

# RED HAT OPENSIFT SUBSCRIPTION AND SIZING GUIDE

## TABLE OF CONTENTS

<b>INTRODUCTION</b> .....	<b>2</b>
<b>RED HAT OPENSIFT SUBSCRIPTION OFFERINGS</b> .....	<b>2</b>
<b>RED HAT OPENSIFT CONTAINER PLATFORM</b> .....	<b>2</b>
Subscription components .....	2
<b>SUBSCRIPTION TYPES</b> .....	<b>3</b>
Cores versus sockets .....	4
Cores versus vCPUs .....	4
OpenShift Container Platform environments .....	5
<b>OPENSIFT ADD-ON</b> .....	<b>5</b>
<b>RED HAT CONSULTING OPENSIFT OFFERING</b> .....	<b>6</b>
<b>RED HAT TRAINING OPENSIFT OFFERINGS</b> .....	<b>6</b>
<b>RED HAT OPENSIFT DEDICATED</b> .....	<b>7</b>
OpenShift Dedicated Base Package .....	7
Additional resources and products .....	7
<b>SUGGESTED INITIAL OPENSIFT DEPLOYMENT</b> .....	<b>8</b>
<b>RED HAT OPENSIFT CONTAINER PLATFORM SIZING</b> .....	<b>8</b>
Infrastructure nodes and masters .....	9
Cores and vCPUs .....	9
Sizing process .....	10
Step 1: Determine the number of OpenShift clusters needed .....	10
Step 2: Determine the number of application instances needed .....	10
Step 3: Determine the total memory footprint .....	11
Step 4: Calculate overhead .....	11
Step 5: Determine VM or hardware size .....	12
Step 6: Calculate totals .....	12
<b>RED HAT JBOSS MIDDLEWARE INTEGRATION</b> .....	<b>13</b>
Portability .....	13
Partial coverage .....	13
Developer access .....	13



## INTRODUCTION

This document is designed to help you understand the subscription model for Red Hat® OpenShift Container Platform 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 OPENSIFT SUBSCRIPTION OFFERINGS

 <b>RED HAT</b> OPENSIFT Online	Multitenant OpenShift environment hosted by Red Hat
 <b>RED HAT</b> OPENSIFT Dedicated	Single-tenant OpenShift environment hosted by Red Hat. Customers work with Red Hat to determine requirements and integrations, and Red Hat implements and maintains the environment via OpenShift Operations.
 <b>RED HAT</b> OPENSIFT Container Platform	OpenShift environment that is implemented and maintained by the customer
 <b>RED HAT</b> OPENSIFT Container Lab	Sales promotional program for first-time OpenShift Container Platform customers for nonproduction workloads
<b>RED HAT</b> CONTAINER DEVELOPMENT KIT	No-cost, local development virtual machine (VM) offered as part of the Red Hat Developer Program. Includes a full installation of OpenShift Container Platform.
<b>RED HAT</b> CLOUD SUITE	Comprehensive integrated solution that provides a complete cloud infrastructure, including Red Hat OpenShift Container Platform. For more information, visit <a href="https://redhat.com/en/technologies/cloud-computing/cloud-suite">redhat.com/en/technologies/cloud-computing/cloud-suite</a> .

## RED HAT OPENSIFT CONTAINER PLATFORM

### SUBSCRIPTION COMPONENTS

- 1. Red Hat Enterprise Linux®/Red Hat Enterprise Linux Atomic Host**  
Each OpenShift subscription includes Red Hat Enterprise Linux/Red Hat Enterprise Linux Atomic Host entitlements.<sup>1</sup>
- 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—memory use, CPU use, network throughput—using Heapster, Cassandra, and Hawkular

---

<sup>1</sup> Customers with existing, excess entitlements may want to consider Red Hat OpenShift Container Platform for Red Hat Enterprise Linux as an alternative. This offering, available at a slightly lower cost, does not include Red Hat Enterprise Linux/Red Hat Enterprise Linux Atomic Host entitlements.

These solutions are supported only in their native integrations with OpenShift, with limited support for customization. These solutions are not supported for general use outside of OpenShift. If you wish to use them outside of OpenShift, several third-party providers support these open source projects.

### 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 technology brief at [redhat.com/en/resources/red-hat-software-collections](https://redhat.com/en/resources/red-hat-software-collections).

