

Self-managed Red Hat OpenShift subscription guide

Table of contents

Introduction	2
Red Hat OpenShift subscription offerings	2
Red Hat OpenShift Kubernetes Engine.....	3
Red Hat OpenShift Container Platform	3
Red Hat OpenShift Platform Plus.....	5
Self-managed OpenShift environments	6
Subscription types.....	7
Disaster recovery.....	7
Migration and swing upgrades.....	8
Entitlements for cores with hyperthreading	8
Entitlements for virtualized compute nodes.....	9
Core bands	9
Bare-metal server considerations.....	9
Alternative architectures (ARM, IBM zSystems, IBM® LinuxONE, and IBM Power)	10
Microsoft Windows Server containers support	11
Red Hat OpenShift Platform Plus component support.....	11
Determining how many Red Hat OpenShift subscriptions you need	13
Example entitlements for an initial self-managed Red Hat OpenShift deployment.....	17
How to estimate your entitlement requirements.....	17
Step 1: Determine standard VM or hardware cores and memory	17
Step 2: Calculate number of application instances needed	18
Step 3: Determine preferred maximum OpenShift node utilization	18
Step 4: Determine total memory footprint	18
Step 5: Calculate totals	19

Introduction

This document will help you understand the subscription model for self-managed [Red Hat® OpenShift®](#) offerings and provide step-by-step instructions for how to approximate the number of entitlements needed for an OpenShift environment. More accurate sizing information is available on request.

Red Hat OpenShift subscription offerings

Red Hat OpenShift provides a consistent application development and management platform across an open hybrid cloud environment, and supports on-premises, virtual, and physical infrastructure, and private cloud, public cloud, and edge deployments. There are 2 ways to operate and use Red Hat OpenShift: self-managed OpenShift and fully managed OpenShift cloud services.

Self-managed OpenShift allows you to install, operate, and manage Red Hat OpenShift environments with maximum control, flexibility, and customization, so you can operate your own environment starting with the infrastructure. Self-managed OpenShift is supported on-premises—using physical servers, virtualization, and in a private cloud environment—and in supported public cloud environments. You control upgrades, manage the lower-level infrastructure, and maintain service-level agreements (SLA).

OpenShift cloud services are fully managed and operated by Red Hat and its public cloud partners in major public clouds. A dedicated site reliability engineering (SRE) team manages and maintains Red Hat OpenShift core services and infrastructure, allowing your DevSecOps teams to concentrate on developing and deploying new applications and modernizing existing ones.

All editions of OpenShift offer a consistent user experience for developers and operations across every environment, allowing you to transfer your skills and applications to the cloud environments where your applications run best.

Self-managed OpenShift software offerings:

- ▶ **Red Hat OpenShift Kubernetes Engine:** A hybrid cloud, enterprise Kubernetes runtime engine that provides core OpenShift functionality to deploy and run applications, which you install and manage in datacenter, public cloud, or edge environments.
- ▶ **Red Hat OpenShift Container Platform:** A hybrid cloud, enterprise Kubernetes platform to build, deploy, and run applications, which you install and manage in datacenter, public cloud, and edge environments.
- ▶ **Red Hat OpenShift Platform Plus:** A single hybrid cloud platform that allows enterprises to build, deploy, run, and manage intelligent applications securely at scale across multiple clusters and cloud footprints. Multiple layers of security, manageability, and automation provide consistency throughout the software supply chain.

OpenShift cloud services offerings:

- ▶ **Red Hat OpenShift Dedicated:** Fully managed Red Hat OpenShift service on Amazon Web Services (AWS) and Google Cloud. [Read more information including pricing on OpenShift.com.](#)
- ▶ **Microsoft Azure Red Hat OpenShift:** Fully managed Red Hat OpenShift service on Microsoft Azure, jointly managed by Red Hat and Microsoft. [Read more.](#)

- ▶ **Red Hat OpenShift Service on AWS:** Fully managed Red Hat OpenShift service on Amazon Web Services, jointly managed by Red Hat and AWS. [Read more.](#)
- ▶ **Red Hat OpenShift Kubernetes Service on IBM Cloud:** Fully managed Red Hat OpenShift service on IBM Cloud, jointly managed by Red Hat and IBM. [Read more.](#)

Red Hat OpenShift Kubernetes Engine

Subscription components:

1. **Red Hat OpenShift Kubernetes Engine** is the Kubernetes runtime engine and infrastructure, and does not include the developer capabilities and advanced features of OpenShift Container Platform. OpenShift Kubernetes Engine includes the OpenShift Kubernetes distribution, Red Hat Enterprise Linux® and Red Hat Enterprise Linux CoreOS (described later in this section), and integrated Kubernetes cluster services components that include the OpenShift installer, monitoring, log forwarding, SDN, ingress router, registry, and more. See [About OpenShift Kubernetes Engine](#) in the OpenShift documentation for details.
2. **Red Hat Enterprise Linux and Red Hat Enterprise Linux CoreOS:** Each OpenShift subscription contains all the software needed for your compute nodes, control plane nodes, and supporting infrastructure nodes. This includes Red Hat Enterprise Linux CoreOS and Red Hat Enterprise Linux software. Red Hat Enterprise Linux CoreOS is required for the OpenShift control plane. Red Hat Enterprise Linux CoreOS is supported for use as a component of OpenShift. OpenShift customers can also choose to use Red Hat Enterprise Linux version 7 or 8 for their OpenShift compute nodes, as an alternative to Red Hat Enterprise Linux CoreOS. Red Hat Enterprise Linux must be installed separately by the customer on those compute nodes. Red Hat Enterprise Linux software is included in OpenShift subscriptions for this purpose. Refer to the [product documentation](#) to determine which version of Red Hat Enterprise Linux is supported with your OpenShift deployment.
3. **Red Hat OpenShift Virtualization:** Accelerate application delivery with a single platform that can manage virtual machines (VMs) and containers with the same tools and teams. Red Hat OpenShift Virtualization lets OpenShift manage and consume both containers and VMs with Kubernetes, using KubeVirt. This includes the entitlement to use Red Hat Enterprise Linux as the guest operating system in all the virtual machines hosted on OpenShift.
4. **Red Hat OpenShift administrator console:** Provides an optimized experience for administrators. The administrative perspective allows the user to view and manage the OpenShift and Kubernetes resources.
5. **Application streams:** OpenShift lets you use the container images provided in application streams (formerly [Software Collections](#)) that are included with Red Hat Enterprise Linux. These images include popular languages and runtimes—such as PHP, Python, Perl, Node.js, and Ruby—as well as databases, such as MySQL, MariaDB, and Redis. This offering also includes an OpenJDK image for Java™ frameworks. For more information, read [this post on application streams](#).

