



# OpenShift Dedicated 4

## Planning your environment

An overview of planning for Dedicated 4



# OpenShift Dedicated 4 Planning your environment

---

An overview of planning for Dedicated 4

## Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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 planning considerations for OpenShift Dedicated cluster deployments.

## Table of Contents

<b>CHAPTER 1. LIMITS AND SCALABILITY</b> .....	<b>3</b>
1.1. CLUSTER MAXIMUMS	3
1.2. OPENSIFT CONTAINER PLATFORM TESTING ENVIRONMENT AND CONFIGURATION	4
1.3. CONTROL PLANE AND INFRASTRUCTURE NODE SIZING AND SCALING	5
1.3.1. Node sizing during installation	5
1.3.2. Node scaling after installation	6
1.3.3. Sizing considerations for larger clusters	6
<b>CHAPTER 2. CUSTOMER CLOUD SUBSCRIPTIONS ON AWS</b> .....	<b>7</b>
2.1. UNDERSTANDING CUSTOMER CLOUD SUBSCRIPTIONS ON AWS	7
2.2. CUSTOMER REQUIREMENTS	7
2.2.1. Account	7
2.2.2. Access requirements	8
2.2.3. Support requirements	8
2.2.4. Security requirements	8
2.3. REQUIRED CUSTOMER PROCEDURE	8
2.4. MINIMUM REQUIRED SERVICE CONTROL POLICY (SCP)	9
2.5. RED HAT MANAGED IAM REFERENCES FOR AWS	13
2.5.1. IAM policies	13
2.5.2. IAM users	14
2.5.3. IAM roles	15
2.6. PROVISIONED AWS INFRASTRUCTURE	15
2.6.1. AWS Elastic Computing (EC2) instances	15
2.6.2. AWS Elastic Block Store (EBS) storage	15
2.6.3. Elastic Load Balancing (ELB) load balancers	16
2.6.4. S3 storage	16
2.6.5. VPC	16
2.6.5.1. Sample VPC Architecture	17
2.6.6. Security groups	17
2.6.6.1. Additional custom security groups	17
2.7. AWS FIREWALL PREREQUISITES	17
2.8. AWS ACCOUNT LIMITS	23
<b>CHAPTER 3. CUSTOMER CLOUD SUBSCRIPTIONS ON GCP</b> .....	<b>26</b>
3.1. UNDERSTANDING CUSTOMER CLOUD SUBSCRIPTIONS ON GCP	26
3.2. CUSTOMER REQUIREMENTS	26
3.2.1. Account	26
3.2.2. Access requirements	26
3.2.3. Support requirements	27
3.2.4. Security requirements	27
3.3. REQUIRED CUSTOMER PROCEDURE	27
3.4. RED HAT MANAGED GOOGLE CLOUD RESOURCES	30
3.4.1. IAM service account and roles	30
3.4.2. IAM group and roles	31
3.5. PROVISIONED GCP INFRASTRUCTURE	32
3.5.1. Compute instances	32
3.5.2. Storage	32
3.5.3. VPC	32
3.5.4. Services	33
3.5.5. Permissions	33
3.6. GCP ACCOUNT LIMITS	34



# CHAPTER 1. LIMITS AND SCALABILITY

This document details the tested cluster maximums for OpenShift Dedicated clusters, along with information about the test environment and configuration used to test the maximums. Information about control plane and infrastructure node sizing and scaling is also provided.

## 1.1. CLUSTER MAXIMUMS

Consider the following tested object maximums when you plan a OpenShift Dedicated cluster installation. The table specifies the maximum limits for each tested type in a OpenShift Dedicated cluster.

These guidelines are based on a cluster of 102 compute (also known as worker) nodes in a multiple availability zone configuration. For smaller clusters, the maximums are lower.



### NOTE

The OpenShift Container Platform version used in all of the tests is OCP 4.8.0.

**Table 1.1. Tested cluster maximums**

Maximum type	4.8 tested maximum
Number of nodes	102
Number of pods <sup>[1]</sup>	20,400
Number of pods per node	250
Number of pods per core	There is no default value
Number of namespaces <sup>[2]</sup>	3,400
Number of pods per namespace <sup>[3]</sup>	20,400
Number of services <sup>[4]</sup>	10,000
Number of services per namespace	10,000
Number of back ends per service	10,000
Number of deployments per namespace <sup>[3]</sup>	1,000

1. The pod count displayed here is the number of test pods. The actual number of pods depends on the application's memory, CPU, and storage requirements.
2. When there are a large number of active projects, etcd can suffer from poor performance if the key space grows excessively large and exceeds the space quota. Periodic maintenance of etcd, including defragmentation, is highly recommended to make etcd storage available.

3. There are a number of control loops in the system that must iterate over all objects in a given namespace as a reaction to some changes in state. Having a large number of objects of a type, in a single namespace, can make those loops expensive and slow down processing the state changes. The limit assumes that the system has enough CPU, memory, and disk to satisfy the application requirements.
4. Each service port and each service back end has a corresponding entry in iptables. The number of back ends of a given service impacts the size of the endpoints objects, which then impacts the size of data that is sent throughout the system.

In OpenShift Container Platform 4.8, half of a CPU core (500 millicore) is reserved by the system compared to previous versions of OpenShift Container Platform.

## 1.2. OPENSIFT CONTAINER PLATFORM TESTING ENVIRONMENT AND CONFIGURATION

The following table lists the OpenShift Container Platform environment and configuration on which the cluster maximums are tested for the AWS cloud platform.

Node	Type	vCPU	RAM(GiB)	Disk type	Disk size(GiB) /IOPS	Count	Region
Control plane/etcd <sup>[1]</sup>	m5.4xlarge	16	64	gp3	350 / 1,000	3	us-west-2
Infrastructure nodes <sup>[2]</sup>	r5.2xlarge	8	64	gp3	300 / 900	3	us-west-2
Workload <sup>[3]</sup>	m5.2xlarge	8	32	gp3	350 / 900	3	us-west-2
Compute nodes	m5.2xlarge	8	32	gp3	350 / 900	102	us-west-2

1. io1 disks are used for control plane/etcd nodes in all versions prior to 4.10.
2. Infrastructure nodes are used to host monitoring components because Prometheus can claim a large amount of memory, depending on usage patterns.
3. Workload nodes are dedicated to run performance and scalability workload generators.

Larger cluster sizes and higher object counts might be reachable. However, the sizing of the infrastructure nodes limits the amount of memory that is available to Prometheus. When creating, modifying, or deleting objects, Prometheus stores the metrics in its memory for roughly 3 hours prior to persisting the metrics on disk. If the rate of creation, modification, or deletion of objects is too high, Prometheus can become overwhelmed and fail due to the lack of memory resources.



## 1.3. CONTROL PLANE AND INFRASTRUCTURE NODE SIZING AND SCALING

When you install a OpenShift Dedicated cluster, the sizing of the control plane and infrastructure nodes are automatically determined by the compute node count.

