WHITE PAPER

Red Hat® OpenShift® on Supermicro® SuperBlade®

Supermicro® SuperBlade® Simplifies and Accelerates Your Private Cloud Deployment Journey

TABLE OF CONTENTS

SUPERMICRO

Supermicro (Nasdaq: SMCI) is the leading innovator in high-performance, high-efficiency server and storage technology and a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its "We Keep IT Green®" initiative by providing customers with the most energy-efficient, environmentally friendly solutions available on the market.

Supermicro SuperBlade with AMD EPYC™ processors powers hybrid cloud deployments

Executive Summary

Organizations must deliver applications, features, and services at the speed their businesses require across increasingly diverse IT environments as technological initiatives such as cloud migration, containerization, and microservices take hold. These technologies offer the promise of more efficient and effective IT services, but it can be challenging for organizations to realize these benefits without proper technological platforms and solutions. Building innovative applications requires cloud-native development to deliver great results with agility. Red Hat customers leverage Red Hat® OpenShift® as a unified development platform to streamline their application development efforts across their heterogeneous IT environments. The Red Hat® OpenShift® platform empowers organizations to deliver more applications and features in less time while providing the functionality to their customers and business lines. These organizations thus derive significant additional value from the higher productivity of their DevOps and application development teams by generating more revenue and making their application development efforts more cost effective.

Server hardware innovations play a crucial role. This paper guides you through deploying and validating basic Red Hat® OpenShift® Container Platform 4.8 functionality on Supermicro® SuperBlade® bare metal systems.

Red Hat® OpenShift® Market Momentum:

Red Hat OpenShift Container Platform (OCP) provides enterprise Kubernetes bundled CI/CD pipelines. It automates builds and deployments, enabling developers to focus on application logic while leveraging best-of-class enterprise infrastructure. Forrester reports that the Red Hat OpenShift Container Platform leads the growing enterprise App Container market.

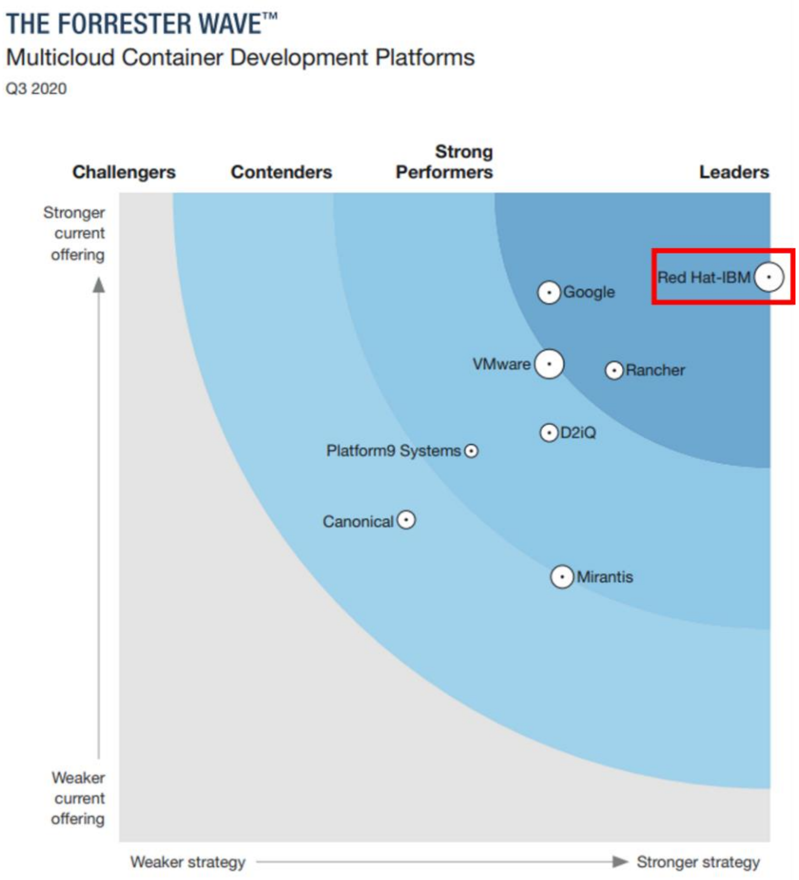

Figure 1: The Forrester Wave

Forrester [1] says that "Red Hat OpenShift is the most widely deployed multi-cloud container platform and boasts powerful development and unified operations experiences across many public and on-premises platforms. Red Hat pioneered the 'operator' model for infrastructure and application management and provides a rich partner ecosystem and popular marketplace."