Red Hat OpenShift Container Platform

Subscription components:

1. **Red Hat OpenShift Kubernetes Engine:** Each OpenShift Container Platform subscription includes all of the components of OpenShift Kubernetes Engine as well as additional layered services described later in this section.

2. **Red Hat JBoss® Web Server:** OpenShift Container Platform subscriptions include Red Hat JBoss Web Server, an enterprise solution that combines the Apache web server with the Apache Tomcat servlet engine, supported by Red Hat. OpenShift Container Platform includes an unlimited right to use JBoss Web Server. [Learn more about JBoss Web Server.](#)
3. **Red Hat single sign-on (SSO) technology:** Red Hat provides web SSO and identity federation based on security assertion markup language (SAML) 2.0, OpenID Connect, and Open Authorization (OAuth) 2.0 specifications. This capability, included in OpenShift subscriptions, may only be deployed inside OpenShift environments. Nevertheless, any application—whether deployed inside or outside of OpenShift—may use Red Hat’s SSO.
4. **Red Hat build of Keycloak:** Red Hat introduces a new era of cloud-native identity and access management named Red Hat build of Keycloak, based on the Keycloak distribution powered by Quarkus. Red Hat build of Keycloak replaces any plans for future feature releases of Red Hat’s single sign-on (SSO) technology. While preserving the power and functionality of Red Hat’s (SSO) technology, Red Hat build of Keycloak is more efficient, more flexible, and optimized for a container-first approach. Red Hat build of Keycloak is included in Red Hat OpenShift subscriptions (Red Hat OpenShift Kubernetes Engine excluded), and may only be deployed inside Red Hat OpenShift environments when using the same subscriptions. A Red Hat build of Keycloak deployment running in Red Hat OpenShift may protect any application, whether deployed inside or outside of Red Hat OpenShift.
5. **Log management:** Adds support for log aggregation and management via Elasticsearch and Kibana integrated with Fluentd for log collection.
6. **Red Hat OpenShift Dev Spaces:** A collaborative Kubernetes-native development environment that delivers OpenShift workspaces and an in-browser integrated development environment (IDE).
7. **Red Hat build of Quarkus:** A full-stack, Kubernetes-native Java framework made for Java virtual machines (JVMs) and native compilation, optimizing Java specifically for containers and allowing it to become an effective platform for serverless, cloud, and Kubernetes environments.
8. **Web console:** Provides an optimized experience for both developers and administrators. The developer perspective grants visibility into application components, and the administrative perspective allows the user to view the OpenShift and Kubernetes resources.
9. **Red Hat OpenShift Pipelines:** Automate and control application delivery across on-premises and public cloud platforms with Kubernetes-native continuous integration/continuous delivery (CI/CD) pipelines based on Tekton.
10. **Red Hat OpenShift GitOps:** An opinionated workflow integrating git repositories, CI/CD tools, and Kubernetes to realize faster, security-focused, scalable software development, without compromising quality, based on Argo CD.
11. **Red Hat OpenShift Serverless:** Event-driven serverless containers and functions that let you deploy and run serverless containers. Powered by a rich ecosystem of event sources, you can manage serverless apps natively in OpenShift. Based on Knative, OpenShift Serverless allows you to run serverless applications anywhere OpenShift runs.

12. **Red Hat OpenShift Service Mesh:** Red Hat OpenShift Service Mesh provides a uniform way to connect, manage, and observe microservice-based applications, including Istio to manage and secure traffic flow across services, Jaeger for distributed tracing, and Kiali to view configuration and monitor traffic.
13. **Red Hat Insights for OpenShift:** Red Hat Insights for OpenShift is a set of hosted services on console.redhat.com, included with a Red Hat subscription, that use configuration and usage data sent from your deployments to console.redhat.com, along with rule-based analytical models, to help you track and optimize expenses, improve stability, and enhance performance.
14. **IBM Cloud Satellite:** Red Hat OpenShift Container Platform customers who choose to purchase and deploy the IBM Cloud Satellite solution can use their OpenShift node subscription to entitle the customer workload-related Red Hat OpenShift Kubernetes Service on IBM Cloud clusters located within their datacenter. Customers can call IBM or Red Hat for support, but ultimately the support experience will start with IBM Cloud Satellite support services. This OpenShift subscription usage is only available to customers deploying IBM Cloud Satellite within their datacenter and not in public cloud environments. Cores are counted the same way as explained elsewhere in this detail for normal OpenShift usage.
15. **Red Hat Support of Spring Boot:** Spring Boot is a popular framework for building stand-alone Java based applications. Red Hat provides developer and production support and guidance for successfully building and deploying Spring Boot workloads in OpenShift.

Red Hat OpenShift Platform Plus

Subscription components:

1. **Red Hat OpenShift Container Platform:** Each OpenShift Platform Plus subscription includes all of the components of OpenShift Container Platform, as well as additional layered products (listed later in this detail) to provide multicluster and hybrid cloud management and security at scale.
2. **Red Hat Advanced Cluster Management for Kubernetes:** Red Hat Advanced Cluster Management for Kubernetes offers end-to-end management visibility and control to manage your cluster and application life cycle, along with security and compliance of your entire OpenShift domain across multiple datacenters and public cloud environments.
3. **Red Hat Advanced Cluster Security for Kubernetes:** Red Hat Advanced Cluster Security for Kubernetes is the industry's first Kubernetes-native security platform that allows organizations to securely build, deploy, and run cloud-native applications anywhere. Red Hat Advanced Cluster Security for Kubernetes delivers lower operational cost, reduced operational risk, and greater developer productivity through a Kubernetes-native approach that supports built-in security across the entire software development life cycle.
4. **Red Hat Quay:** Red Hat Quay is a trusted open source registry platform for efficiently managing containerized content across global datacenters, focusing on cloud-native and DevSecOps development models and environments. With its tight integration into OpenShift and long track record of running Quay.io, one of the largest public-registry Software-as-a-Service (SaaS) in the world, Quay gives customers a reliable and scalable place to centrally manage all software artifacts running on their clusters.