If you change the number of compute nodes in your cluster after installation, the Red Hat Site Reliability Engineering (SRE) team scales the control plane and infrastructure nodes as required to maintain cluster stability.

### 1.3.1. Node sizing during installation

During the installation process, the sizing of the control plane and infrastructure nodes are dynamically calculated. The sizing calculation is based on the number of compute nodes in a cluster.

The following tables list the control plane and infrastructure node sizing that is applied during installation.

AWS control plane and infrastructure node size:


Number of compute nodes	Control plane size	Infrastructure node size
1 to 25	m5.2xlarge	r5.xlarge
26 to 100	m5.4xlarge	r5.2xlarge
101 to 180	m5.8xlarge	r5.4xlarge

GCP control plane and infrastructure node size:

Number of compute nodes	Control plane size	Infrastructure node size
1 to 25	custom-8-32768	custom-4-32768-ext
26 to 100	custom-16-65536	custom-8-65536-ext
101 to 180	custom-32-131072	custom-16-131072-ext

GCP control plane and infrastructure node size for clusters created on or after 21 June 2024:

Number of compute nodes	Control plane size	Infrastructure node size
1 to 25	n2-standard-8	n2-highmem-4
26 to 100	n2-standard-16	n2-highmem-8
101 to 180	n2-standard-32	n2-highmem-16

**NOTE**

The maximum number of compute nodes on OpenShift Dedicated is 180.

### 1.3.2. Node scaling after installation

If you change the number of compute nodes after installation, the control plane and infrastructure nodes are scaled by the Red Hat Site Reliability Engineering (SRE) team as required. The nodes are scaled to maintain platform stability.

Postinstallation scaling requirements for control plane and infrastructure nodes are assessed on a case-by-case basis. Node resource consumption and received alerts are taken into consideration.

#### Rules for control plane node resizing alerts

The resizing alert is triggered for the control plane nodes in a cluster when the following occurs:

- Control plane nodes sustain over 66% utilization on average in a classic ROSA cluster.

**NOTE**

The maximum number of compute nodes on ROSA is 180.

#### Rules for infrastructure node resizing alerts

Resizing alerts are triggered for the infrastructure nodes in a cluster when it has high-sustained CPU or memory utilization. This high-sustained utilization status is:

- Infrastructure nodes sustain over 50% utilization on average in a classic ROSA cluster with a single availability zone using 2 infrastructure nodes.
- Infrastructure nodes sustain over 66% utilization on average in a classic ROSA cluster with multiple availability zones using 3 infrastructure nodes.

**NOTE**

The maximum number of compute nodes on OpenShift Dedicated is 180.

The resizing alerts only appear after sustained periods of high utilization. Short usage spikes, such as a node temporarily going down causing the other node to scale up, do not trigger these alerts.

The SRE team might scale the control plane and infrastructure nodes for additional reasons, for example to manage an increase in resource consumption on the nodes.

### 1.3.3. Sizing considerations for larger clusters

For larger clusters, infrastructure node sizing can become a significant impacting factor to scalability. There are many factors that influence the stated thresholds, including the etcd version or storage data format.

Exceeding these limits does not necessarily mean that the cluster will fail. In most cases, exceeding these numbers results in lower overall performance.

## CHAPTER 2. CUSTOMER CLOUD SUBSCRIPTIONS ON AWS

OpenShift Dedicated provides a Customer Cloud Subscription (CCS) model that allows Red Hat to deploy and manage clusters into a customer's existing Amazon Web Service (AWS) account.

### 2.1. UNDERSTANDING CUSTOMER CLOUD SUBSCRIPTIONS ON AWS

To deploy OpenShift Dedicated into your existing Amazon Web Services (AWS) account using the Customer Cloud Subscription (CCS) model, Red Hat requires several prerequisites be met.

Red Hat recommends the usage of an AWS Organization to manage multiple AWS accounts. The AWS Organization, managed by the customer, hosts multiple AWS accounts. There is a root account in the organization that all accounts will refer to in the account hierarchy.

It is recommended for the OpenShift Dedicated cluster using a CCS model to be hosted in an AWS account within an AWS Organizational Unit. A service control policy (SCP) is created and applied to the AWS Organizational Unit that manages what services the AWS sub-accounts are permitted to access. The SCP applies only to available permissions within a single AWS account for all AWS sub-accounts within the Organizational Unit. It is also possible to apply a SCP to a single AWS account. All other accounts in the customer's AWS Organization are managed in whatever manner the customer requires. Red Hat Site Reliability Engineers (SRE) will not have any control over SCPs within the AWS Organization.

### 2.2. CUSTOMER REQUIREMENTS

OpenShift Dedicated clusters using a Customer Cloud Subscription (CCS) model on Amazon Web Services (AWS) must meet several prerequisites before they can be deployed.

#### 2.2.1. Account

- The customer ensures that [AWS limits](#) are sufficient to support OpenShift Dedicated provisioned within the customer-provided AWS account.
- The customer-provided AWS account should be in the customer's AWS Organization with the applicable service control policy (SCP) applied.



#### NOTE

It is not a requirement that the customer-provided account be within an AWS Organization or for the SCP to be applied, however Red Hat must be able to perform all the actions listed in the SCP without restriction.

- The customer-provided AWS account must not be transferable to Red Hat.
- The customer may not impose AWS usage restrictions on Red Hat activities. Imposing restrictions severely hinders Red Hat's ability to respond to incidents.
- Red Hat deploys monitoring into AWS to alert Red Hat when a highly privileged account, such as a root account, logs into the customer-provided AWS account.
- The customer can deploy native AWS services within the same customer-provided AWS account.

**NOTE**

Customers are encouraged, but not mandated, to deploy resources in a Virtual Private Cloud (VPC) separate from the VPC hosting OpenShift Dedicated and other Red Hat supported services.

### 2.2.2. Access requirements

- To appropriately manage the OpenShift Dedicated service, Red Hat must have the **AdministratorAccess** policy applied to the administrator role at all times.

**NOTE**

This policy only provides Red Hat with permissions and capabilities to change resources in the customer-provided AWS account.

- Red Hat must have AWS console access to the customer-provided AWS account. This access is protected and managed by Red Hat.
- The customer must not utilize the AWS account to elevate their permissions within the OpenShift Dedicated cluster.
- Actions available in [OpenShift Cluster Manager](#) must not be directly performed in the customer-provided AWS account.

### 2.2.3. Support requirements

- Red Hat recommends that the customer have at least [Business Support](#) from AWS.
- Red Hat has authority from the customer to request AWS support on their behalf.
- Red Hat has authority from the customer to request AWS resource limit increases on the customer-provided account.
- Red Hat manages the restrictions, limitations, expectations, and defaults for all OpenShift Dedicated clusters in the same manner, unless otherwise specified in this requirements section.