1 - The Forrester Wave™: Multicloud Container Development Platforms, Q3 2020: The Eight Providers That Matter Most and How They Stack Up Sep 2020

Applications are the competitive differentiator in today's market. Deploying cloud-native applications represents an opportunity to build recurring revenues, high margin services, and platform renewals. According to Gartner [2], **"By 2025, more than 85% of global organizations will be running containerized applications, a significant increase from fewer than 35% in 2019."**

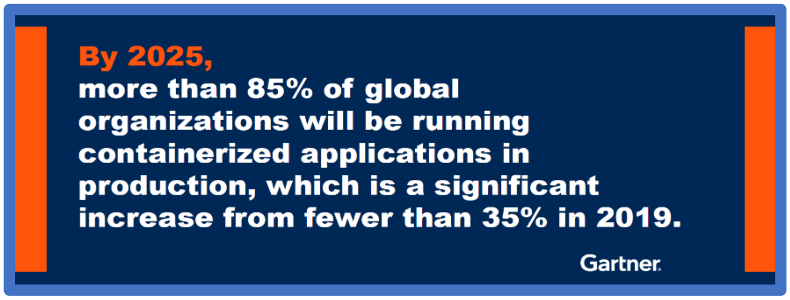

Red Hat OpenShift Container Platform

Red Hat OpenShift Container Platform is an open-source container application platform built primarily on Docker containers and orchestrated using Kubernetes container cluster Red Hat OpenShift management. Red Hat OpenShift supports a broad range of programming languages and services ranging from web frameworks, databases, or connectors to mobile devices and external back ends. The Red Hat OpenShift platform supports both cloud-native, stateless applications and traditional, stateful applications.

The primary IT initiatives that Red Hat OpenShift targets are:

- **Accelerating application delivery with agile and DevOps methodologies**: Red Hat OpenShift offers a common platform for development and operations teams to ensure consistency and standardization of application components, eliminate configuration errors, automate deployment, and enable a controlled rollout of new capabilities into production and rollback in the event of a failure. Additional capabilities enforce policies and role-based access control for environments with elevated security and regulatory requirements.
- **Modernizing application architectures toward microservices:** Red Hat OpenShift provides a common platform for cloud-native microservices applications alongside existing traditional stateful applications. A broad choice of application frameworks, programming languages, and developer tools empowers customers to prototype innovative applications more quickly. Red Hat OpenShift also allows access to a broad range of Red Hat and third-party services, application and middleware services, API management, and storage services.
- **Adopting a consistent application platform for hybrid cloud deployments:** IT organizations that want to decouple application dependencies from the underlying infrastructure adopt container technology to migrate and deploy applications across multiple cloud environments and data centers. Red Hat OpenShift provides a consistent application development and deployment platform regardless of the underlying infrastructure. It also provides operations teams with a scalable, secure, and enterprise-grade application platform and unified container and cloud management capabilities.

Red Hat OpenShift is commonly delivered as either:

- Red Hat OpenShift Container Platform for enterprise customers that want to deploy and manage Red Hat OpenShift in their own datacenter or at a public cloud provider.
- As a cloud-based service.

2 - Gartner Infrastructure, Operations, and Cloud Summit 2020

SUPERMICRO AMD SUPERBLADE

8U SuperBlade Server System Enclosure
SBE-820H/C/J/L-822/622/422

SBA-4114S-C2N
AIOM for Front I/O

SBA-4119SG
GPU /PCIe Cards

The Supermicro® SuperBlade® powered by AMD EPYC™ Processors

Supermicro SuperBlade is built for the most demanding workloads that require high CPU density and the fastest networking available today. Supermicro's new generation blade portfolio is designed to optimize the TCO of key components for today's datacenters, such as cooling, power efficiency, node density, and networking management. In addition, the Supermicro SuperBlade powered by 3rd Gen AMD EPYC™ Processors is a trusted platform that meets enterprise customer demands for on-premises private/hybrid cloud deployments.

The Supermicro SuperBlade comes in an 8U chassis, accommodates up to 20 hot-pluggable single socket nodes, and delivers high performance with AMD EPYC™ Processors and DDR4 3200MHz memory, and fast I/O with PCIe® Gen4. There are 3 AMD powered SuperBlade models available, a SAS model, a SATA model, and a GPU accelerated model, all of which can be mixed in a single 8U enclosure. The 8U SuperBlade can support up to 40 single width GPUs or 20 double width GPUs. SuperBlade SAS/SATA models support AIOM for front I/O, which extends the Open Compute Project 3.0 specification to support a wide range of networking options in a small size form factor. The 8U SuperBlade also provides customers with advanced networking options, such as 200G HDR InfiniBand or 25G Ethernet switches.

