 2.3. CUSTOMER RESPONSIBILITIES FOR DATA AND APPLICATIONS

The customer is responsible for the applications, workloads, and data that they deploy to OpenShift Dedicated. However, Red Hat provides various tools to help the customer manage data and applications on the platform.

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Customer data | • Maintain platform-level standards for data encryption.<br>• Provide OpenShift components to help manage application data, such as secrets.<br>• Enable integration with third-party data services (such as AWS RDS or Google Cloud SQL) to store and manage data outside of the cluster and/or cloud provider. | Maintain responsibility for all customer data stored on the platform and how customer applications consume and expose this data. |

| Resource | Red Hat responsibilities | Customer responsibilities |
|---|---|---|
| Customer applications | • Provision clusters with OpenShift components installed so that customers can access the OpenShift and Kubernetes APIs to deploy and manage containerized applications.<br>• Create clusters with image pull secrets so that customer deployments can pull images from the Red Hat Container Catalog registry.<br>• Provide access to OpenShift APIs that a customer can use to set up Operators to add community, third-party, and Red Hat services to the cluster.<br>• Provide storage classes and plug-ins to support persistent volumes for use with customer applications. | • Maintain responsibility for customer and third-party applications, data, and their complete lifecycle.<br>• If a customer adds Red Hat, community, third-party, their own, or other services to the cluster by using Operators or external images, the customer is responsible for these services and for working with the appropriate provider (including Red Hat) to troubleshoot any issues.<br>• Use the provided tools and features to configure and deploy; keep up-to-date; set up resource requests and limits; size the cluster to have enough resources to run apps; set up permissions; integrate with other services; manage any image streams or templates that the customer deploys; externally serve; save, back up, and restore data; and otherwise manage their highly available and resilient workloads.<br>• Maintain responsibility for monitoring the applications run on OpenShift Dedicated; including installing and operating software to gather metrics and create alerts. |
| Developer services (CodeReady) | Make CodeReady Workspaces available as an add-on through OpenShift Cluster Manager (OCM). | Install, secure, and operate CodeReady Workspaces and the Developer CLI. |

# CHAPTER 3. UNDERSTANDING PROCESS AND SECURITY FOR OPENSHIFT DEDICATED

## 3.1. INCIDENT AND OPERATIONS MANAGEMENT

This documentation details the Red Hat responsibilities for the OpenShift Dedicated managed service.

### 3.1.1. Platform monitoring

A Red Hat Site Reliability Engineer (SRE) maintains a centralized monitoring and alerting system for all OpenShift Dedicated cluster components, SRE services, and underlying cloud provider accounts. Platform audit logs are securely forwarded to a centralized SIEM (Security Information and Event Monitoring) system, where they might trigger configured alerts to the SRE team and are also subject to manual review. Audit logs are retained in the SIEM for one year. Audit logs for a given cluster are not deleted at the time the cluster is deleted.

### 3.1.2. Incident management

An incident is an event that results in a degradation or outage of one or more Red Hat services. An incident can be raised by a customer or Customer Experience and Engagement (CEE) member through a support case, directly by the centralized monitoring and alerting system, or directly by a member of the SRE team.

Depending on the impact on the service and customer, the incident is categorized in terms of severity.

The general workflow of how a new incident is managed by Red Hat:

1. An SRE first responder is alerted to a new incident, and begins an initial investigation.
2. After the initial investigation, the incident is assigned an incident lead, who coordinates the recovery efforts.
3. The incident lead manages all communication and coordination around recovery, including any relevant notifications or support case updates.
4. The incident is recovered.
5. The incident is documented and a root cause analysis is performed within 3 business days of the incident.
6. A root cause analysis (RCA) draft document is shared with the customer within 7 business days of the incident.

### 3.1.3. Notifications

Platform notifications are configured using email. Any customer notification is also sent to the corresponding Red Hat account team and if applicable, the Red Hat Technical Account Manager.