### 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 Mobile Application Platform

Red Hat Mobile Application Platform provides capabilities for developing, integrating, and deploying enterprise mobile applications—whether native, web-based, or hybrid. This offering is included in OpenShift subscriptions, but the included entitlements are not supported for production use.

### 7. Red Hat CloudForms

Red Hat CloudForms provides capacity, trending, showback, security enforcement, and other capabilities for OpenShift. CloudForms entitlements are included in OpenShift subscriptions. However, the right to use is limited to management of and visibility into the OpenShift environment and its underlying hosts/infrastructure.

For more details on the limited acceptable uses of the included CloudForms, refer to the OpenShift Included Red Hat CloudForms FAQ.

## SUBSCRIPTION TYPES

There are two base subscription types for OpenShift Container Platform:

#### 1. OpenShift Container Platform, 2 Core

This subscription option is based on the number of logical cores on the CPUs on the system where OpenShift runs.

#### 2. OpenShift Container Platform, 1-2 Socket

This socket-based subscription option covers physical hosts and/or all VMs on that host. This offering is identical in behavior and coverage to the Red Hat Enterprise Linux socket-based

subscriptions. In other words, a socket-based subscription covers the operating system and/or hypervisor—e.g., Red Hat Virtualization, VMware—running on bare metal, as well as the VMs running on it. The socket-based subscription is available with or without Red Hat Enterprise Linux entitlements.

As with Red Hat Enterprise Linux:

- OpenShift Container Platform subscriptions are stackable to cover larger hosts or VMs with more sockets.
- 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 VMs.

Both subscription types are available with Premium or Standard support.

### CORES VERSUS SOCKETS

The decision on whether to use core-based or socket-based subscriptions depends on the density of OpenShift within your infrastructure.

Socket-based subscriptions follow the scalable virtualization model of Red Hat Enterprise Linux, where all VM guests are covered on a host with a socket-based subscription.

- **High-density, virtualized OpenShift environments:** In an environment with a large number of OpenShift instances compared to the total number of virtual hosts, socket-based OpenShift subscriptions can cover all OpenShift guests on those hosts.
- **Low-density virtualized OpenShift environments:** In an environment with a small number of OpenShift instances compared to the total number of virtual hosts, core-based OpenShift subscriptions can cover only the guests running OpenShift.
- **Public cloud environments:** In certified public clouds, only core-based subscriptions can be used to cover virtual OpenShift instances.
- **For bare-metal environments:** Only socket-based subscriptions can be used for running OpenShift directly on bare metal.

### SPLITTING CORES

Systems that require an odd number of cores will need to consume a full 2-core subscription. For example, a system that is calculated to require only 1 core will end up consuming a full 2-core subscription once it is registered and subscribed. It is recommended that virtual instances be sized so that they require an even number of cores.

### CORES VERSUS VCPUS AND HYPERTHREADING

Making a determination about whether or not a particular system consumes 1 or more cores is currently dependent on whether or not that system has hyperthreading available.

For systems where hyperthreading is enabled and where 1 hyperthread equates to 1 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. For example, 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.

In other words, a 2-core subscription covers 2 vCPUs in a non-hyperthreaded system. For example, a VM might have 4 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 4 vCPUs.

## **OPENSIFT CONTAINER PLATFORM ENVIRONMENTS**

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

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

- Bare metal.
- Virtualized environments, including:
  - VMware.
  - Microsoft Hyper-V.
  - Red Hat Virtualization.
- Private clouds, including Red Hat OpenStack® Platform.

OpenShift can also be installed and used on any Red Hat Enterprise Linux-certified public cloud, such as:

- Amazon Web Services (AWS).
- Google Cloud Platform (GCP).
- Microsoft Azure.

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/cloud-access).