The Supermicro SuperBlade offers the significant advantage of sharing cooling and power between the server blades to reduce power consumption compared to individual rackmount servers that contain their own non-sharable fans and power supplies. Supermicro is also committed to protecting the environment through our “We Keep IT Green®” initiative that provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

Each SuperBlade enclosure contains at least one Chassis Management Module (CMM). The CMM allows administrators to remotely manage and monitor server blades, power supplies, cooling fans, and networking switches. SuperCloud Composer (SCC) is a composable cloud management platform that provides a unified dashboard to administer software-defined data centers. SCC can orchestrate cloud workloads via the streamlined industry-standard Redfish API. SCC can also monitor and manage a broad portfolio of multi-generation Supermicro servers from a single pane of glass, including SuperBlade.

The new 3rd Gen AMD EPYC 7003 Series Processors are built around the "Zen3" core that delivers up to 19% more instructions per cycle than previous generations and contains up to 64 cores per socket for industry-leading performance on popular benchmarks and real-world workloads. As a result, Supermicro and AMD have set multiple world records for performance.[3]

The Supermicro SuperBlade powered by 3rd Gen AMD EPYC processors includes a range of compute systems designed for many of the most demanding applications: Artificial Intelligence (AI), Machine Learning (ML), High Performance Computing (HPC), enterprise applications, and cloud deployments. These systems help organizations reduce time-to-solution for a wide

3 - https://www.amd.com/en/processors/epyc-world-records

range of applications, add advanced security features, and allow all workloads to run either on-prem or in a public or private cloud. The Supermicro SuperBlade offers the highest density, maximum performance, best power efficiency, and lowest Total Cost of Ownership (TCO).

Supermicro and Red Hat

Supermicro and Red Hat are working together to offer proven, validated solutions and reference architectures across multiple Red Hat software platforms, including Red Hat OpenShift. This collaboration involves careful engineering and testing that removes much of the time, risk, and time required to deploy modern hybrid cloud infrastructure, allowing organizations to focus on their applications.

Red Hat OpenShift on SuperBlade – Reference Architecture

Customers requiring a bare metal on-premise container platform solution can use the new Red Hat feature-assisted installer available from https://cloud.redhat.com/openshift/install/metal.

A fully-populated 20-node SuperBlade in an 8U Chassis can be configured as:

- 1 **admin/provisioner** node for the Red Hat OpenShift cluster.
- 3 Red Hat OpenShift **master nodes** to manage the Kubernetes cluster and schedule pods to run on worker nodes.
- 4 Red Hat OpenShift **container-native storage nodes** to provide containers with persistent storage.
- 8 Red Hat OpenShift **worker nodes** to host the application pods.
- 3 Red Hat OpenShift **infrastructure nodes** to run routing services, the container application registry.
- 1 **load balancer** node.