The following activities can trigger notifications:

- Platform incident
- Performance degradation

- Cluster capacity warnings
- Critical vulnerabilities and resolution
- Upgrade scheduling

### 3.1.4. Backup and recovery

All OpenShift Dedicated clusters are backed up using cloud provider snapshots. Notably, this does not include customer data stored on persistent volumes. All snapshots are taken using the appropriate cloud provider snapshot APIs and are uploaded to a secure object storage bucket (S3 in AWS, and GCS in Google Cloud) in the same account as the cluster.

| Component | Snapshot frequency | Retention | Notes |
|---|---|---|---|
| Full object store backup, all SRE-managed cluster persistent volumes (PVs) | Daily | 7 days | This is a full backup of all Kubernetes objects like etcd, as well as all SRE-managed PVs in the cluster. |
| | Weekly | 30 days | |
| Full object store backup | Hourly | 24 hour | This is a full backup of all Kubernetes objects like etcd. No PVs are backed up in this backup schedule. |
| Node root volume | Never | N/A | Nodes are considered to be short-term. Nothing critical should be stored on a node's root volume. |

- Red Hat SRE rehearses recovery processes quarterly.
- Red Hat does not commit to any Recovery Point Objective (RPO) or Recovery Time Objective (RTO).
- Customers should take regular backups of their data.
- Backups performed by SRE are taken as a precautionary measure only. They are stored in the same region as the cluster.
- Customers can access SRE backup data on request by opening a support case.
- Red Hat highly encourages customers to deploy multi-AZ clusters with workloads that follow Kubernetes best practices to ensure high availability within a region.
- In the event an entire cloud region is unavailable, customers must install a new cluster in a different region and restore their apps using their backup data.

### 3.1.5. Cluster capacity

Evaluating and managing cluster capacity is a responsibility that is shared between Red Hat and the customer. Red Hat SRE is responsible for the capacity of all master and infrastructure nodes on the cluster.

Red Hat SRE also evaluates cluster capacity during upgrades and in response to cluster alerts. The impact of a cluster upgrade on capacity is evaluated as part of the upgrade testing process to ensure that capacity is not negatively impacted by new additions to the cluster. During a cluster upgrade, additional worker nodes are added to make sure that total cluster capacity is maintained during the upgrade process.

Capacity evaluations by SRE staff also happen in response to alerts from the cluster, once usage thresholds are exceeded for a certain period of time. Such alerts can also result in a notification to the customer.

## 3.2. CHANGE MANAGEMENT

Cluster changes are initiated in one of two ways:

- A customer initiates changes through self-service capabilities like cluster deployment, worker node scaling, and cluster deletion.
- An SRE initiates a change through Operator-driven capabilities like configuration, upgrade, patching, or configuration changes.

Change history is captured in the **Cluster History** section in OpenShift Cluster Manager (OCM) **Overview** tab and is available to customers. This includes logs from the following changes:

- Adding or removing identity providers
- Adding or removing users to/from the dedicated-admins group
- Scaling the cluster compute nodes
- Scaling the cluster load balancer
- Scaling the cluster persistent storage
- Upgrading the cluster

SRE-initiated changes that require manual intervention generally follow the below procedure:

- Preparing for change
    - Change characteristics are identified and a gap analysis against current state is performed.
    - Change steps are documented and validated.
    - Communication plan and schedule is shared with all stakeholders.
    - CICD and end-to-end tests are updated to automate change validation.
    - Change request capturing change details is submitted for management approval.
- Managing change
    - Automated nightly CI/CD jobs pick up the change and run tests.

    - The change is made to integration and stage environments, and manually validated before updating the customer cluster.
    - Major change notifications are sent before and after the event.

- Reinforcing the change
    - Feedback on the change is collected and analyzed.
    - Potential gaps are diagnosed in order to understand resistance and automate similar change requests.
    - Corrective actions are implemented.

**NOTE**

SREs consider manual changes a failure and this is only used as a fallback process.

### 3.2.1. Configuration management

The infrastructure and configuration of the OpenShift Dedicated environment is managed as code. Red Hat SRE manages changes to the OpenShift Dedicated environment using a GitOps workflow and automated CI/CD pipeline.

Each proposed change undergoes a series of automated verifications immediately upon check-in. Changes are then deployed to a staging environment where they undergo automated integration testing. Finally, changes are deployed to the production environment. Each step is fully automated.