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

## **OPENSIFT ADD-ONS**

Many of Red Hat's middleware solutions are available as add-ons for OpenShift. Some Red Hat storage and management offerings may also be purchased for use with OpenShift. These solutions include:

- Middleware:
  - Red Hat JBoss Enterprise Application Platform (EAP)
  - Red Hat JBoss Data Grid
  - Red Hat JBoss Data Virtualization

- Red Hat JBoss AMQ
- Red Hat Decision Manager
- Red Hat JBoss BPM Suite
- Red Hat JBoss Fuse
- Red Hat Mobile Application Platform<sup>2</sup>
- Red Hat 3scale API Management<sup>2</sup>
- Storage:
  - Container-native storage via Red Hat Gluster Storage<sup>2</sup>
- Management:
  - Red Hat CloudForms<sup>3</sup>

## RED HAT CONSULTING OPENSIFT OFFERINGS

Red Hat has an extensive catalog of consulting offerings, ranging from initial software deployment to complete programs designed to support your organization's digital transformation journey.

- **Discovery Session:** No-cost Discovery Sessions create open discussion on how an organization can adopt OpenShift, take advantage of DevOps practices, and establish innovative application development. Discovery Sessions are designed to produce a summary report presenting the focus areas discussed, scope and priority information, and a proposed approach for continuing to work with Red Hat Consulting to increase adoption.
- **OpenShift Container Platform Pilot:** This eight-week, extended introductory services engagement puts organizations on a path to modernizing application delivery through efficient use of container and container orchestration technologies.
- **Container Adoption Program:** This comprehensive program supports enterprise-scale container adoption, including establishing container platform infrastructure, deployment pipeline automation, and mass migration of applications to containers. The program begins with a pilot offering and incorporates Red Hat Open Innovation Labs to support faster innovation and time to market.
- **Red Hat Open Innovation Labs:** In this immersive experience program, customers visit Red Hat's labs to experiment, become more agile, learn DevOps practices, and catalyze innovation alongside Red Hat experts. Learn more at [redhat.com/en/open-innovation-labs](https://redhat.com/en/open-innovation-labs).

## RED HAT TRAINING OPENSIFT OFFERINGS

Red Hat Training courses can help you and your staff quickly become familiar with OpenShift and related technologies. Red Hat offers a variety of ways to access training—from no-cost online webinars available on demand to scheduled training held on-premise at your location.

### Paid, instructor-led training:

- Introduction to containers, Kubernetes, and Red Hat OpenShift (D0180)

---

<sup>2</sup> Not currently available for use with OpenShift Dedicated

<sup>3</sup> Currently available as a Technology Preview running inside a container on OpenShift

- Red Hat OpenShift Administration I (DO280)
- Developing and deploying applications on OpenShift (DO290)

**No-cost, self-paced training:**

- Deploying containerized applications technical overview (DO080)
- Red Hat OpenShift Interactive Learning Portal: <https://learn.openshift.com>

**RED HAT OPENSIFT DEDICATED**

Red Hat OpenShift Dedicated provides one or more single-tenant, high-availability OpenShift clusters delivered as a public cloud service. OpenShift Dedicated clusters are managed by Red Hat OpenShift operations experts, who use years of experience to increase security while providing a streamlined path to use OpenShift in public cloud environments.

Similar to OpenShift Container Platform, the right number of application nodes for OpenShift Dedicated depends on the size of applications—their memory footprint—and the total number of application instances. Note that application node size and capacity are fixed. See the sizing section of this document for examples.

**OPENSIFT DEDICATED BASE PACKAGE**

The OpenShift Dedicated Base Package is the smallest environment available. It provides an entire OpenShift cluster that is implemented and maintained by Red Hat—and, optionally, securely connected to your internal network. There is a choice of public cloud providers and hosting regions. Currently, AWS and GCP are supported.

**TABLE 1. OPENSIFT DEDICATED BASE PACKAGE OVERVIEW**

ITEM	AMOUNT/TYPE	DESCRIPTION
Base platform	--	Includes 3 x multimaster high availability (HA) VMs and 2 x infrastructure nodes for OpenShift's containerized components
Application nodes	4	Each application node includes 4 vCPUs and 16GB of memory
Network I/O	48TB	Each application node includes 12TB of network I/O per year. 48TB (4 x 12TB) is included.
Persistent storage	100GB	High speed solid-state drive (SSD) persistent storage for application use
Support	Premium	Red Hat Premium Support, which includes 24x7 coverage

**ADDITIONAL RESOURCES AND PRODUCTS**

Application nodes, persistent storage, and network I/O can be added to OpenShift Dedicated to expand the capacity of the cluster as needed. Other Red Hat middleware products can also be added to an OpenShift Dedicated environment (Table 2).