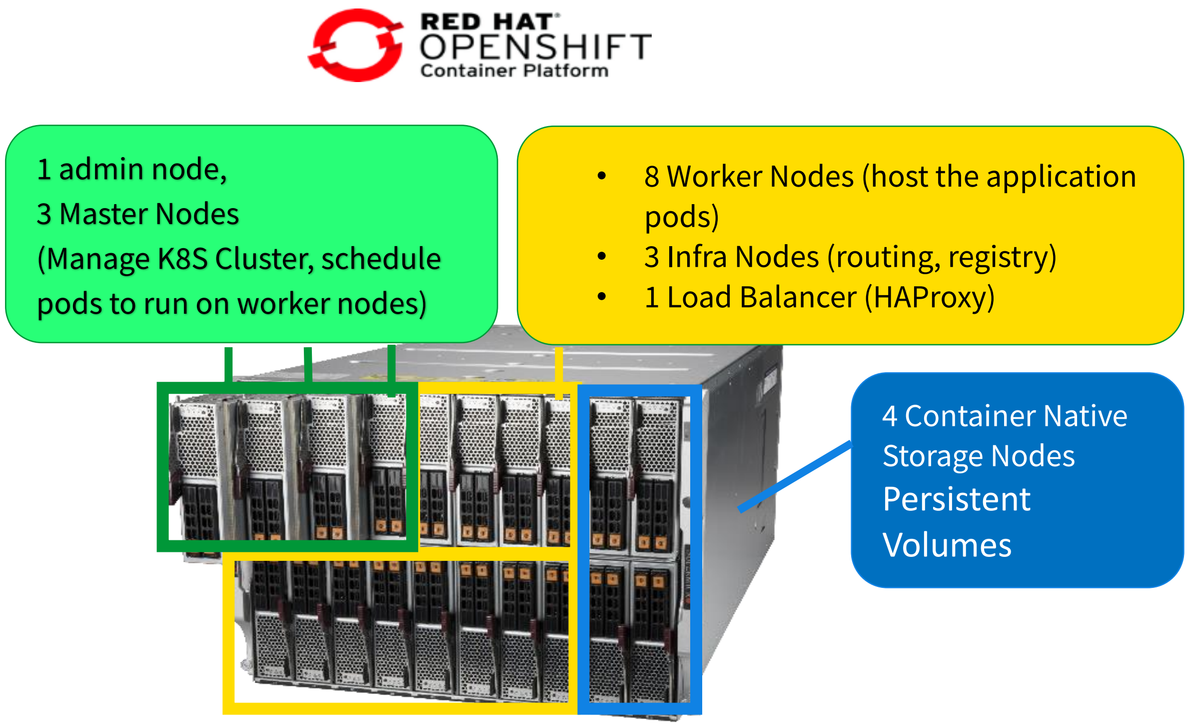

Figure 2: Deploying Red Hat OpenShift Container Platform on Supermicro SuperBlade

Installation Overview

The assisted installer focuses on bare-metal deployments and helps simplify Red Hat OpenShift Container Platform (OCP) installation. The service discovers and validates targeted hardware and significantly improves installation success rates. This section guides you through installing the [Red Hat SaaS portal] to deploy an OCP cluster using Supermicro SuperBlade as control plane nodes for compute/storage resources. The benefits of containerization include:

- Highly available control plane nodes via affinity rules.
- Resource consolidation for physical hosts.
- Simplified backup and restore.

Note: A fully-populated Supermicro SuperBlade can accommodate 20 nodes; however, this white paper uses seven nodes for testing this deployment.

The prerequisites for deploying OCP are:

1. **Shared VLAN:** The assisted installer requires all nodes to be present on the same VLAN to use a virtual IP address for both ingress and API.
2. **DHCP:** The shared VLAN requires DHCP, much like an Installer-Provisioned Installation (IPI).
3. **DNS:** The cluster requires DNS records for ingress and API access. The assisted installer supports creating these records for the user via route53, but in this white paper, the following DNS records have been created.

File: Sample DNSMASQ (/etc/dnsmasq.d) Configuration file used in this white paper
Dnsmasq is a lightweight, easy to configure, DNS forwarder, and DHCP server. It is designed to provide DNS and, optionally, DHCP to a small network.

```
# DHCP and Forward & Reverse DNS configuration
# /etc/dnsmasq.d/domain_name.conf
# Don't use /etc/resolv.conf
no-resolv
interface=baremetal
#interface=lo
listen-address=<Enter Listen_IP_addr here>
bind-interfaces

# Don't use /etc/hosts
no-hosts

# Recursive DNS
server=<Enter server_IP_addr here>

# Set local domain
domain=<Enter domain_name here>

#### DHCP (dnsmasq --help dhcp)
dhcp-range=<Enter start_IP_addr, end_IP_addr here>
#set DHCP options
dhcp-option=<Enter option:netmask,subnet_netmask_24 here>
dhcp-option=<Enter option:router,router_IP_addr here>
dhcp-option=<Enter option:dns-server,dns-server_IP_addr here>

# External API endpoint (External API VIP)
address=<Enter /api.domain_name/endpoint_IP_addr here>

# Internal API endpoint (Internal API VIP)
address=<Enter /api-int.domain_name/endpoint_IP_addr here>

# wildcard domain *.apps.<clusterName>.<baseDomain> (Ingres VIP)
address=<Enter /.apps.domain_name/ingres_IP_addr here>

# master0 / etcd-0
dhcp-host= <Enter master0_MAC_addr,master0_IP_addr,master-0.domain_name here>
address=<Enter /master-0.domain_name/master0_IP_addr here>
ptr-record=<Enter master0_ptr-record_IP_addr.in-addr.arpa,master-0.domain_name
here>

# master1 / etcd-1
dhcp-host= <Enter master1_MAC_addr,master1_IP_addr,master-1.domain_name here>
address=<Enter /master-1.domain_name/master1_IP_addr here>
ptr-record=<Enter master1_ptr-record_IP_addr.in-addr.arpa,master-1.domain_name
here>

# master2 / etcd-2
dhcp-host= <Enter master2_MAC_addr,master2_IP_addr,master-2.domain_name here>
address=<Enter /master-2.domain_name/master2_IP_addr here>
ptr-record=master2_ptr-record_IP_addr.in-addr.arpa,master-2.domain_name here>
```

```
# worker1
dhcp-host= <Enter worker1_MAC_addr,worker1_IP_addr,worker-1.domain_name here>
address=<Enter /worker-1.domain_name/worker1_IP_addr here>
ptr-record=<Enter worker1_ptr-record_IP_addr.in-addr.arpa,worker-1.domain_name
here>

# worker2
dhcp-host= <Enter worker2_MAC_addr,worker2_IP_addr,worker-2.domain_name here>
address=<Enter /worker-2.domain_name/worker2_IP_addr here>
ptr-record=<Enter worker2_ptr-record_IP_addr.in-addr.arpa,worker-2.domain_name
here>

# worker3
dhcp-host= <Enter worker3_MAC_addr,worker3_IP_addr,worker-3.domain_name here>
address=<Enter /worker-3.domain_name/worker3_IP_addr here>
ptr-record=<Enter worker3_ptr-record_IP_addr.in-addr.arpa,worker-3.domain_name
here>

#
# END OF FILE
```