### 2.2.4. Security requirements

- The customer-provided IAM credentials must be unique to the customer-provided AWS account and must not be stored anywhere in the customer-provided AWS account.
- Volume snapshots will remain within the customer-provided AWS account and customer-specified region.
- Red Hat must have ingress access to EC2 hosts and the API server through white-listed Red Hat machines.
- Red Hat must have egress allowed to forward system and audit logs to a Red Hat managed central logging stack.

## 2.3. REQUIRED CUSTOMER PROCEDURE

The Customer Cloud Subscription (CCS) model allows Red Hat to deploy and manage OpenShift Dedicated into a customer's Amazon Web Services (AWS) account. Red Hat requires several prerequisites in order to provide these services.

### Procedure

1. If the customer is using AWS Organizations, you must either use an AWS account within your organization or [create a new one](#).
2. To ensure that Red Hat can perform necessary actions, you must either create a service control policy (SCP) or ensure that none is applied to the AWS account.
3. [Attach](#) the SCP to the AWS account.
4. Within the AWS account, you must [create](#) an **osdCcsAdmin** IAM user with the following requirements:
  - This user needs at least **Programmatic access** enabled.
  - This user must have the **AdministratorAccess** policy attached to it.
5. Provide the IAM user credentials to Red Hat.
  - You must provide the **access key ID** and **secret access key** in [OpenShift Cluster Manager](#).

## 2.4. MINIMUM REQUIRED SERVICE CONTROL POLICY (SCP)

Service control policy (SCP) management is the responsibility of the customer. These policies are maintained in the AWS Organization and control what services are available within the attached AWS accounts.

Required/optional	Service	Actions	Effect
Required	Amazon EC2	All	Allow
	Amazon EC2 Auto Scaling	All	Allow
	Amazon S3	All	Allow
	Identity And Access Management	All	Allow
	Elastic Load Balancing	All	Allow
	Elastic Load Balancing V2	All	Allow
	Amazon CloudWatch	All	Allow
	Amazon CloudWatch Events	All	Allow

Required/optional	Service	Actions	Effect
	Amazon CloudWatch Logs	All	Allow
	AWS Support	All	Allow
	AWS Key Management Service	All	Allow
	AWS Security Token Service	All	Allow
	AWS Resource Tagging	All	Allow
	AWS Route53 DNS	All	Allow
	AWS Service Quotas	ListServices GetRequestedServiceQuotaChange GetServiceQuota RequestServiceQuotaIncrease ListServiceQuotas	Allow
Optional	AWS Billing	ViewAccount ViewBilling ViewUsage	Allow
	AWS Cost and Usage Report	All	Allow
	AWS Cost Explorer Services	All	Allow

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}

```

```

    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "autoscaling:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "s3:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "iam:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "elasticloadbalancing:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "cloudwatch:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "events:*"
    ],
    "Resource": [
      "*"
    ]
  },
},

```

```
{
  "Effect": "Allow",
  "Action": [
    "logs:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "support:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "kms:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "sts:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "tag:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "route53:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Effect": "Allow",
```



```

    "Action": [
      "servicequotas:ListServices",
      "servicequotas:GetRequestedServiceQuotaChange",
      "servicequotas:GetServiceQuota",
      "servicequotas:RequestServiceQuotaIncrease",
      "servicequotas:ListServiceQuotas"
    ],
    "Resource": [
      "*"
    ]
  }
]
}

```

## 2.5. RED HAT MANAGED IAM REFERENCES FOR AWS

Red Hat is responsible for creating and managing the following Amazon Web Services (AWS) resources: IAM policies, IAM users, and IAM roles.

### 2.5.1. IAM policies



#### NOTE

IAM policies are subject to modification as the capabilities of OpenShift Dedicated change.

- The **AdministratorAccess** policy is used by the administration role. This policy provides Red Hat the access necessary to administer the OpenShift Dedicated cluster in the customer-provided AWS account.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "*",
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}

```

- The **CustomerAdministratorAccess** role provides the customer access to administer a subset of services within the AWS account. At this time, the following are allowed:
  - VPC Peering
  - VPN Setup
  - Direct Connect (only available if granted through the service control policy)

```

{
  "Version": "2012-10-17",
  "Statement": [
    {

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:AttachVpnGateway",
      "ec2:DescribeVpnConnections",
      "ec2:AcceptVpcPeeringConnection",
      "ec2>DeleteVpcPeeringConnection",
      "ec2:DescribeVpcPeeringConnections",
      "ec2:CreateVpnConnectionRoute",
      "ec2:RejectVpcPeeringConnection",
      "ec2:DetachVpnGateway",
      "ec2>DeleteVpnConnectionRoute",
      "ec2>DeleteVpnGateway",
      "ec2:DescribeVpcs",
      "ec2:CreateVpnGateway",
      "ec2:ModifyVpcPeeringConnectionOptions",
      "ec2>DeleteVpnConnection",
      "ec2:CreateVpcPeeringConnection",
      "ec2:DescribeVpnGateways",
      "ec2:CreateVpnConnection",
      "ec2:DescribeRouteTables",
      "ec2:CreateTags",
      "ec2:CreateRoute",
      "directconnect:*"
    ],
    "Resource": "*"
  }
]
}

```

- If enabled, the **BillingReadOnlyAccess** role provides read-only access to view billing and usage information for the account.

Billing and usage access is only granted if the root account in the AWS Organization has it enabled. This is an optional step the customer must perform to enable read-only billing and usage access and does not impact the creation of this profile and the role that uses it. If this role is not enabled, users will not see billing and usage information. See this tutorial on [how to enable access to billing data](#).