**TABLE 2. RED HAT OPENSIFT DEDICATED ADD-ONS**

ITEM	AMOUNT	DESCRIPTION
Application node	1	Adds memory and compute resources
Network I/O	12TB	Additional network I/O for data transfer
Persistent storage	500GB	Additional persistent storage
Red Hat JBoss EAP	1 year	JBoss EAP subscription covering 1 node with Premium support
Red Hat JBoss Fuse	1 year	JBoss Fuse subscription covering 1 node with Premium support
Red Hat JBoss AMQ	1 year	JBoss AMQ subscription covering 1 node with Premium support
Red Hat Decision Manager	1 year	Decision Manager subscription covering 1 node with Premium support
Red Hat JBoss BPM Suite	1 year	JBoss BPM Suite subscription covering 1 node with Premium support
Red Hat JBoss Data Grid	1 year	JBoss Data Grid subscription covering 1 node with Premium support

### SUGGESTED INITIAL OPENSIFT DEPLOYMENT

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

- **16 x OpenShift Container Platform, 2-Core Premium subscriptions**, including:
  - Multimaster HA (3 VMs).
  - Redundant infrastructure nodes (2 VMs).
  - Application nodes (16 VMs).
- **2 x Container-Native Storage for OpenShift Container Platform add-on**: Adds scaleable block and file storage for applications inside OpenShift.
- **16 x Red Hat JBoss Enterprise Application Platform for OpenShift Container Platform, 2-Core Premium service**: Adds support for Java EE applications on OpenShift.

### RED HAT OPENSIFT 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 Atomic Host where pods run. OpenShift environments can have many nodes.

- **Masters:** Instances of Red Hat Enterprise Linux or Red Hat Enterprise Linux Atomic Host that act as the orchestration/management layer for OpenShift. Masters are included in OpenShift Container Platform subscriptions. See the “Infrastructure nodes and masters” section for more details.
- **Infrastructure nodes:** Instances of Red Hat Enterprise Linux or Red Hat Enterprise Linux Atomic Host that are running pods supporting OpenShift’s infrastructure. Infrastructure nodes are included in OpenShift Container Platform subscriptions. See the “Infrastructure nodes and masters” section for more details.
- **Cluster:** A group of OpenShift masters and nodes.

In summary:

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

## INFRASTRUCTURE NODES AND MASTERS

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 masters or infrastructure nodes.

### INFRASTRUCTURE NODES

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

- Registry
- Router
- Metrics aggregation
- Logging aggregation
- Red Hat CloudForms

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 an OpenShift node, you must run those instances on regular application nodes.

### MASTERS

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

### 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—1 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 THE NUMBER OF OPENSIFT CLUSTERS NEEDED

First, understand how many separate, distinct clusters of OpenShift you need. For example, to physically separate development and production environments, you would need two separate OpenShift clusters. If all of your software development life cycle (SDLC) environments can exist together in the same OpenShift, you will only need one cluster.

The OpenShift scheduler is extremely flexible. For example, you can configure OpenShift to ensure that workloads for each of your SDLC environments land on specific hosts.

**TABLE 3. OPENSIFT CLUSTER SIZING QUESTIONS**

RELEVANT QUESTIONS	EXAMPLE ANSWERS
What kind of life-cycle environments do you plan to use OpenShift in (e.g., development, testing, staging, production)?	We plan to use OpenShift across our entire pipeline.
Are your environments physically segregated today? For example, do you have different VM management systems for development and production? Do they need to stay segregated given OpenShift's flexibility?	Our SDLC environments are not managed by different VM management systems and do not require separation via multiple OpenShift clusters.

#### STEP 2: DETERMINE THE 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 4. OPENSIFT APPLICATION INSTANCE ESTIMATE 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 THE 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. OPENSIFT MEMORY FOOTPRINT QUESTIONS**

RELEVANT QUESTIONS	EXAMPLE ANSWERS
What is the maximum memory footprint of each application type?	All application instances use 2GB of memory or less, according to our monitoring software
What is the application memory limit—for example, heap size in Java—or VM memory setting (if applications were previously deployed inside VMs)?	We typically allocate 16GB of RAM to each VM running our application.