An authorized SRE reviewer must approve advancement to each step. The reviewer might not be the same individual who proposed the change. All changes and approvals are fully auditable as part of the GitOps workflow.

### 3.2.2. Patch management

OpenShift Container Platform software and the underlying immutable Red Hat Enterprise Linux CoreOS (RHCOS) operating system image are patched for bugs and vulnerabilities as a side effect of regular z-stream upgrades. Read more about RHCOS architecture in the OpenShift Container Platform documentation.

### 3.2.3. Release management

OpenShift Dedicated clusters are upgraded as frequently as weekly to ensure that the latest security patches and bug fixes are applied to OpenShift Dedicated clusters.

Patch-level upgrades, also referred to as z-stream upgrades (for example, 4.3.18 to 4.3.19), are automatically deployed on Tuesdays. New z-stream releases are tested nightly with automated OpenShift Dedicated integration testing and released only once validated in the OSD environment.

Minor version upgrades, also referred to as y-stream upgrades (for example, 4.3 to 4.4), are coordinated with customers by email notification.

Customers can review the history of all cluster upgrade events in their OCM web console.

## 3.3. IDENTITY AND ACCESS MANAGEMENT

Most access by Red Hat site reliability engineering (SRE) teams is done by using cluster Operators through automated configuration management.

### 3.3.1. Subprocessors

For a list of the available subprocessors, see the Red Hat Subprocessor List on the Red Hat Customer Portal.

### 3.3.2. SRE access to all OpenShift Dedicated clusters

SREs access OpenShift Dedicated clusters through the web console or command-line tools. Authentication requires multi-factor authentication (MFA) with industry-standard requirements for password complexity and account lockouts. SREs must authenticate as individuals to ensure auditability. All authentication attempts are logged to a Security Information and Event Management (SIEM) system.

SREs access private clusters using an encrypted tunnel through a hardened SRE Support Pod running in the cluster. Connections to the SRE Support Pod are permitted only from a secured Red Hat network using an IP allow-list. In addition to the cluster authentication controls described above, authentication to the SRE Support Pod is controlled using SSH keys. SSH key authorization is limited to SRE staff and automatically synchronized with Red Hat corporate directory data. Corporate directory data is secured and controlled by HR systems, including management review, approval, and audits.

### 3.3.3. Privileged access controls in OpenShift Dedicated

Red Hat SRE adheres to the principle of least privilege when accessing OpenShift Dedicated and public cloud provider components. There are four basic categories of manual SRE access:

- SRE admin access through the Red Hat Customer Portal with normal two-factor authentication and no privileged elevation.
- SRE admin access through the Red Hat corporate SSO with normal two-factor authentication and no privileged elevation.
- OpenShift elevation, which is a manual elevation using Red Hat SSO. It is limited to 2 hours, is fully audited, and requires management approval.
- Cloud provider access or elevation, which is a manual elevation for cloud provider console or CLI access. Access is limited to 60 minutes and is fully audited.

Each of these access types has different levels of access to components:

| Component | Typical SRE admin access (Red Hat Customer Portal) | Typical SRE admin access (Red Hat SSO) | OpenShift elevation | Cloud provider access |
|---|---|---|---|---|
| OpenShift Cluster Manager (OCM) | R/W | No access | No access | No access |
| OpenShift web console | No access | R/W | R/W | No access |

| Component | Typical SRE admin access (Red Hat Customer Portal) | Typical SRE admin access (Red Hat SSO) | OpenShift elevation | Cloud provider access |
|---|---|---|---|---|
| Node operating system | No access | A specific list of elevated OS and network permissions. | A specific list of elevated OS and network permissions. | No access |
| AWS Console | No access | No access, but this is the account used to request cloud provider access. | No access | All cloud provider permissions using the SRE identity. |

### 3.3.4. SRE access to cloud infrastructure accounts

Red Hat personnel do not access cloud infrastructure accounts in the course of routine OpenShift Dedicated operations. For emergency troubleshooting purposes, Red Hat SRE have well-defined and auditable procedures to access cloud infrastructure accounts.