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "aws-portal:ViewAccount",
        "aws-portal:ViewBilling"
      ],
      "Resource": "*"
    }
  ]
}

```

### 2.5.2. IAM users

The **osdManagedAdmin** user is created immediately after taking control of the customer-provided AWS account. This is the user that will perform the OpenShift Dedicated cluster installation.

### 2.5.3. IAM roles

- The **network-mgmt** role provides customer-federated administrative access to the AWS account through a separate AWS account. It also has the same access as a read-only role. The **network-mgmt** role only applies to non-Customer Cloud Subscription (CCS) clusters. The following policies are attached to the role:
  - AmazonEC2ReadOnlyAccess
  - CustomerAdministratorAccess
- The **read-only** role provides customer-federated read-only access to the AWS account through a separate AWS account. The following policies are attached to the role:
  - AWSAccountUsageReportAccess
  - AmazonEC2ReadOnlyAccess
  - AmazonS3ReadOnlyAccess
  - IAMReadOnlyAccess
  - BillingReadOnlyAccess

## 2.6. PROVISIONED AWS INFRASTRUCTURE

This is an overview of the provisioned Amazon Web Services (AWS) components on a deployed OpenShift Dedicated cluster. For a more detailed listing of all provisioned AWS components, see the [OpenShift Container Platform documentation](#).

### 2.6.1. AWS Elastic Computing (EC2) instances

AWS EC2 instances are required to deploy the control plane and data plane functions of OpenShift Dedicated in the AWS public cloud. Instance types might vary for control plane and infrastructure nodes depending on worker node count.

- Single availability zone
  - 3 m5.2xlarge minimum (control plane nodes)
  - 2 r5.xlarge minimum (infrastructure nodes)
  - 2 m5.xlarge minimum but highly variable (worker nodes)
- Multiple availability zones
  - 3 m5.2xlarge minimum (control plane nodes)
  - 3 r5.xlarge minimum (infrastructure nodes)
  - 3 m5.xlarge minimum but highly variable (worker nodes)

### 2.6.2. AWS Elastic Block Store (EBS) storage

Amazon EBS block storage is used for both local node storage and persistent volume storage.

Volume requirements for each EC2 instance:

- Control plane volumes
  - Size: 350 GB
  - Type: io1
  - Input/output operations per second: 1000
- Infrastructure volumes
  - Size: 300 GB
  - Type: gp2
  - Input/output operations per second: 900
- Worker volumes
  - Size: 300 GB
  - Type: gp2
  - Input/output operations per second: 900

### 2.6.3. Elastic Load Balancing (ELB) load balancers

Up to two Network Load Balancers for API and up to two Classic Load Balancers for application router. For more information, see the [ELB documentation for AWS](#).

### 2.6.4. S3 storage

The image registry and Elastic Block Store (EBS) volume snapshots are backed by AWS S3 storage. Pruning of resources is performed regularly to optimize S3 usage and cluster performance.




#### NOTE

Two buckets are required with a typical size of 2 TB each.

### 2.6.5. VPC

Customers should expect to see one VPC per cluster. Additionally, the VPC needs the following configurations:

- **Subnets:** Two subnets for a cluster with a single availability zone, or six subnets for a cluster with multiple availability zones.



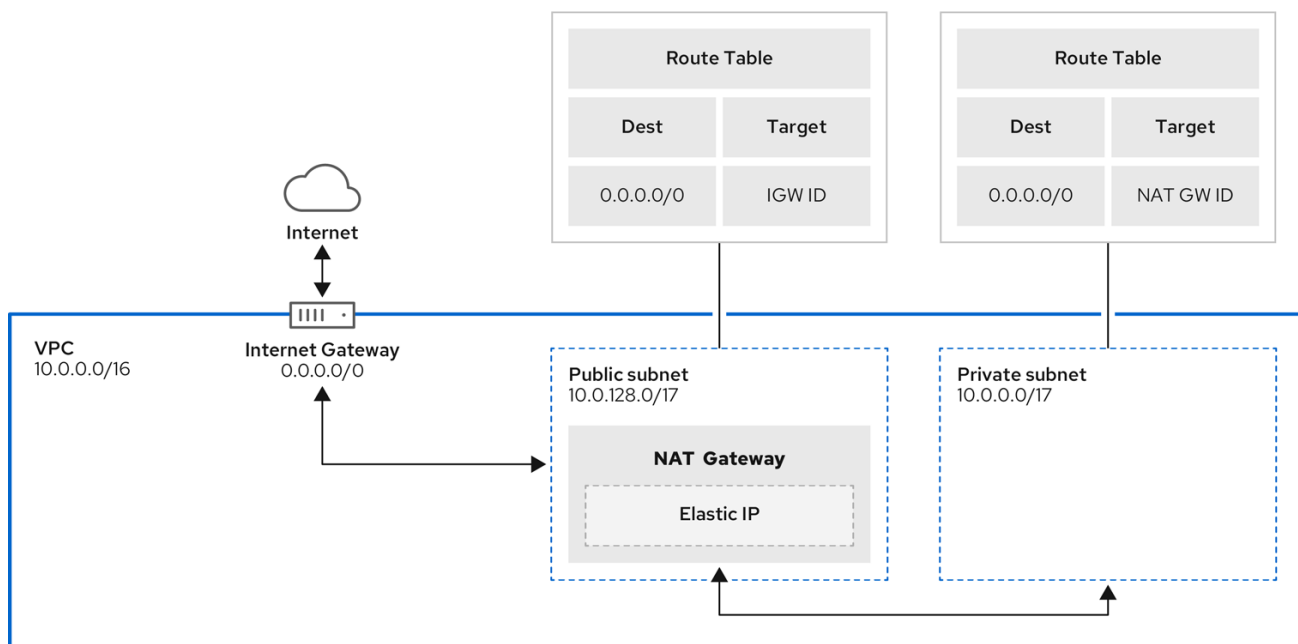
#### NOTE

A **public subnet** connects directly to the internet through an internet gateway. A **private subnet** connects to the internet through a network address translation (NAT) gateway.

- **Route tables:** One route table per private subnet, and one additional table per cluster.
- **Internet gateways:** One Internet Gateway per cluster.

- **NAT gateways:** One NAT Gateway per public subnet.

### 2.6.5.1. Sample VPC Architecture



204\_OpenShift\_0122

## 2.6.6. Security groups

AWS security groups provide security at the protocol and port-access level; they are associated with EC2 instances and Elastic Load Balancing. Each security group contains a set of rules that filter traffic coming in and out of an EC2 instance. You must ensure the ports required for the [OpenShift Container Platform installation](#) are open on your network and configured to allow access between hosts.

### 2.6.6.1. Additional custom security groups

When you create a cluster by using a non-managed VPC, you can add custom security groups during cluster creation. Custom security groups are subject to the following limitations:

- You must create the custom security groups in AWS before you create the cluster. For more information, see [Amazon EC2 security groups for Linux instances](#).
- You must associate the custom security groups with the VPC that the cluster will be installed into. Your custom security groups cannot be associated with another VPC.
- You might need to request additional quota for your VPC if you are adding additional custom security groups. For information on requesting an AWS quota increase, see [Requesting a quota increase](#).

## 2.7. AWS FIREWALL PREREQUISITES

If you are using a firewall to control egress traffic from OpenShift Dedicated, you must configure your firewall to grant access to the certain domain and port combinations below. OpenShift Dedicated requires this access to provide a fully managed OpenShift service.

### Prerequisites

- You have configured an Amazon S3 gateway endpoint in your AWS Virtual Private Cloud (VPC). This endpoint is required to complete requests from the cluster to the Amazon S3 service.

## Procedure

- Allowlist the following URLs that are used to install and download packages and tools:

Domain	Port	Function
<b>registry.redhat.io</b>	443	Provides core container images.
<b>quay.io</b>	443	Provides core container images.
<b>cdn01.quay.io</b>	443	Provides core container images.
<b>cdn02.quay.io</b>	443	Provides core container images.
<b>cdn03.quay.io</b>	443	Provides core container images.
<b>sso.redhat.com</b>	443	Required. The <a href="https://console.redhat.com/openshift">https://console.redhat.com/openshift</a> site uses authentication from <b>sso.redhat.com</b> to download the pull secret and use Red Hat SaaS solutions to facilitate monitoring of your subscriptions, cluster inventory, chargeback reporting, and so on.
<b>quay-registry.s3.amazonaws.com</b>	443	Provides core container images.
<b>ocm-quay-production-s3.s3.amazonaws.com</b>	443	Provides core container images.
<b>quayio-production-s3.s3.amazonaws.com</b>	443	Provides core container images.
<b>cart-rhcos-ci.s3.amazonaws.com</b>	443	Provides Red Hat Enterprise Linux CoreOS (RHCOS) images.
<b>openshift.org</b>	443	Provides Red Hat Enterprise Linux CoreOS (RHCOS) images.
<b>registry.access.redhat.com</b>	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog. Additionally, the registry provides access to the <b>odo</b> CLI tool that helps developers build on OpenShift and Kubernetes.

