

# Red Hat OpenShift sizing and subscription guide

## Table of contents

<b>Introduction</b>	2
<b>Red Hat OpenShift Container Platform</b>	2
Subscription components	2
Subscription types	3
Disaster recovery	4
OpenShift Container Platform environments (x86)	4
Cores vs vCPUs and hyperthreading	4
Splitting cores	5
<b>Red Hat OpenShift Dedicated</b>	5
<b>Suggested initial OpenShift deployment</b>	6
<b>Red Hat OpenShift Container Platform sizing</b>	6
Infrastructure nodes and supervisors	6
Cores and vCPUs	7
Sizing process	7
Step 1: Determine standard VM or hardware cores and memory	8
Step 2: Calculate number of application instances needed	8
Step 3: Determine preferred maximum OpenShift node utilization	8
Step 4: Determine total memory footprint	9
Step 5: Calculate totals	9



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

## Introduction

This document will help you understand the subscription model for 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 Container Platform:** OpenShift environment that is implemented and maintained by the customer.
- ▶ **Red Hat OpenShift Dedicated:** Highly available private OpenShift clusters hosted by Red Hat. Customers work with Red Hat to determine requirements and integrations, and Red Hat implements and fully manages the environment.
- ▶ **Microsoft Azure Red Hat OpenShift:** Flexible, fully managed Red Hat OpenShift service on Microsoft Azure.
- ▶ **Managed Red Hat OpenShift on AWS:** Amazon Red Hat OpenShift is a jointly managed and jointly supported enterprise Kubernetes service by Red Hat and AWS.
- ▶ **Red Hat OpenShift Kubernetes Service on IBM Cloud:** This OpenShift Service is a managed service supported by IBM running on IBM Cloud.
- ▶ **Red Hat OpenShift Kubernetes Engine:** OpenShift environment implemented and maintained by the customer, without support for advanced networking, management, and DevOps features.

## Red Hat OpenShift Container Platform

### Subscription components

1. **Red Hat Enterprise Linux® CoreOS:** Each OpenShift subscription includes support for Red Hat Enterprise Linux CoreOS. Red Hat Enterprise Linux CoreOS is only supported for use as a component of Red Hat OpenShift Container Platform.
2. **Red Hat OpenShift Container Platform:** Each subscription includes entitlements for OpenShift and its integrated components, including the following integrated solutions:
  - ▶ **Log aggregation:** Aggregates container logs and platform logs using Elasticsearch, Fluentd, and Kibana
  - ▶ **Metrics aggregation:** Aggregates container performance metrics such as memory use, CPU use, and network throughput using Prometheus and Grafana

These solutions are supported only in their native integrations with OpenShift, with limited support for customization; they are not supported for general use outside of OpenShift.

3. **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, MongoDB, and Redis. This offering also includes an OpenJDK image for Java™ frameworks, such as Spring Boot. For more information, read the [Red Hat Software Collections technology brief](#).

4. **Red Hat JBoss® Web Server:** OpenShift 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 includes an unlimited right to use JBoss Web Server. Learn more at [redhat.com/en/technologies/jboss-middleware/web-server](https://redhat.com/en/technologies/jboss-middleware/web-server).
5. **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.
6. **Red Hat CodeReady Workspaces:** A collaborative Kubernetes-native development solution that delivers OpenShift workspaces and an in-browser integrated development environment (IDE).
7. **Quarkus:** A full-stack, Kubernetes-native Java framework made for Java virtual machines (JVMs) and native compilation, optimizing Java specifically for containers and enabling 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. Add VMs to new and existing applications. Modernize legacy VM applications over time or maintain them as VMs.
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 drill down to the OpenShift and Kubernetes resources.
10. **Red Hat OpenShift Pipelines:** Automate and control application delivery across on-premise and public cloud platforms with Kubernetes-native CI/CD pipelines. No more CI server maintenance overhead.
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.

### Subscription types

Red Hat OpenShift Container Platform 2-core subscription is based on the number of logical cores on the CPUs in the system where OpenShift runs.

As with Red Hat Enterprise Linux:

- ▶ OpenShift Container Platform subscriptions are stackable to cover larger hosts.
- ▶ Cores can be distributed across as many VMs as needed. For example, 10 2-core subscriptions will provide 20 cores that can be used across any number of virtual machines (VMs).