In AWS, SREs generate a short-lived AWS access token for the **BYOCAdminAccess** user using the AWS Security Token Service (STS). Access to the STS token is audit logged and traceable back to individual users. The **BYOCAdminAccess** has the **AdministratorAccess** IAM policy attached.

In Google Cloud, SREs access resources after being authenticated against a Red Hat SAML identity provider (IDP). The IDP authorizes tokens that have time-to-live expirations. The issuance of the token is auditable by corporate Red Hat IT and linked back to an individual user.

### 3.3.5. Red Hat support access

Members of the Red Hat CEE team typically have read-only access to parts of the cluster. Specifically, CEE has limited access to the core and product namespaces and does not have access to the customer namespaces.

| Role | Core namespace | Layered product namespace | Customer namespace | Cloud infrastructure account* |
|---|---|---|---|---|
| OpenShift SRE | Read: All<br>Write: Very Limited [1] | Read: All<br>Write: None | Read: None[2]<br>Write: None | Read: All [3]<br>Write: All [3] |
| CEE | Read: All<br>Write: None | Read: All<br>Write: None | Read: None[2]<br>Write: None | Read: None<br>Write: None |

| Role | Core namespace | Layered product namespace | Customer namespace | Cloud infrastructure account* |
|---|---|---|---|---|
| Customer administrator | Read: None<br>Write: None | Read: None<br>Write: None | Read: All<br>Write: All | Read: Limited[4]<br>Write: Limited[4] |
| Customer user | Read: None<br>Write: None | Read: None<br>Write: None | Read: Limited[5]<br>Write: Limited[5] | Read: None<br>Write: None |
| Everybody else | Read: None<br>Write: None | Read: None<br>Write: None | Read: None<br>Write: None | Read: None<br>Write: None |

Cloud Infrastructure Account refers to the underlying AWS or Google Cloud account

1. Limited to addressing common use cases such as failing deployments, upgrading a cluster, and replacing bad worker nodes.
2. Red Hat associates have no access to customer data by default.
3. SRE access to the cloud infrastructure account is a "break-glass" procedure for exceptional troubleshooting during a documented incident.
4. Customer administrator has limited access to the cloud infrastructure account console through Cloud Infrastructure Access.
5. Limited to what is granted through RBAC by the customer administrator, as well as namespaces created by the user.

### 3.3.6. Customer access

Customer access is limited to namespaces created by the customer and permissions that are granted using RBAC by the customer administrator role. Access to the underlying infrastructure or product namespaces is generally not permitted without **cluster-admin** access. More information on customer access and authentication can be found in the Understanding Authentication section of the documentation.

### 3.3.7. Access approval and review

New SRE user access requires management approval. Separated or transferred SRE accounts are removed as authorized users through an automated process. Additionally, SRE performs periodic access review including management sign-off of authorized user lists.

## 3.4. SECURITY AND REGULATION COMPLIANCE

Security and regulation compliance includes tasks, such as the implementation of security controls and compliance certification.

### 3.4.1. Data classification

Red Hat defines and follows a data classification standard to determine the sensitivity of data and highlight inherent risk to the confidentiality and integrity of that data while it is collected, used, transmitted stored, and processed. Customer-owned data is classified at the highest level of sensitivity and handling requirements.

### 3.4.2. Data management

OpenShift Dedicated uses cloud provider services to help securely manage keys for encrypted data (AWS KMS and Google Cloud KMS). These keys are used for control plane data volumes which are encrypted by default. Persistent volumes for customer applications also use these cloud services for key management.

When a customer deletes their OpenShift Dedicated cluster, all cluster data is permanently deleted, including control plane data volumes, customer application data volumes (PVs), and backup data.

### 3.4.3. Vulnerability management

Red Hat performs periodic vulnerability scanning of OpenShift Dedicated using industry standard tools. Identified vulnerabilities are tracked to their remediation according to timelines based on severity. Vulnerability scanning and remediation activities are documented for verification by third-party assessors in the course of compliance certification audits.

