

Self-managed Red Hat OpenShift sizing and 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	4
Self-managed OpenShift environments	5
Subscription types	6
Disaster recovery	6
Migration and swing upgrades	7
Cores versus vCPUs and hyperthreading	7
Core Bands	7
Bare metal server considerations	7
Alternative architectures (IBM Z, Power)	8
Microsoft Windows Server containers support	9
Red Hat OpenShift Platform Plus component support	9
Example initial self-managed Red Hat OpenShift deployment	11
Self-managed Red Hat OpenShift sizing	11
Sizing process	14
Step 1: Determine standard VM or hardware cores and memory	14
Step 2: Calculate number of application instances needed	15
Step 3: Determine preferred maximum OpenShift node utilization	15
Step 4: Determine total memory footprint	15
Step 5: Calculate totals	16



Introduction

This document will help you understand the subscription model for self-managed Red Hat® OpenShift® offerings and provide easy-to-follow, step-by-step instructions for how to approximate the size of 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, private cloud, public cloud, and edge deployments. There are two ways to operate and consume Red Hat OpenShift: self-managed OpenShift and fully managed OpenShift cloud services.

Self-managed OpenShift allows the customer to install, operate, and manage Red Hat OpenShift environments with maximum control, flexibility, and customization, operating their own environment starting with the infrastructure. Self-managed OpenShift is supported on-premises on physical servers, virtualization, and private cloud, as well as in supported public clouds. The customer controls upgrades, manages the lower level infrastructure, and maintains service-level agreements (SLA).

Fully **managed OpenShift** cloud services are 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 a customer's DevSecOps team 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 footprint, allowing customers to transfer their skills and applications to the clouds where their 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, that is installed and managed by customers in datacenter, public cloud, or edge environments.
- **Red Hat OpenShift Container Platform:** A hybrid cloud, enterprise Kubernetes platform to build, deploy, and run applications that is installed and managed by customers in datacenter, public cloud, or 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.

Fully managed OpenShift cloud service offerings:

- **Red Hat OpenShift Dedicated:** Fully-managed Red Hat OpenShift service on Amazon Web Services (AWS) and Google Cloud. [More info 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. [More info here.](#)
- **Red Hat OpenShift Service on AWS:** Fully-managed Red Hat OpenShift service on Amazon Web Services, jointly managed by Red Hat and AWS. [More info here.](#)

- **Red Hat OpenShift Kubernetes Service on IBM Cloud:** Fully-managed Red Hat OpenShift service on IBM Cloud, jointly managed by Red Hat and IBM. [More info here.](#)

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 below), 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 worker nodes, control plane nodes, and supporting infrastructure nodes. This includes the 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 for their OpenShift worker nodes, as an alternative to Red Hat Enterprise Linux CoreOS. Red Hat Enterprise Linux version 7 must be installed separately by the customer on those worker nodes. The Red Hat Enterprise Linux software is included in OpenShift subscriptions for this purpose.

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 described above) as well as additional layered services described below.
2. **Red Hat Software Collections:** OpenShift lets you use the container images provided in Red Hat Software Collections. 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 the [Red Hat Software Collections technology brief](#).
3. **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.](#)
4. **Red Hat single sign-on (SSO):** 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. However, any application—whether deployed inside or outside of OpenShift—may use Red Hat’s SSO.
5. **Log management:** Adds support for log aggregation and management via Elasticsearch and Kibana integrated with Fluentd for log collection.

6. **Red Hat CodeReady Workspaces:** 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. **Red Hat OpenShift Virtualization:** Accelerate application delivery with a single platform that can manage VMs and containers with the same tools and teams. OpenShift Virtualization lets OpenShift manage and consume both containers and VMs with Kubernetes, using KubeVirt.
9. **Red Hat OpenShift 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.
10. **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.
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 cloud.redhat.com, included with a Red Hat subscription, that use configuration and utilization data sent from customer deployments to cloud.redhat.com, along with rule-based analytical models, to help customers track and optimize spend, 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 eligible to customers deploying IBM Cloud Satellite within their datacenter and not on public clouds. Cores are counted the same way explained in this document for normal OpenShift usage.

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 (listed above) as well as additional layered products listed below 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 clouds.
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 lifecycle.
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 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.

Self-managed OpenShift environments

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

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

- Platform specific installer-provisioned infrastructure (IPI). Provides full integration, with underlying infrastructure platforms listed below, 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.
- 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.
- Platform agnostic UPI. 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.

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 clouds.
 - Red Hat OpenStack® Platform.
- Public clouds, including:
 - Amazon Web Services, Azure, Google Cloud Platform.
 - 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 clouds. For more information, visit the [Red Hat Cloud Access](#) page.

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

Subscription types

Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform and Red Hat OpenShift Kubernetes Engine subscriptions are available in two options each with two 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 is for x86 bare metal physical servers (IBM Z and Power not supported) with no virtualization installed, primarily for limited use cases like OpenShift Virtualization and workloads like artificial intelligence/machine learning (AI/ML) that benefit from direct hardware access without a virtualization layer. Available with Standard 8x5 or Premium 24x7 support SLA.