5. **Red Hat OpenShift Data Foundation Essentials:** Red Hat OpenShift Data Foundation provides persistent software-defined file, block, and object storage and data services for applications running on OpenShift and OpenShift infrastructure services. It is integrated with and optimized for Red Hat OpenShift.

Self-managed OpenShift environments

Self-managed OpenShift (Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform, and Red Hat OpenShift Kubernetes Engine) can be used anywhere 64-bit Red Hat Enterprise Linux is certified and supported.

Red Hat OpenShift 4 supports 3 primary [deployment](#) methods:

- ▶ **Platform-specific installer-provisioned infrastructure (IPI).** Provides full integration, with underlying infrastructure platforms (listed later in this section), to automate the cluster provisioning and installation process. The installer provisions all resources necessary for cluster installation and configures integration between the OpenShift cluster and the infrastructure provider. OpenShift clusters deployed using the installer provisioned infrastructure (IPI) method via the command-line interface (CLI) or Red Hat Advanced Cluster Manager have full platform integration for supported infrastructure types.
- ▶ **Platform-specific user-provisioned infrastructure (UPI).** Depending on the infrastructure platform, a varying amount of integration between OpenShift and the underlying platform is available. The administrator provisions the resources necessary for cluster installation. Depending on the platform, the installer may configure infrastructure integration or the administrator may add integration post-deployment. OpenShift clusters deployed using the user provisioned infrastructure (UPI) method via the CLI or Assisted Installer for OpenShift have this level of integration. User provisioned infrastructure clusters are deployed using the CLI or Assisted Installer for OpenShift, with specific infrastructure platform integration levels available during and after deployment varying based on the provider type.
- ▶ **Platform-agnostic UPI or nonintegrated clusters.** This deployment type provides no integration with the underlying infrastructure. This install method offers the broadest range of compatibility, but the administrator is responsible for creating and managing cluster infrastructure resources. Nonintegrated clusters are deployed using the CLI or Assisted Installer for OpenShift to any hardware or hypervisor certified for Red Hat Enterprise Linux.
- ▶ **Hosted Control Planes:** A form factor of Red Hat OpenShift that decouples the control plane from the data plane (workers), separates management and workload network domains, and provides a shared interface through which administrators and site reliability engineers (SREs) can operate a fleet of clusters. OpenShift clusters with hosted control planes can either be deployed via the HyperShift CLI or the multicluster console. The multicluster console is accessible when either the multicluster engine for Kubernetes operator or Red Hat Advanced Cluster Manager is installed.

For self-managed deployments, OpenShift can be installed on:

- ▶ Bare-metal servers.
- ▶ Virtualized environments, including:
 - ▶ VMware vSphere.

- ▶ Red Hat Virtualization.
- ▶ Other certified virtualization platforms. Other [platforms](#) are supported via the platform-agnostic UPI install method.
- ▶ Private cloud environments.
 - ▶ Red Hat OpenStack® Platform and Azure Stack Hub.
- ▶ Public cloud environments, including:
 - ▶ Amazon Web Services, Azure, Google Cloud Platform, IBM Cloud, Alibaba Cloud, and VMware Cloud on AWS.
 - ▶ Other certified public cloud platforms. Other platforms are supported via the platform-agnostic UPI install method.

For more information about which platforms are supported, visit [the official OpenShift Container Platform documentation](#) page.

Registration for Red Hat Cloud Access is required to use your OpenShift subscriptions on certified public cloud environments. For more information, visit the [Red Hat Cloud Access](#) page.

Find out more about [platforms and cloud environments on which Red Hat OpenShift has been tested and certified](#).

Subscription types

Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform, and Red Hat OpenShift Kubernetes Engine subscriptions are available in 2 options, each with 2 support levels:

- ▶ Core-based (2 Cores or 4 vCPUs). This is based on the aggregate number of physical cores or virtual cores (vCPUs) across all the OpenShift worker nodes running across all OpenShift clusters. Available with Standard 8x5 or Premium 24x7 support SLA.
- ▶ Bare-metal socket pair (1-2 sockets with up to 64 cores). This subscription is available only for x86 bare-metal physical nodes where OpenShift is installed directly to the hardware, with the exception of IBM zSystems and IBM Power architectures, which must use core-based subscriptions.

As with Red Hat Enterprise Linux:

- ▶ OpenShift subscriptions (Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform, and Red Hat OpenShift Kubernetes Engine) are stackable to cover larger hosts.
- ▶ Core-based subscriptions can be distributed to cover all OpenShift worker nodes across all OpenShift clusters. For example, 100 2-core Red Hat OpenShift Platform Plus subscriptions will provide 200 cores (400 vCPUs) that can be used across any number of worker nodes, across all running OpenShift clusters.

Disaster recovery

Red Hat defines 3 types of disaster recovery (DR) environments—hot, warm, and cold. Paid OpenShift subscriptions are needed for hot DR only.

- ▶ Hot DR systems are defined as fully functional and running concurrently with the production systems. They are ready to immediately receive traffic and take over in the event of a disaster within the primary environment.
- ▶ Warm DR systems are defined as already prepared to deploy and host containerized workload representing a reasonable facsimile of that found in the primary site, but containing no customer workload from the source cluster(s).
- ▶ Cold DR systems are defined as having the infrastructure in place, but not the full technology (hardware, software, data) needed to restore service.