Node Requirements

- **CPU architecture:** All nodes must use x86_64 CPU architecture.
- **Identical nodes:** Red Hat recommends that all nodes in a single cluster share an identical configuration (brand, model, CPU, RAM, and storage).
- **Baseboard Management Controller access:** The provisioner node must be able to access the Baseboard Management Controller (BMC) of each OCP cluster node using either IPMI, Redfish, or a proprietary protocol.
- **Latest generation:** IPI relies on BMC protocols that must be compatible across nodes. All nodes must be recent enough to support RHEL 8 for the provisioner node and RHCOS 8 for the control plane and worker nodes. Additionally, RHEL 8 ships with the most recent RAID controller drivers.
- **Registry node:** (Optional) If you are setting up a disconnected mirrored registry, then Red Hat recommends having the registry reside in its own node.
- **Provisioner node:** IPI requires one provisioner node.
- **Control plane:** IPI requires three control plane nodes for high availability.
- **Worker nodes:** While not required, a typical production cluster has one or more worker nodes. Smaller clusters are more resource-efficient for administrators and developers during development, production, and testing.
- **Network interfaces:** Each node must have at least one 10 Gigabit Ethernet (GbE) network interface for the routable bare metal network. By default, each node must also have one 10GbE network interface for

a provisioning network when using the provisioning network for deployment. Network interface names must follow the same naming convention across all nodes. For example, the first NIC name on a node, such as `eth0` or `eno1`, must be the same name on all the other nodes. The same principle applies to the remaining NICs on each node.

- **Unified Extensible Firmware Interface (UEFI):** Installer-provisioned installation requires UEFI boot on all OCP nodes when using IPv6 addressing on the provisioning network. UEFI Device PXE settings must also be set to use the IPv6 protocol on the provisioning network NIC. (Omitting the provisioning network removes this requirement.)

One of the master nodes will be used to bootstrap the cluster and then converted into another master node. Both the kubeconfig and a link to the cluster console will be available for download once the cluster installation is complete.

Configure the node BIOS settings for the virtual CD mount and enable virtualization for the hyper-converged nodes. Next, configure the nodes by assigning a hostname and an IP address along with the desired IPMI user name and password, as shown below.

Blade #	Hostname	IPMI IP	IPMI User	IPMI PW
B1	admin	172.17.41.80	SMCI	Super123
B2	master-0	172.17.41.81	SMCI	Super123
B3	master-1	172.17.41.82	SMCI	Super123
B4	master-2	172.17.41.83	SMCI	Super123
B5	worker-1	172.17.41.84	SMCI	Super123
B6	worker-2	172.17.41.85	SMCI	Super123
B7	worker-3	172.17.41.86	SMCI	Super123

<table>
<tr>
<td>

Step 1

Set Boot Mode Select to **UEFI**.