OpenShift Container Platform subscriptions are available with Premium or Standard support.

### **Disaster recovery**

Red Hat OpenShift does not offer disaster recovery (DR), cold backup, or other subscription types. Any system with OpenShift installed, powered-on or powered-off, running workload or not, requires an active subscription. See the sections titled “Infrastructure nodes and supervisors” to understand more about subscription requirements.

### **OpenShift Container Platform environments (x86)**

OpenShift Container Platform can be used anywhere that 64-bit x86 Red Hat Enterprise Linux is certified and supported.

For on-premise deployments, OpenShift can be installed on:

- ▶ Bare metal
- ▶ Virtualized environments, including:
  - ▶ VMware
  - ▶ Red Hat Virtualization
  - ▶ Other virtualized platforms—other platforms are supported via the Platform Agnostic UPI install method.
- ▶ Private clouds
  - ▶ Red Hat OpenStack® Platform

OpenShift can also be installed and used on any public cloud that supports Red Hat Enterprise Linux. OpenShift cloud installations come with full integration with the underlying cloud platform. Installations not needing this integration can use the Platform Agnostic UPI install method. For more information about which clouds 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 [redhat.com/en/technologies/cloud-computing/cloud-access](https://access.redhat.com/en/technologies/cloud-computing/cloud-access).

For more information on platforms and clouds that Red Hat OpenShift has been tested and certified on, refer to OpenShift Container Platform Tested Integrations at <https://access.redhat.com/articles/2176281>.

### **Cores vs 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, then 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.

### Splitting cores

Systems that require an odd number of cores will need to use a full 2-core subscription. For example, a system that is calculated to require only 1 core will end up using a full 2-core subscription once it is registered and subscribed.

When a single VM with 2 vCPUs uses hyperthreading (see prior section), resulting in 1 calculated vCPU, a full 2-core subscription is required; a single 2-core subscription may not be split across two VMs with 2 vCPUs using hyperthreading.

We recommend that virtual instances be sized so that they require an even number of cores.

### Alternative architectures (IBM Z. Power)

Red Hat OpenShift is also available to 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.

### Red Hat OpenShift Dedicated

Red Hat OpenShift Dedicated provides single-tenant, high-availability, and fully managed OpenShift clusters on Amazon Web Services (AWS) and Google Cloud. OpenShift Dedicated clusters are managed by a global Red Hat OpenShift site reliability engineering (SRE) team and include a 99.9% uptime SLA with Premium Support.

Similar to OpenShift Container Platform, the right type and number of application nodes for OpenShift Dedicated depends on the anticipated resource needs of applications that will run on the platform—their memory footprint and CPU load—and the total number of application instances. However, OpenShift Dedicated can easily scale to accommodate more nodes.

For the Red Hat compute nodes, the customer will pay Red Hat. For CCS clusters, the customer will pay the cloud provider directly for infrastructure used to run OpenShift Dedicated.

**Table 1: Red Hat OpenShift Dedicated package overview**

	<b>Red Hat account</b> Deploy in cloud provider accounts owned by Red Hat	<b>Customer account (CCS)</b> Use your existing cloud provider infrastructure
Single-availability-zone cluster	Minimum of 4 compute nodes in Red Hat account	Minimum of 2 compute nodes in customer account
Multiple-availability-zone cluster	Minimum of 9 compute nodes in Red Hat account	Minimum of 3 compute nodes in customer account
Choice of general purpose, memory-optimized, and compute-optimized instances	Yes	Yes
Choice of application node sizes	Yes	Yes

## Suggested initial OpenShift deployment

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

- ▶ **16 x OpenShift Container Platform, 2-Core Premium subscriptions**, including:
  - ▶ Supervisors (3 VMs)
  - ▶ Redundant infrastructure nodes (3 VMs)
  - ▶ Application nodes (16 VMs)
- ▶ **18 x Red Hat OpenShift Container Storage**: Adds scalable block and file storage for applications inside OpenShift.