As with Red Hat Enterprise Linux:

- OpenShift subscriptions (Red Hat OpenShift Platform Plus, Red Hat OpenShift Container Platform, or Red Hat OpenShift Kubernetes Engine) are stackable to cover larger hosts.
- Core-based subscriptions can be distributed across 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 three 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 to 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 stocked with hardware representing a reasonable facsimile of that found in the primary site, but containing no customer data. To restore service, the last backups from the off-site storage facility must be delivered and bare metal must be restored before recovery can begin.

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

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

Migration and swing upgrades

Red Hat OpenShift 4 provides in-place upgrades between minor versions. However, for customers who are upgrading from Red Hat OpenShift 3 or need to perform a swing upgrade due to maintenance windows or other considerations, your Red Hat OpenShift subscription will cover both the original and destination infrastructure of a one-way migration until such migration is complete. During the migration, Red Hat's subscription management tools will show your environment as being out-of-compliance based on 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](#).

Cores versus vCPUs and hyperthreading

Making a determination about whether or not a particular system uses one or more cores is currently dependent on whether or not that system has hyperthreading available. Note that hyperthreading is only a feature of Intel CPUs; to determine whether a particular system supports hyperthreading, visit <https://access.redhat.com/solutions/7714>.

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

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

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

Core Bands

Red Hat OpenShift subscriptions use a system of measure called Core Bands. That means subscriptions (entitlements to consume OpenShift) are applied and consumed at the OpenShift cluster level and apply to all OpenShift worker nodes on that cluster. If a customer has multiple OpenShift clusters, they would aggregate the sum of cores consumed by the OpenShift worker nodes across all clusters to determine how many subscriptions are needed. For example, if a customer has 100 2-core Red Hat OpenShift Container Platform subscriptions, a total of 200 cores (400 vCPUs) are available to be applied to the OpenShift worker nodes across all running OpenShift clusters.

Bare metal server considerations

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

In addition to core-based subscriptions, Red Hat offers OpenShift socket-based subscriptions as well. 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. Today, the socket-based subscriptions are available for x86 servers only and not for the IBM Z or Power architectures.

To entitle a physical server, stack one 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 two 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. The best practice would be for these to be bare metal physical servers that take advantage of our infrastructure automation which works in homogenous clusters. You have the option to use a separate virtualized server as the control plane (creating a “mixed” cluster), but if you do so, you must use the more manual agnostic installer method, and will not be able to take advantage of the infrastructure automation that is possible with an all-bare metal cluster (i.e., all worker and infrastructure nodes bare metal).

Finally, use of the bare metal socket-pair subscriptions does not change the limitation of the number of containers per node (currently 250–500). Each physical bare metal server is a single node, and cannot use virtualization to carve it into a larger number of smaller nodes (for this you would need to use the core-based subscriptions). This means that the bare metal socket-pair model is best suited for a smaller number of “fatter” workloads like OpenShift Virtualization (where each workload is running a full virtual machine) or AI/ML (where each workload consumes a large amount of CPU and GPU).

Alternative architectures (IBM Z, Power)

Red Hat OpenShift Container Platform and OpenShift Platform Plus can also run on IBM Z and IBM LinuxONE systems and on IBM Power Systems 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 Z and IBM Power.

For IBM Z customers, Red Hat OpenShift does not require the entire physical node to be entitled, only the cores used by OpenShift. IBM Z customers know this as “sub-capacity” entitlement. Customers using only a subset of the available cores (compute capacity) on their IBM Z 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. In short, one Integrated Facility for Linux (IFL) requires one OpenShift core-based subscription.

Red Hat OpenShift Platform Plus components beyond OpenShift Container Platform are not supported on IBM Z and IBM LinuxONE systems or on IBM Power Systems 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 Z and IBM Power clusters. Red Hat Quay itself will not run on IBM Z or Power.
- A standalone subscription of Red Hat Advanced Cluster Management for Kubernetes running on x86 architecture can manage IBM Z or IBM Power environments. Red Hat Advanced Cluster Management itself cannot run on an IBM Z or IBM Power environment.

Red Hat OpenShift Kubernetes Engine is not supported on IBM Z and IBM LinuxONE systems, nor on IBM Power Systems.

Microsoft Windows Server containers support

Self-managed Red Hat OpenShift supports Windows containers. Support is limited to a supported subset of installation infrastructures and OpenShift features. Only Windows container worker nodes are supported. The control and infrastructure planes of the OpenShift environment must be running on x86 infrastructure. For this reason, Windows 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 worker workers, but the subscriptions must be purchased separately.

Red Hat Advanced Cluster Management for Kubernetes and Red Hat Advanced Cluster Security for Kubernetes are not supported for managing Windows nodes, but Red Hat Quay running on x86 architectures can manage container images for Windows container worker nodes.