</td>
<td>

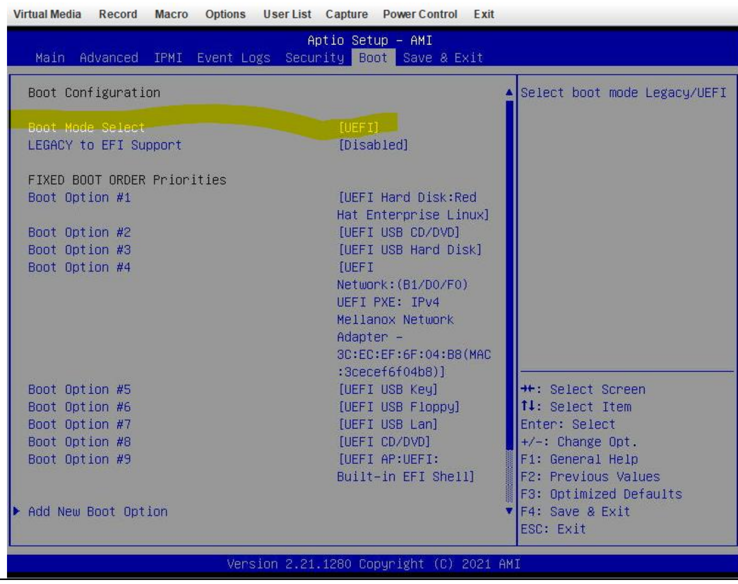

</td>
</tr>
<tr>
<td>

Step 2

Set **Boot Option #1** to the correct UEFI boot disk.

</td>
<td>

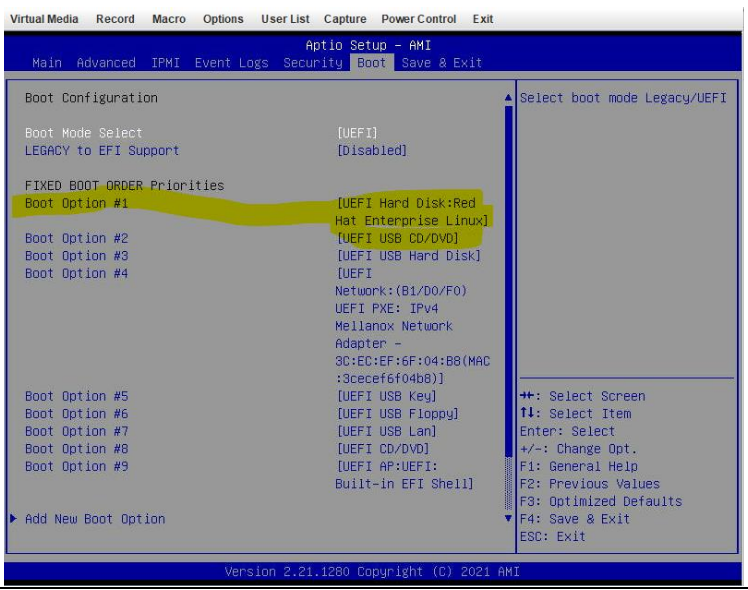

</td>
</tr>
</table>

Step 3

Set **SVM Mode** to **Enabled** to enable CPU virtualization.

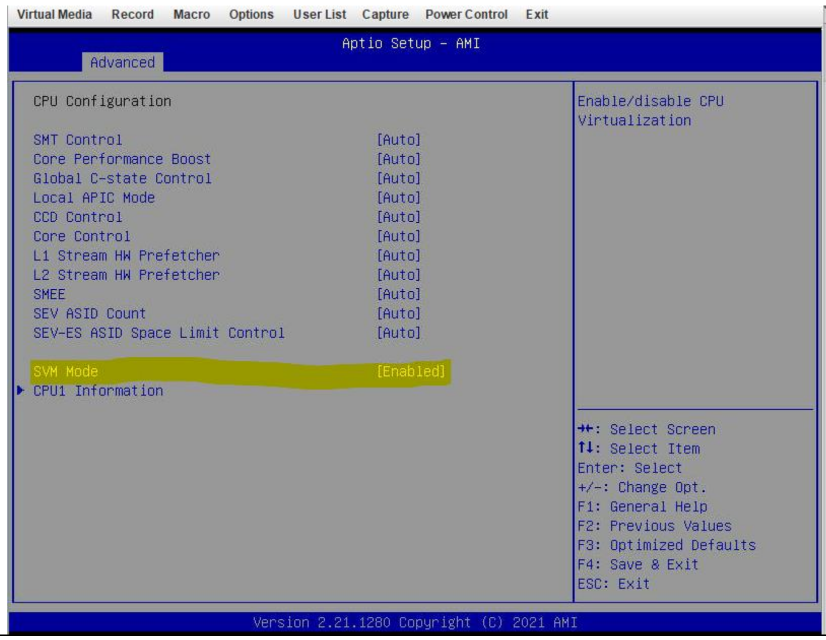

Step 4

Set **SR-IOV Support** to **Enabled**.