Domain	Port	Function
<b>access.redhat.com</b>	443	Required. Hosts a signature store that a container client requires for verifying images when pulling them from <b>registry.access.redhat.com</b> .
<b>registry.connect.redhat.com</b>	443	Required for all third-party images and certified Operators.
<b>console.redhat.com</b>	443	Required. Allows interactions between the cluster and OpenShift Console Manager to enable functionality, such as scheduling upgrades.
<b>sso.redhat.com</b>	443	The <a href="https://console.redhat.com/openshift">https://console.redhat.com/openshift</a> site uses authentication from <b>sso.redhat.com</b> .
<b>pull.q1w2.quay.rhcloud.com</b>	443	Provides core container images as a fallback when quay.io is not available.
<b>.q1w2.quay.rhcloud.com</b>	443	Provides core container images as a fallback when quay.io is not available.
<b>www.okd.io</b>	443	The <b>openshift.org</b> site redirects through <b>www.okd.io</b> .
<b>www.redhat.com</b>	443	The <b>sso.redhat.com</b> site redirects through <b>www.redhat.com</b> .
<b>aws.amazon.com</b>	443	The <b>iam.amazonaws.com</b> and <b>sts.amazonaws.com</b> sites redirect through <b>aws.amazon.com</b> .
<b>catalog.redhat.com</b>	443	The <b>registry.access.redhat.com</b> and <a href="https://registry.redhat.io">https://registry.redhat.io</a> sites redirect through <b>catalog.redhat.com</b> .
<b>dvbwgdztaeq9o.cloudfront.net</b> <sup>[1]</sup>	443	Used by ROSA for STS implementation with managed OIDC configuration.

1. The string of alphanumeric characters before **cloudfront.net** could change if there is a major cloudfront outage that requires redirecting the resource.
2. Allowlist the following telemetry URLs:

Domain	Port	Function
<b>cert-api.access.redhat.com</b>	443	Required for telemetry.
<b>api.access.redhat.com</b>	443	Required for telemetry.
<b>infogw.api.openshift.com</b>	443	Required for telemetry.
<b>console.redhat.com</b>	443	Required for telemetry and Red Hat Insights.
<b>cloud.redhat.com/api/ingress</b>	443	Required for telemetry and Red Hat Insights.
<b>observatorium-mst.api.openshift.com</b>	443	Required for managed OpenShift-specific telemetry.
<b>observatorium.api.openshift.com</b>	443	Required for managed OpenShift-specific telemetry.

Managed clusters require enabling telemetry to allow Red Hat to react more quickly to problems, better support the customers, and better understand how product upgrades impact clusters. For more information about how remote health monitoring data is used by Red Hat, see *About remote health monitoring* in the *Additional resources* section.

3. Allowlist the following Amazon Web Services (AWS) API URIs:

Domain	Port	Function
<b>.amazonaws.com</b>	443	Required to access AWS services and resources.

Alternatively, if you choose to not use a wildcard for Amazon Web Services (AWS) APIs, you must allowlist the following URLs:

Domain	Port	Function
<b>ec2.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.
<b>events. &lt;aws_region&gt;.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.
<b>iam.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.
<b>route53.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.



Domain	Port	Function
<b>sts.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment, for clusters configured to use the global endpoint for AWS STS.
<b>sts.&lt;aws_region&gt;.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment, for clusters configured to use regionalized endpoints for AWS STS. See <a href="#">AWS STS regionalized endpoints</a> for more information.
<b>tagging.us-east-1.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment. This endpoint is always us-east-1, regardless of the region the cluster is deployed in.
<b>ec2.&lt;aws_region&gt;.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.
<b>elasticloadbalancing.&lt;aws_region&gt;.amazonaws.com</b>	443	Used to install and manage clusters in an AWS environment.
<b>servicequotas.&lt;aws_region&gt;.amazonaws.com</b>	443	Required. Used to confirm quotas for deploying the service.
<b>tagging.&lt;aws_region&gt;.amazonaws.com</b>	443	Allows the assignment of metadata about AWS resources in the form of tags.

4. Allowlist the following OpenShift URLs:

Domain	Port	Function
<b>mirror.openshift.com</b>	443	Used to access mirrored installation content and images. This site is also a source of release image signatures, although the Cluster Version Operator (CVO) needs only a single functioning source.
<b>storage.googleapis.com/openshift-release</b> (Recommended)	443	Alternative site to mirror.openshift.com/. Used to download platform release signatures that are used by the cluster to know what images to pull from quay.io.
<b>api.openshift.com</b>	443	Used to check if updates are available for the cluster.

5. Allowlist the following site reliability engineering (SRE) and management URLs:

Domain	Port	Function
<b>api.pagerduty.com</b>	443	This alerting service is used by the in-cluster alertmanager to send alerts notifying Red Hat SRE of an event to take action on.
<b>events.pagerduty.com</b>	443	This alerting service is used by the in-cluster alertmanager to send alerts notifying Red Hat SRE of an event to take action on.
<b>api.deadmanssnitch.com</b>	443	Alerting service used by OpenShift Dedicated to send periodic pings that indicate whether the cluster is available and running.
<b>nosnch.in</b>	443	Alerting service used by OpenShift Dedicated to send periodic pings that indicate whether the cluster is available and running.
<b>.osdsecuritylogs.splunkcloud.com OR inputs1.osdsecuritylogs.splunkcloud.cominputs2.osdsecuritylogs.splunkcloud.cominputs4.osdsecuritylogs.splunkcloud.cominputs5.osdsecuritylogs.splunkcloud.cominputs6.osdsecuritylogs.splunkcloud.cominputs7.osdsecuritylogs.splunkcloud.cominputs8.osdsecuritylogs.splunkcloud.cominputs9.osdsecuritylogs.splunkcloud.cominputs10.osdsecuritylogs.splunkcloud.cominputs11.osdsecuritylogs.splunkcloud.cominputs12.osdsecuritylogs.splunkcloud.cominputs13.osdsecuritylogs.splunkcloud.cominputs14.osdsecuritylogs.splunkcloud.cominputs15.osdsecuritylogs.splunkcloud.com</b>	9997	Used by the <b>splunk-forwarder-operator</b> as a logging forwarding endpoint to be used by Red Hat SRE for log-based alerting.
<b>http-inputs-osdsecuritylogs.splunkcloud.com</b>	443	Required. Used by the <b>splunk-forwarder-operator</b> as a logging forwarding endpoint to be used by Red Hat SRE for log-based alerting.
<b>sftp.access.redhat.com</b> (Recommended)	22	The SFTP server used by <b>must-gather-operator</b> to upload diagnostic logs to help troubleshoot issues with the cluster.

- Allowlist the following URLs for optional third-party content:

Domain	Port	Function
<b>registry.connect.redhat.com</b>	443	Required for all third-party-images and certified operators.