### 3.4.4. Network security

#### 3.4.4.1. Firewall and DDoS protection

Each OpenShift Dedicated cluster is protected by a secure network configuration at the cloud infrastructure level using firewall rules (AWS Security Groups or Google Cloud Compute Engine firewall rules). OpenShift Dedicated customers on AWS are also protected against DDoS attacks with AWS Shield Standard.

#### 3.4.4.2. Private clusters and network connectivity

Customers can optionally configure their OpenShift Dedicated cluster endpoints (web console, API, and application router) to be made private so that the cluster control plane or applications are not accessible from the Internet.

For AWS, customers can configure a private network connection to their OpenShift Dedicated cluster through AWS VPC peering, AWS VPN, or AWS Direct Connect.

**NOTE**

At this time, private clusters are not supported for OpenShift Dedicated clusters on Google Cloud.

#### 3.4.4.3. Cluster network access controls

Fine-grained network access control rules can be configured by customers per project by using **NetworkPolicy** objects and the OpenShift SDN.

### 3.4.5. Penetration testing

Red Hat performs periodic penetration tests against OpenShift Dedicated. Tests are performed by an independent internal team using industry standard tools and best practices.

Any issues that are discovered are prioritized based on severity. Any issues found belonging to open source projects are shared with the community for resolution.

### 3.4.6. Compliance

OpenShift Dedicated follows common industry best practices for security and controls. The certifications are outlined in the following table.

**Table 3.1. Security and control certifications for OpenShift Dedicated**

| Certification | OpenShift Dedicated on AWS | OpenShift Dedicated on GCP |
|---|---|---|
| ISO 27001 | Yes | Yes |
| PCI DSS | Yes | Yes |
| SOC 1 | Yes | Yes |
| SOC 2 Type 1 | Yes | Yes |
| SOC 2 Type 2 | Yes | Yes |

## 3.5. DISASTER RECOVERY

OpenShift Dedicated provides disaster recovery for failures that occur at the pod, worker node, infrastructure node, master node, and availability zone levels.

All disaster recovery requires that the customer use best practices for deploying highly available applications, storage, and cluster architecture (for example, single-zone deployment vs. multi-zone deployment) to account for the level of desired availability.

One single-zone cluster will not provide disaster avoidance or recovery in the event of an availability zone or region outage. Multiple single-zone clusters with customer-maintained failover can account for outages at the zone or region levels.

One multi-zone cluster will not provide disaster avoidance or recovery in the event of a full region outage. Multiple multi-zone clusters with customer-maintained failover can account for outages at the region level.

# CHAPTER 4. UNDERSTANDING AVAILABILITY FOR OPENSHIFT DEDICATED

Availability and disaster avoidance are extremely important aspects of any application platform. OpenShift Dedicated provides many protections against failures at several levels, but customer-deployed applications must be appropriately configured for high availability. In addition, to account for cloud provider outages that might occur, other options are available, such as deploying a cluster across multiple availability zones or maintaining multiple clusters with failover mechanisms.

## 4.1. POTENTIAL POINTS OF FAILURE

OpenShift Container Platform provides many features and options for protecting your workloads against downtime, but applications must be architected appropriately to take advantage of these features.

OpenShift Dedicated can help further protect you against many common Kubernetes issues by adding Red Hat Site Reliability Engineer (SRE) support and the option to deploy a multi-zone cluster, but there are a number of ways in which a container or infrastructure can still fail. By understanding potential points of failure, you can understand risks and appropriately architect both your applications and your clusters to be as resilient as necessary at each specific level.

**NOTE**

An outage can occur at several different levels of infrastructure and cluster components.

### 4.1.1. Container or pod failure

By design, pods are meant to exist for a short time. Appropriately scaling services so that multiple instances of your application pods are running protects against issues with any individual pod or container. The node scheduler can also ensure that these workloads are distributed across different worker nodes to further improve resiliency.

When accounting for possible pod failures, it is also important to understand how storage is attached to your applications. Single persistent volumes attached to single pods cannot leverage the full benefits of pod scaling, whereas replicated databases, database services, or shared storage can.