Hibernating clusters that are not expressly configured and designed for warm or cold DR—such as clusters running on cloud services that are temporarily hibernating owing to lower demand—require subscriptions. When warm or cold DR clusters are brought out of hibernation for running workloads, subscriptions are required. Bringing a cluster out of hibernation temporarily for routine maintenance or tests does not require an additional subscription for any of the components in OpenShift software offerings.

When setting up a warm DR cluster, add the infrastructure node label to it. This will prevent it from being counted against your total number of paid subscriptions.

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

Migration and swing upgrades

Red Hat OpenShift 4 provides in-place upgrades between minor versions. If you are upgrading from Red Hat OpenShift 3, or need to perform a swing upgrade between OpenShift 4 minor versions as a result of maintenance windows or other considerations, your Red Hat OpenShift subscription will cover both the original and destination infrastructure of a 1-way migration until such migration is complete. During the migration, Red Hat's subscription management tools will show your environment as being out-of-compliance on the basis of the number of OpenShift subscriptions you purchased. Red Hat allows this for major version upgrades and will not require the purchase of additional subscriptions to get back into compliance during the migration. Finally, OpenShift provides tooling to assist in these migrations, along with Red Hat consulting services, if desired. See documentation on the [migration toolkit for containers](#).

Entitlements for cores with hyperthreading

Making a determination about whether or not a particular OpenShift node uses 1 or more physical cores is determined by whether or not that system has hyperthreading enabled. Note that hyperthreading is only a feature of Intel central processing units (CPUs). [Discover how](#) to determine whether a particular system supports hyperthreading.

For systems where hyperthreading is enabled, and where 1 hyperthread equates to 1 schedulable system core, a calculation of cores at a ratio of 2 cores = 4 logical CPUs or vCPUs is used.

In other words, a 2-core subscription covers 4 logical CPUs in a hyperthreaded system. Red Hat's subscription management tools assume that hyperthreading is turned on for systems using Intel central processing units.

Entitlements for virtualized compute nodes

When deploying OpenShift compute nodes to a hypervisor, such as VMware vSphere or Red Hat OpenStack Platform, the number of entitlements required is the lesser of the cores/threads assigned to the virtual compute nodes or the sum of cores of the physical servers.

For example, an OpenShift cluster composed of virtual compute nodes totaling 200 cores would require 100 2-core/4-thread entitlements when the underlying physical hypervisor nodes have more than 200 cores. The same 200 core OpenShift cluster, when deployed to a physical hypervisor cluster totaling just 120 cores - resulting in overcommitment of OpenShift CPU resources at the hypervisor - would need only 120 cores of entitlements (60 2-core/4-thread).

Core bands

Red Hat OpenShift subscriptions use a system of measure called core bands. This means subscriptions (entitlements to deploy and use OpenShift) are applied and consumed at the OpenShift cluster level and apply to all eligible OpenShift compute nodes on that cluster. If you have multiple OpenShift clusters, you would aggregate the sum of cores consumed by the OpenShift compute nodes across all clusters to determine how many subscriptions are needed. For example, if you have 100 two-core Red Hat OpenShift Container Platform subscriptions, a total of 200 cores (400 vCPUs with hyperthreading) are available to be applied to the OpenShift compute nodes across all running OpenShift clusters.

Bare-metal server considerations

A physical server can be entitled using either core-based (2-core/4 logical CPU) or socket-based (1-2 socket, 64 cores) Red Hat OpenShift subscriptions. If core-based subscriptions are used, stack a sufficient number of them to cover the total number of physical cores in the server.

In addition to core-based subscriptions, Red Hat also offers OpenShift socket-based subscriptions. For certain deployment types, this is a more economical option. The socket-based subscriptions are limited to entitling an x86 server with up to 2 sockets with a total of 64 cores across them. The socket-based subscriptions are currently available for x86 and ARM servers only and not for the IBM zSystems and IBM Power architectures.

To entitle a physical server, stack 1 or more subscriptions to cover either the total number of sockets or physical cores in it (whichever is greater). For example, a server has 2 sockets and 48 cores. One subscription is needed because the server has 2 sockets and less than 64 cores, while a server with 2 sockets and 96 cores would need 2 subscriptions. Two subscriptions are needed to cover 96 cores because a single subscription covers a maximum of 64 cores.

Bare-metal socket-pair subscriptions also come with infrastructure subscriptions for the control plane and infrastructure. Qualified control plane and infrastructure workloads (explained in detail below) may be deployed to either, or a mix of, virtual and/or physical servers when using socket-based subscriptions. A mixed cluster, composed of virtual and physical nodes, is a supported deployment architecture when deploying a platform agnostic, non-integrated cluster without cloud provider or machine application programming interface (API) integrations.

Each physical, bare-metal server using socket-based entitlements can only be used as a single OpenShift node. Using a hypervisor, including OpenShift Virtualization, to create virtual OpenShift compute nodes will require entitling the virtual compute nodes using core-based subscriptions. This

means that the bare-metal socket-pair model is best suited for resource intensive workloads like OpenShift Virtualization (where each workload is running a full virtual machine) or artificial intelligence and machine learning (AI/ML) (where each workload consumes a large amount of CPU and graphics processing unit (GPU).

Finally, using the bare-metal socket-pair subscriptions does not change the limitation of the number of containers per node (currently 250-500).

Alternative architectures (ARM, IBM zSystems, IBM® LinuxONE , IBM Power)

Note: While this document refers only to IBM zSystems from here on, all information that references IBM zSystems also applies to IBM® LinuxONE.

Red Hat OpenShift Container Platform can also run on ARM, IBM zSystems, and IBM Power for customers using these platforms as the standard for building and deploying cloud-native applications and microservices. Only the core-based subscription model is supported for IBM zSystems and IBM Power platforms.

ARM clusters are entitled using the same rules as x86.

For IBM zSystems customers, Red Hat OpenShift does not require the entire physical node to be entitled, only the cores used by OpenShift. IBM zSystems customers know this as “subcapacity” entitlement. Customers using only a subset of the available cores (compute capacity) on their IBM zSystems environment for OpenShift Container Platform only require subscriptions for the subset that is used for the compute nodes. This applies regardless of how CPU partitioning is achieved, whether by CPU pooling, capping, separate logical partitions (LPARs), or other means.