<b>rhc4tp-prod-z8cxf-image-registry-us-east-1-evenkyleffocxqvofrk.s3.dualstack.us-east-1.amazonaws.com</b>	443	Provides access to container images hosted on <b>registry.connect.redhat.com</b>
<b>oso-rhc4tp-docker-registry.s3-us-west-2.amazonaws.com</b>	443	Required for Sonatype Nexus, F5 Big IP operators.

7. Allowlist any site that provides resources for a language or framework that your builds require.
8. Allowlist any outbound URLs that depend on the languages and frameworks used in OpenShift. See [OpenShift Outbound URLs to Allow](#) for a list of recommended URLs to be allowed on the firewall or proxy.

#### Additional resources

- [About remote health monitoring](#)

## 2.8. AWS ACCOUNT LIMITS

The OpenShift Dedicated cluster uses a number of Amazon Web Services (AWS) components, and the default [service limits](#) affect your ability to install OpenShift Dedicated clusters. If you use certain cluster configurations, deploy your cluster in certain AWS regions, or run multiple clusters from your account, you might need to request additional resources for your AWS account.

The following table summarizes the AWS components whose limits can impact your ability to install and run OpenShift Dedicated clusters.

Component	Number of clusters available by default	Default AWS limit	Description
-----------	---	-------------------	-------------

Component	Number of clusters available by default	Default AWS limit	Description
Instance Limits	Varies	Varies	<p>At a minimum, each cluster creates the following instances:</p> <ul style="list-style-type: none"> <li>● One bootstrap machine, which is removed after installation</li> <li>● Three control plane nodes</li> <li>● Two infrastructure nodes for a single availability zone; three infrastructure nodes for multi-availability zones</li> <li>● Two worker nodes for a single availability zone; three worker nodes for multi-availability zones</li> </ul> <p>These instance type counts are within a new account's default limit. To deploy more worker nodes, deploy large workloads, or use a different instance type, review your account limits to ensure that your cluster can deploy the machines that you need.</p> <p>In most regions, the bootstrap and worker machines use <b>m4.large</b> machines and the control plane machines use <b>m4.xlarge</b> instances. In some regions, including all regions that do not support these instance types, <b>m5.large</b> and <b>m5.xlarge</b> instances are used instead.</p>
Elastic IPs (EIPs)	0 to 1	5 EIPs per account	<p>To provision the cluster in a highly available configuration, the installation program creates a public and private subnet for each <a href="#">availability zone within a region</a>. Each private subnet requires a <a href="#">NAT Gateway</a>, and each NAT gateway requires a separate <a href="#">elastic IP</a>. Review the <a href="#">AWS region map</a> to determine how many availability zones are in each region. To take advantage of the default high availability, install the cluster in a region with at least three availability zones. To install a cluster in a region with more than five availability zones, you must increase the EIP limit.</p> <div data-bbox="863 1823 968 1989" style="background-color: #333; color: #fff; padding: 5px; width: 60px; height: 70px; margin-bottom: 10px;"> </div> <p><b>IMPORTANT</b></p> <p>To use the <b>us-east-1</b> region, you must increase the EIP limit for your account.</p>

Component	Number of clusters available by default	Default AWS limit	Description
Virtual Private Clouds (VPCs)	5	5 VPCs per region	Each cluster creates its own VPC.
Elastic Load Balancing (ELB)	3	20 per region	By default, each cluster creates internal and external Network Load Balancers for the primary API server and a single Classic Load Balancer for the router. Deploying more Kubernetes LoadBalancer Service objects will create additional <a href="#">load balancers</a> .
NAT Gateways	5	5 per availability zone	The cluster deploys one NAT gateway in each availability zone.
Elastic Network Interfaces (ENIs)	At least 12	350 per region	<p>The default installation creates 21 ENIs and an ENI for each availability zone in your region. For example, the <b>us-east-1</b> region contains six availability zones, so a cluster that is deployed in that zone uses 27 ENIs. Review the <a href="#">AWS region map</a> to determine how many availability zones are in each region.</p> <p>Additional ENIs are created for additional machines and load balancers that are created by cluster usage and deployed workloads.</p>
VPC Gateway	20	20 per account	Each cluster creates a single VPC Gateway for S3 access.
S3 buckets	99	100 buckets per account	Because the installation process creates a temporary bucket and the registry component in each cluster creates a bucket, you can create only 99 OpenShift Dedicated clusters per AWS account.
Security Groups	250	2,500 per account	Each cluster creates 10 distinct security groups.

## CHAPTER 3. CUSTOMER CLOUD SUBSCRIPTIONS ON GCP

OpenShift Dedicated provides a Customer Cloud Subscription (CCS) model that allows Red Hat to deploy and manage clusters in a customer's existing Google Cloud Platform (GCP) account.

### 3.1. UNDERSTANDING CUSTOMER CLOUD SUBSCRIPTIONS ON GCP

Red Hat OpenShift Dedicated provides a Customer Cloud Subscription (CCS) model that allows Red Hat to deploy and manage OpenShift Dedicated into a customer's existing Google Cloud Platform (GCP) account. Red Hat requires several prerequisites be met in order to provide this service.

Red Hat recommends the usage of GCP project, managed by the customer, to organize all of your GCP resources. A project consists of a set of users and APIs, as well as billing, authentication, and monitoring settings for those APIs.

It is recommended for the OpenShift Dedicated cluster using a CCS model to be hosted in a GCP project within a GCP organization. The Organization resource is the root node of the GCP resource hierarchy and all resources that belong to an organization are grouped under the organization node. An IAM service account with certain roles granted is created and applied to the GCP project. When you make calls to the API, you typically provide service account keys for authentication. Each service account is owned by a specific project, but service accounts can be provided roles to access resources for other projects.

### 3.2. CUSTOMER REQUIREMENTS

OpenShift Dedicated clusters using a Customer Cloud Subscription (CCS) model on Google Cloud Platform (GCP) must meet several prerequisites before they can be deployed.

#### 3.2.1. Account

- The customer ensures that [Google Cloud limits](#) are sufficient to support OpenShift Dedicated provisioned within the customer-provided GCP account.