To avoid disruption to your applications during planned maintenance, such as upgrades, it is important to define a pod disruption budget. These are part of the Kubernetes API and can be managed with the OpenShift CLI (**oc**) like other object types. They allow the specification of safety constraints on pods during operations, such as draining a node for maintenance.

### 4.1.2. Worker node failure

Worker nodes are the virtual machines that contain your application pods. By default, an OpenShift Dedicated cluster has a minimum of four worker nodes for a single availability-zone cluster. In the event of a worker node failure, pods are relocated to functioning worker nodes, as long as there is enough capacity, until any issue with an existing node is resolved or the node is replaced. More worker nodes means more protection against single node outages, and ensures proper cluster capacity for rescheduled pods in the event of a node failure.

NOTE

When accounting for possible node failures, it is also important to understand how storage is affected.

### 4.1.3. Cluster failure

OpenShift Dedicated clusters have at least three master nodes and three infrastructure nodes that are preconfigured for high availability, either in a single zone or across multiple zones depending on the type of cluster you have selected. This means that master and infrastructure nodes have the same resiliency of worker nodes, with the added benefit of being managed completely by Red Hat.

In the event of a complete master outage, the OpenShift APIs will not function, and existing worker node pods will be unaffected. However, if there is also a pod or node outage at the same time, the masters will have to recover before new pods or nodes can be added or scheduled.

All services running on infrastructure nodes are configured by Red Hat to be highly available and distributed across infrastructure nodes. In the event of a complete infrastructure outage, these services will be unavailable until these nodes have been recovered.

### 4.1.4. Zone failure

A zone failure from a public cloud provider affects all virtual components, such as worker nodes, block or shared storage, and load balancers that are specific to a single availability zone. To protect against a zone failure, OpenShift Dedicated provides the option for clusters that are distributed across three availability zones, called multi-availability zone clusters. Existing stateless workloads are redistributed to unaffected zones in the event of an outage, as long as there is enough capacity.

### 4.1.5. Storage failure

If you have deployed a stateful application, then storage is a critical component and must be accounted for when thinking about high availability. A single block storage PV is unable to withstand outages even at the pod level. The best ways to maintain availability of storage are to use replicated storage solutions, shared storage that is unaffected by outages, or a database service that is independent of the cluster.

# CHAPTER 5. OPENSHIFT DEDICATED UPDATE LIFE CYCLE

## 5.1. OVERVIEW

Red Hat provides a published product life cycle for OpenShift Dedicated in order for customers and partners to effectively plan, deploy, and support their applications running on the platform. Red Hat publishes this life cycle in order to provide as much transparency as possible and might make exceptions from these policies as conflicts arise.

OpenShift Dedicated is a managed instance of Red Hat OpenShift and maintains an independent release schedule. More details about the managed offering can be found in the OpenShift Dedicated service definition. The availability of Security Advisories and Bug Fix Advisories for a specific version are dependent upon the Red Hat OpenShift Container Platform life cycle policy and subject to the OpenShift Dedicated maintenance schedule.

### Additional resources

- OpenShift Dedicated service definition

## 5.2. DEFINITIONS

**Table 5.1. Version reference**

| Version format | Major | Minor | Patch | Major.minor.patch |
|---|---|---|---|---|
| | x | y | z | x.y.z |
| Example | 4 | 5 | 21 | 4.5.21 |

### Major releases or X-releases

Referred to only as *major releases* or *X-releases* (X.y.z).

#### Examples

- "Major release 5" → 5.y.z
- "Major release 4" → 4.y.z
- "Major release 3" → 3.y.z

### Minor releases or Y-releases

Referred to only as *minor releases* or *Y-releases* (x.Y.z).

#### Examples

- "Minor release 4" → 4.4.z
- "Minor release 5" → 4.5.z
- "Minor release 6" → 4.6.z

**Patch releases or Z-releases**

Referred to only as *patch releases* or *Z-releases* (x.y.Z).

**Examples**

- "Patch release 14 of minor release 5" → 4.5.14
- "Patch release 25 of minor release 5" → 4.5.25
- "Patch release 26 of minor release 6" → 4.6.26

## 5.3. MAJOR VERSIONS (X.Y.Z)

Major versions of OpenShift Dedicated, for example version 4, are supported for one year following the release of a subsequent major version or the retirement of the product.