For IBM zSystems, one Integrated Facility for Linux (IFL) requires one OpenShift core-based subscription. When no partitioning is used, up to 3 IFLs per cluster may be identified per OpenShift cluster for control plane or infrastructure services running on the host. These must be actively used for control plane and/or infra services to qualify and do not require OpenShift entitlements. 3-node compact cluster deployments require all IFLs to be entitled.

Red Hat OpenShift Platform Plus components beyond OpenShift Container Platform are not supported on IBM zSystems and IBM Power at this time, with the following exceptions:

- ▶ A standalone subscription of Red Hat Quay running on x86 architectures provides a global registry for multiple architectures, including IBM zSystems and IBM Power clusters. Red Hat Quay itself will not run on IBM zSystems or IBM Power.
- ▶ Red Hat Advanced Cluster Management for Kubernetes can be installed on IBM zSystems or IBM Power environments and manage Red Hat OpenShift nodes running on IBM zSystems or IBM Power environments.
- ▶ With Red Hat Advanced Cluster Security for Kubernetes, you can secure clusters running on Red Hat OpenShift on IBM zSystems or IBM Power by using the RHACS Operator.
- ▶ Red Hat OpenShift Data Foundation fully supports installation on IBM zSystems or IBM Power.

Red Hat OpenShift Kubernetes Engine is not supported on IBM zSystems and IBM Power.

Microsoft Windows Server containers support

Self-managed Red Hat OpenShift supports a subset of installation infrastructures and OpenShift features using Microsoft Windows Server containers. Windows Server containers are only supported on Microsoft Windows Server-based compute nodes. The control and infrastructure planes of the OpenShift environment must be running on x86 infrastructure using Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS. For this reason, Windows Server container support is sold as a standalone subscription priced by core.

Red Hat OpenShift Platform Plus and Red Hat OpenShift Container Platform infrastructure can be used to deploy and manage Windows Server compute nodes. Microsoft Windows Server container support for Red Hat OpenShift subscriptions must be purchased as a separate add-on.

Red Hat Advanced Cluster Management for Kubernetes and Red Hat Advanced Cluster Security for Kubernetes do not support managing Microsoft Windows nodes, but Red Hat Quay running on x86 architectures can manage container images for Microsoft Windows Server-based workloads.

Red Hat OpenShift Platform Plus component support

The components of the Red Hat OpenShift Platform Plus subscription have different levels of support for alternative (non-x86) architectures and for Windows. Table 1 provides an overview of that support.

Table 1: Overview of Red Hat OpenShift Platform Plus support

Red Hat OpenShift Platform Plus Component	IBM zSystems		IBM Power		Microsoft Windows		ARM	
	Installs on	Manages to	Installs on	Manages to	Installs on	Manages to	Installs on	Manages to
Red Hat OpenShift	Yes, infrastructure, control, and workers		Yes, infrastructure, control, and workers		Workers only (with separate subscription and Windows license)		Yes, infrastructure, control, and workers	
Red Hat Advanced Cluster Management for Kubernetes	Yes	Yes	Yes	Yes	No	No	Yes*	Yes*
Red Hat Advanced Cluster Security for Kubernetes	No	Yes	No	Yes	No	No	No	No
Red Hat OpenShift Data Foundation	Yes	Yes	Yes	Yes	No	No	No	No
Red Hat Quay	No	Yes	No	Yes	No	Yes	No	No

*ARM support is available in Advanced Cluster Management 2.5 or later versions.

** Advanced Cluster Security 3.74 or later versions support the securing of OpenShift on IBM zSystems and IBM Power.

For more details, see the compatibility matrices for [Red Hat OpenShift Container Platform](#), [Red Hat Advanced Cluster Management](#), [Red Hat Advanced Cluster Security](#), [Red Hat Quay](#) and [Red Hat OpenShift Data Foundation](#).

Red Hat OpenShift Platform Plus includes additional software beyond the core OpenShift Container Platform to help you manage and secure your OpenShift environment at scale across multiple clusters and multiple clouds. Red Hat OpenShift Platform Plus is available both in the core-based and bare-metal socket-pair subscription models with the previously mentioned limitations.

The additional software included with Red Hat OpenShift Platform Plus is generally limited to managing the nodes entitled with OpenShift Platform Plus subscriptions. For example, the subscription for Red Hat Advanced Cluster Management included with OpenShift Platform Plus can only be used to manage OpenShift Platform Plus entitled nodes and clusters. Customers who also wish to manage non-OpenShift Platform Plus entitled nodes and clusters, for example Red Hat OpenShift Services on AWS clusters, would need to purchase additional Red Hat Advanced Cluster Management add-on subscriptions to cover those clusters.

The additional software subscriptions are also inseparable from the OpenShift Platform Plus subscription. For example, you cannot purchase 100 OpenShift Platform Plus subscriptions, install 200 cores of Red Hat OpenShift Container Platform subscriptions, and separately use Red Hat Advanced Cluster Management to manage 200 cores of Azure Red Hat OpenShift with the same subscription. The additional software can only be used to manage the same 200 cores on which the core OpenShift Platform Plus software is installed.