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

## Red Hat OpenShift Container Platform sizing

To conduct a more thorough sizing exercise to determine how many OpenShift Container Platform 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 deployed unit in OpenShift. A running instance of an application—for example, an app server or database.
- ▶ **Application instance**: Effectively the same as pod and used interchangeably.
- ▶ **Node**: Instances of Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS where pods run. OpenShift environments can have many nodes.
- ▶ **Supervisors (supervisor nodes)**: Instances of Red Hat Enterprise Linux CoreOS that act as the orchestration/management layer for OpenShift. Supervisors are included in OpenShift Container Platform subscriptions. See the “Infrastructure nodes and supervisors” section for more details.
- ▶ **Infrastructure nodes**: Instances of Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS that are running pods supporting OpenShift’s infrastructure. Infrastructure nodes are included in OpenShift Container Platform subscriptions. See the “Infrastructure nodes and supervisors” section for more details.
- ▶ **Cluster**: A group of OpenShift supervisors and nodes.

In summary:

- ▶ Applications are packaged in container images.
- ▶ Containers are grouped in pods.
- ▶ Pods run on nodes, which are managed by supervisors.

## Infrastructure nodes and supervisors

Each OpenShift Container Platform subscription provides extra entitlements for OpenShift, Red Hat Enterprise Linux, and other OpenShift-related components. These extra entitlements are included for the purpose of running either OpenShift Container Platform infrastructure nodes or supervisors.

### Infrastructure nodes

To qualify as an infrastructure node and use the included entitlement, only these included OpenShift components may be run as application instances:

- ▶ Red Hat OpenShift-included registry
- ▶ Router
- ▶ OpenShift cluster monitoring
- ▶ OpenShift log aggregation
- ▶ Red Hat Quay
- ▶ Red Hat OpenShift Container Storage
- ▶ Red Hat Advanced Cluster Manager

In addition, customers are permitted to deploy and run node-level monitoring, node enablement or provider enablement agents, on supervisor and infrastructure nodes. These agents must be scoped only to the node level and not provide external-facing services themselves, and that end users do not interact with them directly. Examples of these may include:

- ▶ Monitoring agents
- ▶ CNI/CSI providers
- ▶ Hardware or virtualization enablement agents

No other application instances or types may be run on an infrastructure node using the included entitlement. To run other infrastructure workloads as application instances on OpenShift, you must run those instances on regular application nodes. Verify infrastructure status qualifications with Red Hat.

### Supervisors

Supervisors generally are not used as nodes and, by default, will not run application instances. However, you could use a supervisor as a functional node. Whether a supervisor requires a full OpenShift Container Platform subscription depends on the application instances it runs. See the “Infrastructure nodes” section above.

In a Compact 3-Node cluster, Worker workloads are run on the supervisors. There is no special pricing for this and you would count the cores on the 3 nodes regardless of the role they play.

### Cores and vCPUs

Because of the way that Red Hat Enterprise Linux recognizes CPUs—and due to how modern CPUs work—it often appears that there are twice as many CPUs present. Because of this effect and how virtualization works, Red Hat implements a 2:1 mapping of subscription cores to vCPUs.

In the case of a VM—whether in a public cloud, private cloud, or local virtualized environment—one subscription core would cover 2 vCPUs. In other words, if a VM has 4 vCPUs assigned, a 2-core subscription would be required.

### Sizing process

OpenShift subscriptions do not limit application instances. You can run as many application instances in the 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, 2 vCPUs is equivalent to 1 core.

**Table 2: VM and hardware sizing questions**

Relevant questions	Example answer
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 OpenShift—such as a database, front-end static server, or message broker instance—is considered an application instance.

This figure can simply be an approximation to help you calculate a gross estimate of your 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 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 question	Example answer
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 question	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 based on the data gathered in steps 1-5.

- ▶ 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 Container Platform subscriptions<sup>1</sup>  
= Total cores / 2 OR  
= Effective virtual cores / 2

<sup>1</sup> If hyperthreading is in use, two virtual cores count only as one core of a subscription. See the section, “Cores versus vCPUs and hyperthreading” for details on whether to use effective or actual cores in this calculation.

## Example calculation for virtualized environments

System sizing (from steps 1-6 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 Container Platform subscriptions would be needed.

**Note:** 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 and/or smaller total environment size.



### 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 integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.



facebook.com/redhatinc  
@RedHat  
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  
#F26601\_1220

Copyright © 2020 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.