- The customer-provided GCP account should be in the customer's Google Cloud Organization with the applicable Service Account applied.
- The customer-provided GCP account must not be transferable to Red Hat.
- The customer may not impose GCP usage restrictions on Red Hat activities. Imposing restrictions severely hinders Red Hat's ability to respond to incidents.
- Red Hat deploys monitoring into GCP to alert Red Hat when a highly privileged account, such as a root account, logs into the customer-provided GCP account.
- The customer can deploy native GCP services within the same customer-provided GCP account.



#### NOTE

Customers are encouraged, but not mandated, to deploy resources in a Virtual Private Cloud (VPC) separate from the VPC hosting OpenShift Dedicated and other Red Hat supported services.

#### 3.2.2. Access requirements

- To appropriately manage the OpenShift Dedicated service, Red Hat must have the **AdministratorAccess** policy applied to the administrator role at all times.



#### NOTE

This policy only provides Red Hat with permissions and capabilities to change resources in the customer-provided GCP account.

- Red Hat must have GCP console access to the customer-provided GCP account. This access is protected and managed by Red Hat.
- The customer must not utilize the GCP account to elevate their permissions within the OpenShift Dedicated cluster.
- Actions available in the [OpenShift Cluster Manager](#) must not be directly performed in the customer-provided GCP account.

### 3.2.3. Support requirements

- Red Hat recommends that the customer have at least [Production Support](#) from GCP.
- Red Hat has authority from the customer to request GCP support on their behalf.
- Red Hat has authority from the customer to request GCP resource limit increases on the customer-provided account.
- Red Hat manages the restrictions, limitations, expectations, and defaults for all OpenShift Dedicated clusters in the same manner, unless otherwise specified in this requirements section.

### 3.2.4. Security requirements

- The customer-provided IAM credentials must be unique to the customer-provided GCP account and must not be stored anywhere in the customer-provided GCP account.
- Volume snapshots will remain within the customer-provided GCP account and customer-specified region.
- Red Hat must have ingress access to the API server through white-listed Red Hat machines.
- Red Hat must have egress allowed to forward system and audit logs to a Red Hat managed central logging stack.

## 3.3. REQUIRED CUSTOMER PROCEDURE

The Customer Cloud Subscription (CCS) model allows Red Hat to deploy and manage OpenShift Dedicated into a customer's Google Cloud Platform (GCP) project. Red Hat requires several prerequisites to provide these services.

**WARNING**

To use OpenShift Dedicated in your GCP project, the following GCP organizational policy constraints cannot be in place:

- **constraints/iam.allowedPolicyMemberDomains** (This policy constraint is supported only if Red Hat's **DIRECTORY\_CUSTOMER\_ID C02k0I5e8** is included in the allow list. Use this policy constraint with caution).
- **constraints/compute.restrictLoadBalancerCreationForTypes**
- **constraints/compute.requireShieldedVm** (This policy constraint is supported only if the cluster is installed with "Enable Secure Boot support for Shielded VMs" selected during the initial cluster creation).
- **constraints/compute.vmExternalIpAccess** (This policy constraint is supported only after installation).

**Procedure**

1. [Create a Google Cloud project](#) to host the OpenShift Dedicated cluster.

**NOTE**

The project name must be 10 characters or less.

2. [Enable](#) the following required APIs in the project that hosts your OpenShift Dedicated cluster:

**Table 3.1. Required API services**

API service	Console service name
<a href="#">Cloud Deployment Manager V2 API</a>	<b>deploymentmanager.googleapis.com</b>
<a href="#">Compute Engine API</a>	<b>compute.googleapis.com</b>
<a href="#">Google Cloud APIs</a>	<b>cloudapis.googleapis.com</b>
<a href="#">Cloud Resource Manager API</a>	<b>cloudresourcemanager.googleapis.com</b>
<a href="#">Google DNS API</a>	<b>dns.googleapis.com</b>
<a href="#">Network Security API</a>	<b>networksecurity.googleapis.com</b>
<a href="#">IAM Service Account Credentials API</a>	<b>iamcredentials.googleapis.com</b>
<a href="#">Identity and Access Management (IAM) API</a>	<b>iam.googleapis.com</b>



API service	Console service name
<a href="#">Service Management API</a>	<b>servicemanagement.googleapis.com</b>
<a href="#">Service Usage API</a>	<b>serviceusage.googleapis.com</b>
<a href="#">Google Cloud Storage JSON API</a>	<b>storage-api.googleapis.com</b>
<a href="#">Cloud Storage</a>	<b>storage-component.googleapis.com</b>
<a href="#">Organization Policy API</a>	<b>orgpolicy.googleapis.com</b>

- To ensure that Red Hat can perform necessary actions, you must create an **osd-ccs-admin** IAM [service account](#) user within the GCP project.

The following roles must be [granted to the service account](#) :

**Table 3.2. Required roles**

Role	Console role name
Compute Admin	<b>roles/compute.admin</b>
DNS Administrator	<b>roles/dns.admin</b>
Organization Policy Viewer	<b>roles/orgpolicy.policyViewer</b>
Service Management Administrator	<b>roles/servicemanagement.admin</b>
Service Usage Admin	<b>roles/serviceusage.serviceUsageAdmin</b>
Storage Admin	<b>roles/storage.admin</b>
Compute Load Balancer Admin	<b>roles/compute.loadBalancerAdmin</b>
Role Viewer	<b>roles/viewer</b>
Role Administrator	<b>roles/iam.roleAdmin</b>
Security Admin	<b>roles/iam.securityAdmin</b>
Service Account Key Admin	<b>roles/iam.serviceAccountKeyAdmin</b>
Service Account Admin	<b>roles/iam.serviceAccountAdmin</b>

Role	Console role name
Service Account User	<b>roles/iam.serviceAccountUser</b>

4. [Create the service account key](#) for the **osd-ccs-admin** IAM service account. Export the key to a file named **osServiceAccount.json**; this JSON file will be uploaded in Red Hat OpenShift Cluster Manager when you create your cluster.

## 3.4. RED HAT MANAGED GOOGLE CLOUD RESOURCES

Red Hat is responsible for creating and managing the following IAM Google Cloud Platform (GCP) resources.

### 3.4.1. IAM service account and roles

The **osd-managed-admin** IAM service account is created immediately after taking control of the customer-provided GCP account. This is the user that will perform the OpenShift Dedicated cluster installation.

The following roles are attached to the service account:

**Table 3.3. IAM roles for osd-managed-admin**

Role	Console role name	Description
Compute Admin	<b>roles/compute.admin</b>	Provides full control of all Compute Engine resources.
DNS Administrator	<b>roles/dns.admin</b>	Provides read-write access to all Cloud DNS resources.