**STEP 4: CALCULATE OVERHEAD**

If you intend to install monitoring, security, or other software onto the operating system where OpenShift runs, these components may use up some of the available system memory. Additionally, OpenShift node processes and software use some memory—approximately 0.5GB.

Simply calculate the total overhead of components installed on a host. This figure will be subtracted later to estimate the total capacity of the environment.

**TABLE 6. OPENSIFT OVERHEAD QUESTIONS**

RELEVANT QUESTIONS	EXAMPLE ANSWERS
What monitoring or management solutions are you using? How much memory overhead will you need?	Our host applications use 0.5GB of memory, and OpenShift uses 0.5GB, for a total of 1GB of memory overhead.
Is there any additional software you will need to run on each of the nodes?	(Answers will vary.)

**STEP 5: DETERMINE VM OR HARDWARE SIZE**

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.

**TABLE 7. VM AND HARDWARE SIZING QUESTIONS**

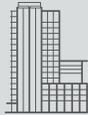
RELEVANT QUESTIONS	EXAMPLE ANSWERS
What is the memory capacity of the VMs you will use for nodes?  What is the number of vCPUs for the VMs you will use for nodes?	Our VMs have 64GB of memory and 4 vCPUs.
What is the memory capacity of the physical hardware you will use for nodes?  How many sockets do your physical servers have?	Our physical servers have 56GB of memory and 2 sockets.

**STEP 6: CALCULATE TOTALS**

Finally, determine the number of OpenShift subscriptions needed based on the data gathered in steps 1- 5.

**Review:**

- **Number of clusters** = The total amount of separate clusters you plan to deploy
- **Memory footprint** = Number of application instances \* Memory of applications
- **Number of nodes** = Memory Footprint / (VM or hardware memory - Overhead)
- **Virtual: Number of cores** = Number of nodes \* Number of standard vCPUs
- **Physical: Number of sockets** = Number of nodes \* Number of standard sockets
- **Number of OpenShift Container Platform subscriptions** = Number of cores or sockets / 2



### ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.

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



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

redhat.com  
#f1189\_0218

### Example calculation for virtualized environments:

- Number of clusters: 1 OpenShift cluster with 4 life-cycle environments
- Memory footprint:
  - Development memory footprint: 1,250 instances \* 2GB = 2,500GB
  - Production memory footprint: 250 instances \* 2GB = 500GB
  - Total memory footprint: 500GB + 2,500GB = 3,000GB
- VM Size: 64GB
- VM vCPUs: 4
- Overhead: 1GB
- Number of nodes:  $3,000\text{GB} / (64\text{GB} - 1\text{GB}) = 47.6$  (rounds to 48 nodes)
- Number of vCPUs:  $48 \times 4 = 192$
- Number of cores:  $192 / 2 = 96$
- Number of node subscriptions:  $96 / 2 = 48$

In this example, 48 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.

## RED HAT JBOSS MIDDLEWARE INTEGRATION

### PORTABILITY

Red Hat JBoss Middleware subscriptions not specifically designed for OpenShift may be used in OpenShift environments or across OpenShift and non-OpenShift environments.

If you have existing JBoss Middleware subscriptions, you can use them in OpenShift without having to purchase OpenShift-specific subscriptions. If you purchase OpenShift-specific JBoss Middleware subscriptions, you may use them outside of OpenShift. All cases are valid with correct core or socket accounting.

For more details on JBoss Middleware portability, refer to Section 1.5.2 of the Red Hat Global Subscription Services Appendix, Appendix 1 of the Enterprise Agreement, available at [redhat.com/en/about/licenses](https://redhat.com/en/about/licenses).

### PARTIAL COVERAGE

You are not required to provide middleware subscription coverage for an entire OpenShift Container Platform environment. Platform administrators can work to ensure that the middleware workload runs only on the correct, subscribed hosts inside your OpenShift environment.

### DEVELOPER ACCESS

OpenShift does not have a developer access model for JBoss Middleware products. Customers who wish to use Red Hat JBoss Middleware products on OpenShift must purchase subscriptions, regardless of whether the use is for development, test, or production environments.

Copyright © 2018 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack® Word Mark and OpenStack Logo are either registered trademarks / service marks or trademarks / service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Java is a trademark of Oracle America, Inc.