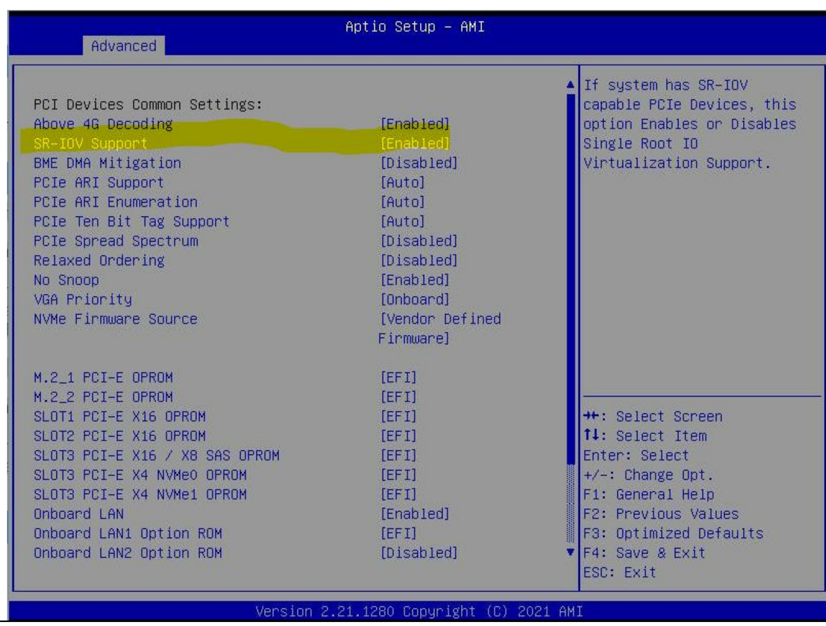

<table>
<tr>
<td>Step 5

Set IOMMU to Enabled.</td>
<td>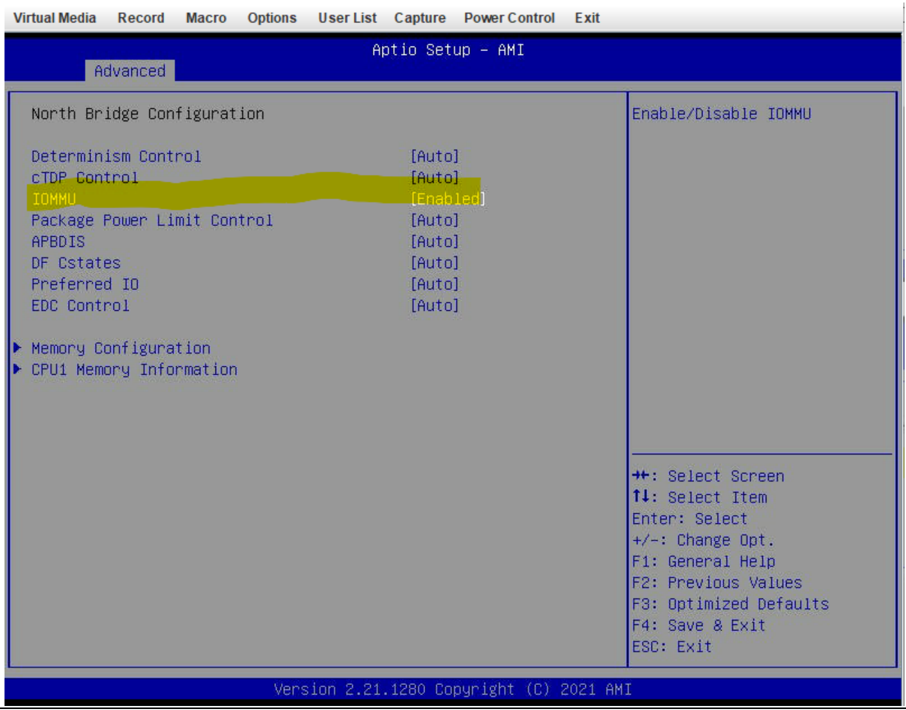
</td>
</tr>
<tr>
<td>Step 6