Security Admin	<b>roles/iam.securityAdmin</b>	Security admin role, with permissions to get and set any IAM policy.
Storage Admin	<b>roles/storage.admin</b>	Grants full control of objects and buckets.  When applied to an individual <b>bucket</b> , control applies only to the specified bucket and objects within the bucket.
Service Account Admin	<b>roles/iam.serviceAccountAdmin</b>	Create and manage service accounts.
Service Account Key Admin	<b>roles/iam.serviceAccountKeyAdmin</b>	Create and manage (and rotate) service account keys.

Role	Console role name	Description
Service Account User	<b>roles/iam.serviceAccountUser</b>	Run operations as the service account.
Role Administrator	<b>roles/iam.roleAdmin</b>	Provides access to all custom roles in the project.

### 3.4.2. IAM group and roles

The **sd-sre-platform-gcp-access** Google group is granted access to the GCP project to allow Red Hat Site Reliability Engineering (SRE) access to the console for emergency troubleshooting purposes.

The following roles are attached to the group:

**Table 3.4. IAM roles for sd-sre-platform-gcp-access**

Role	Console role name	Description
Compute Admin	<b>roles/compute.admin</b>	Provides full control of all Compute Engine resources.
Editor	<b>roles/editor</b>	Provides all viewer permissions, plus permissions for actions that modify state.
Organization Policy Viewer	<b>roles/orgpolicy.policyViewer</b>	Provides access to view Organization Policies on resources.
Project IAM Admin	<b>roles/resourceManager.projectIamAdmin</b>	Provides permissions to administer IAM policies on projects.
Quota Administrator	<b>roles/serviceManagement.quotaAdmin</b>	Provides access to administer service quotas.
Role Administrator	<b>roles/iam.roleAdmin</b>	Provides access to all custom roles in the project.
Service Account Admin	<b>roles/iam.serviceAccountAdmin</b>	Create and manage service accounts.
Service Usage Admin	<b>roles/serviceusage.serviceUsageAdmin</b>	Ability to enable, disable, and inspect service states, inspect operations, and consume quota and billing for a consumer project.

Role	Console role name	Description
Tech Support Editor	<b>roles/cloudsupport.techSupportEditor</b>	Provides full read-write access to technical support cases.

## 3.5. PROVISIONED GCP INFRASTRUCTURE

This is an overview of the provisioned Google Cloud Platform (GCP) components on a deployed OpenShift Dedicated cluster. For a more detailed listing of all provisioned GCP components, see the [OpenShift Container Platform documentation](#).

### 3.5.1. Compute instances

GCP compute instances are required to deploy the control plane and data plane functions of OpenShift Dedicated in GCP. Instance types might vary for control plane and infrastructure nodes depending on worker node count.

- Single availability zone
  - 2 infra nodes (custom machine type: 4 vCPU and 32 GB RAM)
  - 3 control plane nodes (custom machine type: 8 vCPU and 32 GB RAM)
  - 2 worker nodes (custom machine type: 4 vCPU and 16 GB RAM)
- Multiple availability zones
  - 3 infra nodes (custom machine type: 4 vCPU and 32 GB RAM)
  - 3 control plane nodes (custom machine type: 8 vCPU and 32 GB RAM)
  - 3 worker nodes (custom machine type: 4 vCPU and 16 GB RAM)

### 3.5.2. Storage

- Infrastructure volumes:
  - 300 GB SSD persistent disk (deleted on instance deletion)
  - 110 GB Standard persistent disk (kept on instance deletion)
- Worker volumes:
  - 300 GB SSD persistent disk (deleted on instance deletion)
- Control plane volumes:
  - 350 GB SSD persistent disk (deleted on instance deletion)

### 3.5.3. VPC

- **Subnets:** One master subnet for the control plane workloads and one worker subnet for all others.

- **Router tables:** One global route table per VPC.
- **Internet gateways:** One internet gateway per cluster.
- **NAT gateways:** One master NAT gateway and one worker NAT gateway per cluster.

### 3.5.4. Services

The following services must be enabled on a GCP CCS cluster:

- **deploymentmanager**
- **compute**
- **cloudapis**
- **cloudresourcemanager**
- **dns**
- **iamcredentials**
- **iam**
- **servicemanagement**
- **serviceusage**
- **storage-api**
- **storage-component**
- **orgpolicy**
- **networksecurity**

### 3.5.5. Permissions

The following roles must be added to the support service account:

- **compute.admin**
- **dns.admin**
- **orgpolicy.policyViewer**
- **servicemanagement.admin**
- **serviceusage.serviceUsageAdmin**
- **storage.admin**
- **compute.loadBalancerAdmin**
- **viewer**
- **iam.roleAdmin**

- `iam.securityAdmin`
- `iam.serviceAccountKeyAdmin`
- `iam.serviceAccountAdmin`
- `iam.serviceAccountUser`

### 3.6. GCP ACCOUNT LIMITS

The OpenShift Dedicated cluster uses a number of Google Cloud Platform (GCP) components, but the default [quotas](#) do not affect your ability to install an OpenShift Dedicated cluster.

A standard OpenShift Dedicated cluster uses the following resources. Note that some resources are required only during the bootstrap process and are removed after the cluster deploys.

**Table 3.5. GCP resources used in a default cluster**

Service	Component	Location	Total resources required	Resources removed after bootstrap
Service account	IAM	Global	5	0
Firewall Rules	Compute	Global	11	1
Forwarding Rules	Compute	Global	2	0
In-use global IP addresses	Compute	Global	4	1
Health checks	Compute	Global	3	0
Images	Compute	Global	1	0
Networks	Compute	Global	2	0
Static IP addresses	Compute	Region	4	1
Routers	Compute	Global	1	0
Routes	Compute	Global	2	0
Subnetworks	Compute	Global	2	0
Target Pools	Compute	Global	3	0
CPUs	Compute	Region	28	4

Service	Component	Location	Total resources required	Resources removed after bootstrap
Persistent Disk SSD (GB)	Compute	Region	896	128

**NOTE**

If any of the quotas are insufficient during installation, the installation program displays an error that states both which quota was exceeded and the region.

Be sure to consider your actual cluster size, planned cluster growth, and any usage from other clusters that are associated with your account. The CPU, Static IP addresses, and Persistent Disk SSD (Storage) quotas are the ones that are most likely to be insufficient.

If you plan to deploy your cluster in one of the following regions, you will exceed the maximum storage quota and are likely to exceed the CPU quota limit:

- asia-east2
- asia-northeast2
- asia-south1
- australia-southeast1
- europe-north1
- europe-west2
- europe-west3
- europe-west6
- northamerica-northeast1
- southamerica-east1
- us-west2

You can increase resource quotas from the [GCP console](#), but you might need to file a support ticket. Be sure to plan your cluster size early so that you can allow time to resolve the support ticket before you install your OpenShift Dedicated cluster.