**Example**

- If version 5 were made available on OpenShift Dedicated on January 1, version 4 would be allowed to continue running on managed clusters for 12 months, until December 31. After this time, clusters would need to be upgraded or migrated to version 5.

## 5.4. MINOR VERSIONS (X.Y.Z)

Red Hat supports two minor versions of the major release.

- Y: The latest available minor release. For example, 4.8.
- Y-1: The previous minor version. For example, 4.7.

After an upgrade path from the previous minor version (Y-1) to the latest minor version (Y) is available, clusters running Y-2 must upgrade their cluster within a 30 day grace period. Any cluster remaining on Y-2 30 days after notification of upgrade availability are classified as being in limited support status until the cluster is upgraded to a supported release.

**Example**

1. A customer's cluster is currently running on 4.5.18. The latest version for 4.6 is 4.6.27.
2. On February 25, 4.7.2 is released as an available upgrade path from 4.6.27 and the customer is notified.
3. The cluster must be upgraded to 4.6.27 or later by March 25.
4. If the upgrade has not been performed, then the cluster will have SRE alerting disabled and will be unsupported until it is upgraded to 4.6.27 or later.

## 5.5. PATCH VERSIONS (X.Y.Z)

During the period in which a minor release is supported, all OpenShift Container Platform patch releases are supported unless otherwise specified.

For reasons of platform security and stability, a patch release might be deprecated, which would prevent installations of that release and trigger mandatory upgrades off that release.

**Example**

1. 4.7.6 is found to contain a critical CVE.
2. Any releases impacted by the CVE will be removed from the supported patch release list. In addition, any clusters running 4.7.6 will be scheduled for automatic upgrades within 48 hours.

## 5.6. LIMITED SUPPORT STATUS

While operating outside of the supported versions list, you might be asked to upgrade the cluster to a supported version when requesting support, unless you are within the 30-day grace period after version deprecation. Additionally, Red Hat does not make any runtime or SLA guarantees for clusters outside of the supported versions list at the end of the 30-day grace period.

Red Hat provides best effort to ensure an upgrade path from an unsupported release to a supported release is available. However, if a supported upgrade path is no longer available, you might be required to create a new cluster and migrate your workloads.

## 5.7. SUPPORTED VERSIONS EXCEPTION POLICY

Red Hat reserves the right to add or remove new or existing versions, or delay upcoming minor release versions, that have been identified to have one or more critical production impacting bugs or security issues without advance notice.

## 5.8. INSTALLATION POLICY

While Red Hat recommends installation of the latest support release, OpenShift Dedicated supports installation of any supported release as covered by the preceding policy.

## 5.9. MANDATORY UPGRADES

In the event that a Critical or Important CVE, or other bug identified by Red Hat, significantly impacts the security or stability of the cluster, the customer must upgrade to the next supported patch release within 48 hours.

In extreme circumstances and based on Red Hat's assessment of the CVE criticality to the environment, if the upgrade to the next supported patch release has not been performed within 48 hours of notification, the cluster will be automatically updated to the latest patch release to mitigate potential security breach or instability.

## 5.10. LIFE CYCLE DATES

| Version | General availability | End of life |
|---|---|---|
| 4.8 | Jul 27, 2021 | Release of 4.10 + 30 days |
| 4.7 | Feb 24, 2021 | Release of 4.9 + 30 days |

| Version | General availability | End of life |
|---|---|---|
| 4.6 | Oct 27, 2020 | Aug 26, 2021 |
| 4.5 | Sep 23, 2020 | Mar 26, 2021 |
| 4.4 | Sep 15, 2020 | Nov 26, 2020 |
| 4.3 | Feb 19, 2020 | Oct 23, 2020 |
| 4.2 | Nov 12, 2019 | Oct 15, 2020 |
| 4.1 | Jun 11, 2019 | Mar 20, 2020 |
| 3.11 | Oct 10, 2018 | Jul 31, 2021 [a] |

[a] https://access.redhat.com/articles/5254001