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 System Z		IBM Power		Windows containers	
	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)	
Red Hat Advanced Cluster Management for Kubernetes	No	Yes	No	Yes	No	No
Red Hat Advanced Cluster Security for Kubernetes	No	No	No	No	No	No
Red Hat Quay	No	Yes	No	Yes	No	Yes

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

Red Hat OpenShift Platform Plus includes additional software beyond the core OpenShift Container Platform to help the customer manage and secure their 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 limitations listed above.

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 be used to manage any OpenShift Platform Plus nodes and clusters. If a customer wishes to also manage some Red Hat OpenShift Service on AWS clusters, they would need to purchase additional Red Hat Advanced Cluster Management add-on licenses to cover those clusters.

The additional software subscriptions are also inseparable from the OpenShift Platform Plus subscription. For example, the customer 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:** OpenShift Platform Plus subscription allows the customer to install as many Red Hat Advanced Cluster Management central instances as needed to manage their environment, and covers the management of all nodes and clusters entitled with OpenShift Platform Plus, including control plane and infrastructure nodes. If the customer wants to manage nodes and clusters without OpenShift Platform Plus entitlements (for example, the customer also has 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 the customer needs to purchase Red Hat Advanced Cluster Management add-on subscriptions to cover those environments. They 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 their requirement. More information on Red Hat Advanced Cluster Management licenses, Red Hat Advanced Cluster Management supported environments, and Red Hat Advanced Cluster Management best practices can be found on redhat.com.
- **Red Hat Advanced Cluster Security for Kubernetes:** The OpenShift Platform Plus subscription allows the customer to install as many Red Hat Advanced Cluster Security central applications as needed to manage their environment, and covers the management of all nodes and clusters entitled with OpenShift Platform Plus, including control plane and infrastructure nodes. If the customer wants to manage nodes and clusters without OpenShift Platform Plus entitlements (for example, the customer also has 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), then the customer needs 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. More information on [Red Hat Advanced Cluster Security supported environments](#) can be found on redhat.com.

- **Red Hat Quay:** The OpenShift Platform Plus subscription allows the customer to install Red Hat Quay on any cluster that has a OpenShift Platform Plus entitlement. There is no limit on the number of Quay deployments you can install to your Red Hat OpenShift Platform Plus clusters. Quay then can serve any supported Kubernetes environment the customer wishes, including the OpenShift Platform Plus environment, other self-managed OpenShift clusters, managed OpenShift services, and supported third-party Kubernetes. If the customer wishes to install Quay in a non-OpenShift Platform Plus environment, they will need to purchase a separate Red Hat Quay subscription. Red Hat Quay is also available as a fully managed SaaS offering at quay.io.

Example 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 virtual machines 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.
- **16 x Red Hat OpenShift Data Foundation:** Adds scalable block and file storage for applications inside Red Hat OpenShift. This is an optional add-on for customers running stateful applications that require storage.

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

Self-managed Red Hat OpenShift sizing

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 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 provides extra entitlements for Red Hat OpenShift, Red Hat Enterprise Linux, and other OpenShift-related components. These extra entitlements are included for the purpose of running OpenShift control plane and infrastructure nodes.

In the case of Red Hat OpenShift Platform Plus subscriptions, this includes 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, customers 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/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.

Customers 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 (plugins).
- 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. Verify infrastructure status qualifications with Red Hat.

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 plugins).
- 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 a customer does 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, the customer may want to augment these management

features with additional solutions. In these instances, Red Hat allows the software components of those solutions (regardless if 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 a number of 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, please verify your software's infrastructure node status qualifications with Red Hat [via support services](#).

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 three nodes regardless of the role they play.

Red Hat Quay bootstrap registry for mirroring OpenShift container images

Included as part of Red Hat OpenShift subscriptions is a Quay entitlement for the single purpose of mirroring content required for bootstrapping disconnected OpenShift clusters, called OpenShift Mirror Registry. This is a limited support entitlement for Quay that would allow customers to bootstrap and operate a Quay registry for a preprovisioned and customer-managed Red Hat Enterprise Linux 8 host.

Supported to be used only as part of the documented OpenShift Mirror Registry installer, customers are permitted to use Quay as a registry mirror, limited to mirroring the OpenShift release payload, OperatorHub content, sample operator images, and Cincinnati graph image.

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

- Monitoring agents.
- CNI/CSI providers.
- Hardware or virtualization enablement agents.
- Operators supporting ISV services.
- Custom Operators as deployment controllers.

Sizing process

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 an 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, two vCPUs is equivalent to one 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: VM and hardware 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, based on 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, 1 GB 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 2 GB of memory or less. OR We typically allocate 2 GB for JVM heap.

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.

Step 5: Calculate totals

Finally, determine the number of OpenShift subscriptions needed based on 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 above)

- 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
- 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 scalability, overcommitment, idling, and resource quota/limiting features. The calculations above 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.



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

redhat.com
#F28271_0521

North America
1 888 REDHAT1
www.redhat.com

**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

Copyright © 2021 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. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the OpenStack community. Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle America, Inc. in the U.S. and other countries.