Specific rules for each layered product are:

- ▶ **Red Hat Advanced Cluster Management for Kubernetes:** An OpenShift Platform Plus subscription allows you to install as many Red Hat Advanced Cluster Management central instances as needed to manage your environment, and covers the management of all nodes and clusters entitled with OpenShift Platform Plus, including control plane and infrastructure nodes. If you wish to manage nodes and clusters without OpenShift Platform Plus entitlements (for example, if you also have self-managed OpenShift Container Platform or Red Hat OpenShift Kubernetes Engine entitled clusters, clusters running in a fully managed OpenShift cloud, or third-party Kubernetes environments supported by Red Hat Advanced Cluster Management), then you need to purchase Red Hat Advanced Cluster Management add-on subscriptions to cover those environments. You can choose to manage them centrally from the Red Hat Advanced Cluster Management console installed on OpenShift Platform Plus, or from a separate central application if that meets your requirement. [Learn more about Red Hat Advanced Cluster Management subscriptions, Red Hat Advanced Cluster Management supported environments, and Red Hat Advanced Cluster Management best practices.](#)
- ▶ **Red Hat Advanced Cluster Security for Kubernetes:** The OpenShift Platform Plus subscription allows you to install as many Red Hat Advanced Cluster Security central applications as needed to manage your environment, and covers the management of all nodes and clusters entitled with OpenShift Platform Plus, including control plane and infrastructure nodes. If you want to manage nodes and clusters without OpenShift Platform Plus entitlements (for example, if you also have self-managed OpenShift Container Platform or OpenShift Kubernetes Engine entitled clusters, clusters running in a fully managed Red Hat OpenShift cloud, or third-party Kubernetes environments supported by Red Hat Advanced Cluster Security) you need to purchase Red Hat Advanced Cluster Security add-on subscriptions to cover those environments. Red Hat's

suggested practice is to manage each environment with a separate Red Hat Advanced Cluster Security central application. [Learn more about Red Hat Advanced Cluster Security supported environments.](#)

- ▶ **Red Hat Quay:** The OpenShift Platform Plus subscription allows you to install Red Hat Quay on any cluster that has an OpenShift Platform Plus entitlement. There is no limit on the number of Quay deployments you can install on your OpenShift Platform Plus clusters. Quay can then serve any supported Kubernetes environment you wish, including the OpenShift Platform Plus environment, other self-managed OpenShift clusters, managed OpenShift services, and supported third-party Kubernetes. If you wish to install Quay in a non-OpenShift Platform Plus environment, you will need to purchase a separate Red Hat Quay subscription. Red Hat Quay is also available as a fully managed [SaaS offering](#).
- ▶ **Red Hat OpenShift Data Foundation.** The OpenShift Platform Plus subscription allows you to install Red Hat OpenShift Data Foundation Essentials on any cluster that has an OpenShift Platform Plus entitlement. The Red Hat Data Foundation entitlement is limited to the features available in Essentials, and to 256TB of data storage per OpenShift cluster. You can choose to extend functionality and capacity through additional subscriptions. Please see the [OpenShift Data Foundation Subscription Guide](#) (customer portal login required) or consult with a Red Hat sales representative for further guidance.

Determining how many Red Hat OpenShift subscriptions you need

To conduct a more thorough sizing exercise to determine how many self-managed OpenShift (Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform, or Red Hat OpenShift Kubernetes Engine) or add-on subscriptions you need, use the following questions and examples.

A few basic OpenShift terms are used in these sizing exercises:

- ▶ **Pod:** The smallest deployable Kubernetes unit in OpenShift. A Kubernetes pod instance could have a single container or multiple containers running as sidecars.
- ▶ **Application instance:** An “application” may be a single-pod instance or may be deployed across multiple-pod instances that make up an application service. For example, a highly available Tomcat application service may consist of two or more Tomcat pods.
- ▶ **Worker node:** Instances (VMs or bare-metal hosts) of Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS where end-user application pods run. OpenShift environments can have many worker nodes.
- ▶ **Control plane nodes:** Instances (VMs or bare-metal hosts) of Red Hat Enterprise Linux CoreOS that act as the Kubernetes orchestration/management layer for OpenShift. Control plane nodes are included in self-managed OpenShift subscriptions. See the Red Hat OpenShift control plane and infrastructure nodes section for more details.
- ▶ **Infrastructure nodes:** Instances (virtual or physical hosts) of Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS that are running pods supporting OpenShift’s cluster infrastructure or running the HAProxy-based load balancer for ingress traffic. Infrastructure nodes are included in self-managed OpenShift subscriptions. See the Red Hat OpenShift control plane and infrastructure nodes section below for more details.
- ▶ **Cluster:** An OpenShift Kubernetes cluster consisting of a control plane and one or more worker nodes.

In summary:

- ▶ Applications are packaged in container images.
- ▶ Containers are deployed as pods.
- ▶ Pods run on Kubernetes worker nodes, which are managed by the Kubernetes control plane nodes.

Red Hat OpenShift control plane and infrastructure nodes

Each self-managed Red Hat OpenShift subscription includes entitlements for Red Hat OpenShift, Red Hat Enterprise Linux, and other OpenShift-related components. These entitlements are included for running OpenShift control plane and infrastructure workloads and do not need to be accounted for during sizing.

Red Hat OpenShift Platform Plus subscriptions include the management of these control plane and infrastructure nodes by Red Hat Advanced Cluster Security for Kubernetes and Red Hat Advanced Cluster Management for Kubernetes.

Infrastructure nodes

The OpenShift installer deploys a highly available OpenShift control plane composed of three control plane nodes, in addition to OpenShift worker nodes, to run end-user applications. By default, Kubernetes control plane components (e.g. API server, etcd, scheduler) and supporting cluster services (e.g. monitoring, registry) will be deployed on the OpenShift control plane nodes. However, you may decide to move some of these supporting cluster services to dedicated infrastructure nodes.

To qualify as an infrastructure node and use the included entitlement, only components that are supporting the cluster, and not part of an end-user application, may be running on those instances. Examples include:

- ▶ OpenShift registry.
- ▶ OpenShift Ingress Router (local and global and multicluster ingress).
- ▶ OpenShift monitoring.
- ▶ OpenShift log management.
- ▶ HAProxy-based instances used for cluster ingress.
- ▶ Red Hat Quay.
- ▶ Red Hat OpenShift Data Foundation (previously Red Hat OpenShift Container Storage).
- ▶ Red Hat Advanced Cluster Management for Kubernetes.
- ▶ Red Hat Advanced Cluster Security for Kubernetes.
- ▶ Red Hat OpenShift GitOps.
- ▶ Red Hat OpenShift Pipelines.
- ▶ Hosted control planes for Red Hat OpenShift.

You may deploy and run custom and third-party agents and tools for monitoring, log data collection and forwarding, hardware drivers, infrastructure integration such as virtualization agents, etc., to infrastructure nodes without disqualifying the node for infrastructure licensing. However, this

is limited only to agents and related components, including controller pods for operators and does not include the custom or third-party software suite. Examples of non-Red Hat software that qualify as infrastructure workload include:

- ▶ Custom and third-party monitoring agents.
- ▶ Container network interface (CNI) and container storage interface (CSI) drivers and controllers (plug-ins).
- ▶ Hardware or virtualization enablement accelerators.
- ▶ Controller pods used for Kubernetes CRD or Operators (custom or third-party software).

No other end-user application instances or types may be run on an infrastructure node using the included entitlement. To run other infrastructure workloads as application instances on Red Hat OpenShift, you must run those instances on regular application nodes. You can verify with Red Hat whether an app or service qualifies as an infrastructure workload.

Additional approved usage of the infrastructure node

As end users increase their usage of Red Hat OpenShift, they may begin using some of the more sophisticated application deployment patterns. As a result, they may need to add additional software components to the platform. As a general principle, Red Hat OpenShift subscriptions are based on the total capacity of the Red Hat OpenShift worker nodes that are required to run the application workloads and supporting application services deployed to those worker nodes. Red Hat OpenShift control plane nodes and components that are used to augment the capabilities of the platform, or its ability to deploy application workloads, can run on Red Hat OpenShift control plane nodes or additional infrastructure nodes that users may configure that do not require an entitlement. Where applicable, end users can use infrastructure nodes without disqualifying the node for infrastructure licensing to house these software components. Examples may include:

- ▶ CNI and CSI drivers and controllers (also known as plug-ins).
- ▶ Hardware or virtualization enablement accelerators (related to the Special Resource Operator or Node Feature Discovery operator).
- ▶ Cloud or virtualization agents.

Third-party management and monitoring products

Sometimes you may not want to use the Red Hat-provided monitoring and management features to manage Red Hat OpenShift, such as cluster monitoring, cluster logging, advanced cluster management, advanced cluster security. Or, you may want to augment these management features with additional solutions. In these instances, Red Hat allows the software components of those solutions (regardless of whether they are custom or purchased from a third-party vendor) to use the infrastructure label within Red Hat OpenShift so they do not incur the use of worker-node-cores counts for their framework's load. These software solutions can be related to several use cases from monitoring, alerting, security scanning, configuration management, and other Day 2 management tasks of Red Hat OpenShift. They must be exclusively used for the management and monitoring of Red Hat OpenShift and not end-user applications running on the platform.

No other end-user applications may be run on an infrastructure node that falls outside of the descriptions put forth in this section. If you need to, you can verify your software's infrastructure node status qualifications with [Red Hat Technical Support](#).

Control plane nodes

OpenShift Kubernetes control plane nodes generally are not used as worker nodes, and by default, will not run application instances. However, you may choose to use a control plane node as a node for hosting end-user applications. Whether a control plane node requires a full OpenShift subscription depends on whether it runs supporting OpenShift cluster components only or end-user applications. See the Infrastructure nodes section.

In a compact 3-node cluster, end-user application workloads are run on the control plane nodes. There is no special pricing for this, and you would count the cores on the 3 nodes regardless of the role they play.

Single node OpenShift

A single node OpenShift instance deploys all OpenShift services and end-user applications to a single physical or virtual node, with optimizations to reduce the footprint and maximize resources available to applications. As with compact 3-node clusters above, there is no special accommodations for this deployment model, all cores on the node need to be entitled.

Red Hat Enterprise Linux entitlements for supporting services

Red Hat Enterprise Linux entitlements for OpenShift compute nodes are included with the OpenShift entitlement. OpenShift subscriptions do not include other entitlements for Red Hat Enterprise Linux nodes with the following exception:

- ▶ A Red Hat Enterprise Linux node used specifically for bare-metal IPI provisioning.

Red Hat Enterprise Linux entitlements are not included for external nodes hosting services that OpenShift depends on, such as internet proxies, load balancers, or the mirror registry.

Bootstrap container registry for mirroring OpenShift container images

The mirror registry for OpenShift is a Quay entitlement for the single purpose of easing the process of mirroring content required for bootstrapping disconnected OpenShift clusters and is included as part of the OpenShift subscription. This is a limited support entitlement for a minimal Quay deployment created by a specific installer, which allows you to run a Quay registry on a preprovisioned and customer-managed Red Hat Enterprise Linux 8 host.

Note: You are permitted to use Quay as a registry mirror limited to the purpose of mirroring the OpenShift release payload, OperatorHub content, sample Operator images, and Cincinnati graph image.

The mirror registry for OpenShift is not intended to be a general-purpose registry working at arbitrary scale. Nonetheless, a limited set of custom images is permitted to be stored there, which contains required auxiliary software-like agents. These agents must be scoped only to the node level and not provide external-facing application services themselves, and end users may not interact with them directly. Examples of these may include:

- ▶ Monitoring agents.
- ▶ CNI and CSI providers.
- ▶ Hardware or virtualization enablement agents.
- ▶ Operators supporting independent software vendor (ISV) services.
- ▶ Custom operators as deployment controllers.

Example entitlements for an initial self-managed Red Hat OpenShift deployment

The following example bill of materials provides an extremely flexible, scalable Red Hat OpenShift environment designed to run as VMs and support hundreds of application containers:

- ▶ **16 x OpenShift Platform Plus, 2-Core Premium subscriptions**, including:
 - ▶ Control plane nodes (3 VMs).
 - ▶ Additional infrastructure nodes (3 VMs).
 - ▶ Worker nodes (16 VMs sized at 2 cores or 4 vCPUs).
 - ▶ Multicluster management, advanced observability, and policy compliance.
 - ▶ Declarative security and active threat detection and response.
 - ▶ Scalable global container registry.
 - ▶ Persistent storage for applications and Red Hat OpenShift infrastructure services.

Optional:

- ▶ **16 x Red Hat OpenShift Data Foundation Advanced**: Adds enhanced scalability, granular encryption, disaster recovery functionality, data security, and resilient block and file for file, block, and object storage services for workloads deployed inside Red Hat OpenShift as well as OpenShift infrastructure services. This is an optional add-on for customers running stateful applications that require persistent storage, or who want to build and operate a dedicated external storage cluster shared by multiple OpenShift clusters. Red Hat OpenShift Data Foundation Advanced is also available as part of a bundle called *Red Hat OpenShift Platform Plus with Red Hat OpenShift Data Foundation Advanced*.

Red Hat offers many additional application services and runtimes that have their own subscription and consumption models.

How to estimate your entitlement requirements

Red Hat OpenShift subscriptions do not limit application instances. You can run as many application instances in the Red Hat OpenShift environment as the underlying hardware and infrastructure will support. Larger-capacity hardware can run many application instances on a small number of hosts, while smaller-capacity hardware will require many hosts to run many application instances. The primary factor in determining the size of a Red Hat OpenShift environment is how many pods, or application instances, will be running at any given time.

Step 1: Determine standard VM or hardware cores and memory

You may have a standard VM size for application instances or, if you typically deploy on bare metal, a standard server configuration. The following questions will help you more accurately understand your VM and hardware needs. Remember that, in most cases, 2 vCPUs are equivalent to core.

Table 2: VM and hardware sizing questions

Relevant questions	Example answers
What is the memory capacity of the VMs you will use for nodes?	Our VMs have 64GB of memory and 4 vCPUs and hyperthreading is used.
What is the number of vCPUs for the VMs you will use for nodes?	
Is hyperthreading in use?	

Step 2: Calculate number of application instances needed

Next, determine how many application instances, or pods, you plan to deploy. When sizing the environment, any application component deployed on Red Hat OpenShift—such as a database, front-end static server, or message broker instance—is considered an application instance.

This figure can be an approximation to help you calculate a gross estimate of your Red Hat OpenShift environment size. CPU, memory oversubscription, quotas and limits, and other features can be used to further refine this estimate.

Table 3: Application and instance sizing questions

Relevant questions	Example answers
How many application instances do you anticipate deploying in each Red Hat OpenShift environment?	We have around 1,250 application instances in our development environment and around 250 application instances in production.
What type of applications are they (e.g., language, framework, database)?	We mainly deploy Java but have some Microsoft .NET Core and Ruby applications as well. We also use a lot of MySQL.

Step 3: Determine preferred maximum OpenShift node utilization

We recommend reserving some space in case of increased demand, especially if autoscaling is enabled for workloads. Your preferred utilization will vary on the basis of historical load for the applications that will run on OpenShift.

Table 4: Preferred maximum OpenShift node utilization questions

Relevant questions	Example answers
How much space do I want to reserve for increased demand?	We want to run nodes at a maximum average of 80% of total capacity (leaving 20% in reserve).

Step 4: Determine total memory footprint

Next, calculate the total memory footprint of the deployed applications. If you are considering a completely greenfield environment, memory use data may not be available, but you can use educated approximations—for example, 1GB of memory per Java application instance—to make an estimate.

Table 5: OpenShift memory footprint questions

Relevant questions	Example answers
What is the average memory footprint of applications?	Our application instances use 2GB of memory or less. OR We typically allocate 2GB for JVM heap.

Step 5: Calculate totals

Finally, determine the number of OpenShift subscriptions needed on the basis of the data gathered in steps 1-4.

- ▶ Effective per-node memory capacity (GB)
= Preferred maximum OpenShift node utilization (%) * standard VM or hardware memory
- ▶ Total memory utilization
= Application instances * average application memory footprint
- ▶ Number of nodes required to cover utilization
= Total memory utilization / standard VM or hardware memory
- ▶ Total required cores
= Number of nodes required to cover utilization * standard VM or hardware cores
- ▶ Effective virtual cores
= Total required cores / 2
- ▶ Number of OpenShift Platform Plus subscriptions¹
= Total cores / 2 OR
= Effective virtual cores / 2

Example calculation for virtualized environments

System sizing (from steps 1-5)

- ▶ Standard number of VM cores = 4 (hyperthreading used, 2 effective virtual cores)
- ▶ Standard VM memory = 64GB
- ▶ Preferred maximum node utilization = 80%
- ▶ Average application memory footprint = 2GB
- ▶ Number of application instances = 1500

Subscription calculations

- ▶ Effective node memory capacity
= 80% preferred maximum node utilization * 64GB standard VM memory
= 51GB

¹ If hyperthreading is in use, 2 virtual cores count only as 1 core of a subscription. See the section on Cores versus vCPUs and hyperthreading for details on whether to use effective or actual cores in this calculation.

- ▶ Total memory utilization
 - = 1500 application instances * 2GB average application memory footprint
 - = 3000GB
- ▶ Nodes required to cover utilization
 - = 3000GB total memory utilization / 51GB effective node memory capacity
 - = 59 nodes
- ▶ Total cores
 - = 59 nodes required * 2 cores per node
 - = 118 total cores
- ▶ Total subscriptions
 - = 118 total cores / 2 cores per subscription
 - = **59 subscriptions**

In this example, 59 2-core OpenShift Platform Plus 2-core subscriptions would be needed.

Notes: Red Hat OpenShift supports many features and functions which affect scalability, Pod scheduling, idling, and resource quota/limiting features. The previous calculations are guidelines, and you may be able to tune your actual environment for better resource use or smaller total environment size. OpenShift Platform Plus customers should take into account the needs of the additional software applications (Red Hat Advanced Cluster Management, Red Hat Advanced Cluster Security, and Quay) including storage and compute resources, even though they may not require additional worker subscriptions.

If working with a third-party reseller, please refer to their specific terms and agreements for Red Hat products and services.



About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

f facebook.com/redhatinc
t @RedHat
in linkedin.com/company/red-hat

North America
1 888 REDHAT1

**Europe, Middle East,
and Africa**
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com

redhat.com
#639118_1223

Copyright © 2023 Red Hat, Inc. Red Hat, the Red Hat logo, JBoss, and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission. Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle America, Inc. in the U.S. and other countries. All other trademarks are the property of their respective owners.