Set PCI AER Support to Enabled.</td>
<td>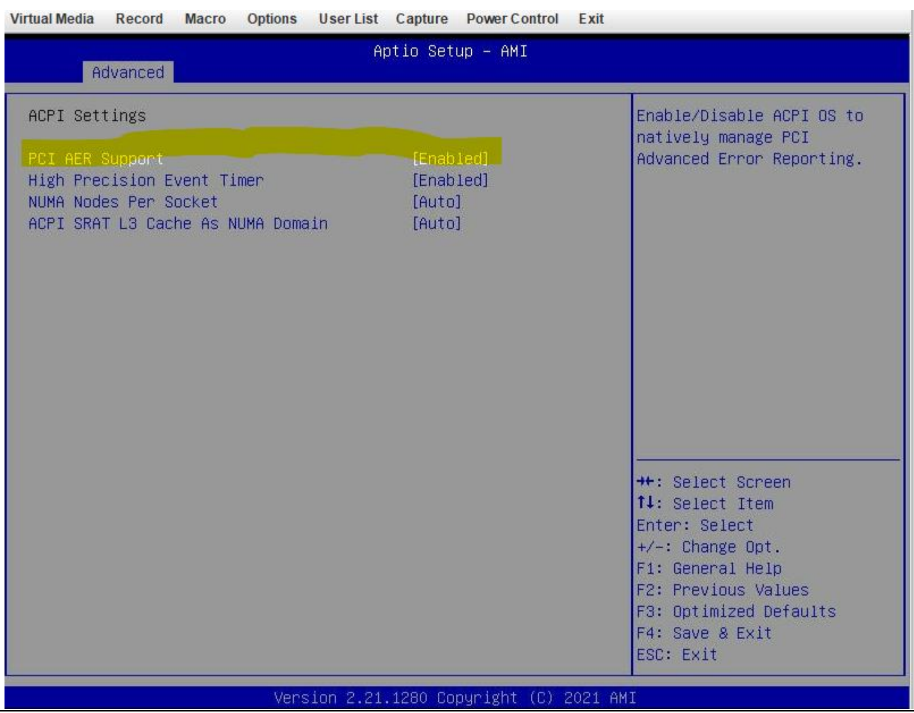
</td>
</tr>
<tr>
<td>Step 7

Download the Supermicro IPMI tool to</td>
<td>SMCIPMITool is an out-of-band Supermicro utility that allows users to interface with SuperBlade systems and IPMI devices via the Command Line Interface (CLI). It is available here.
https://www.supermicro.com/SwDownload/SwSelect_Free.aspx?cat=IPMI</td>
</tr>
</table>

mount/unmount USB virtual media and System Power Control.	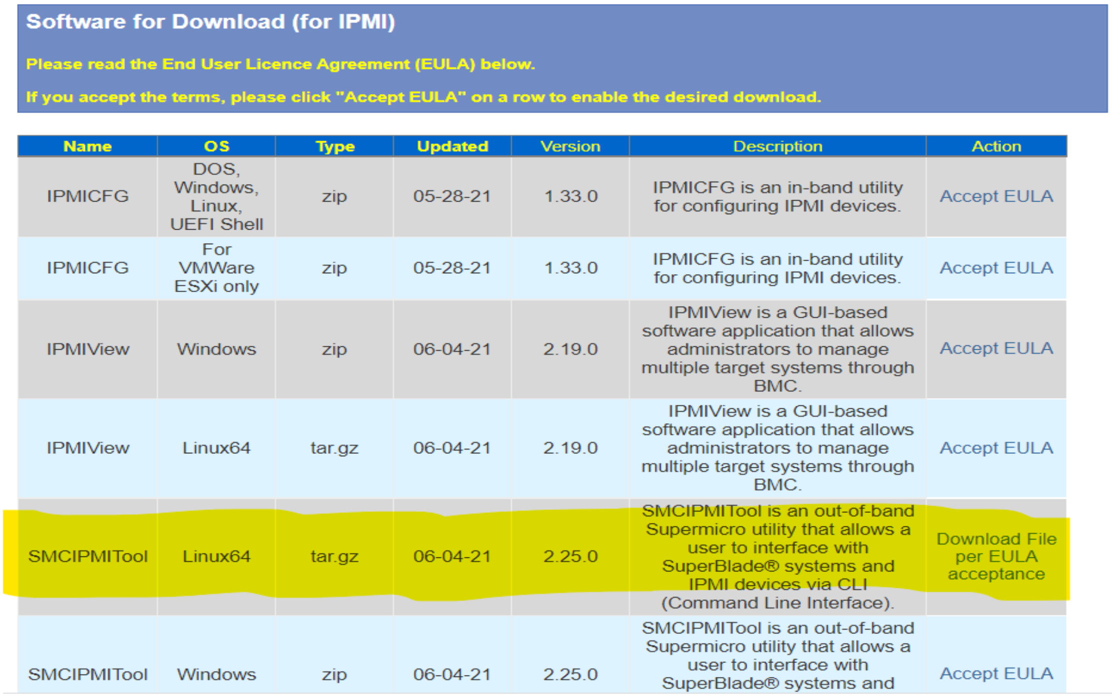
Step 8 Log in to the portal using your Red Hat credentials. This will auto-populate any pull secrets required to pull container images associated with the account. Once logged in, click the **Create Cluster** button to start the new OCP installation.	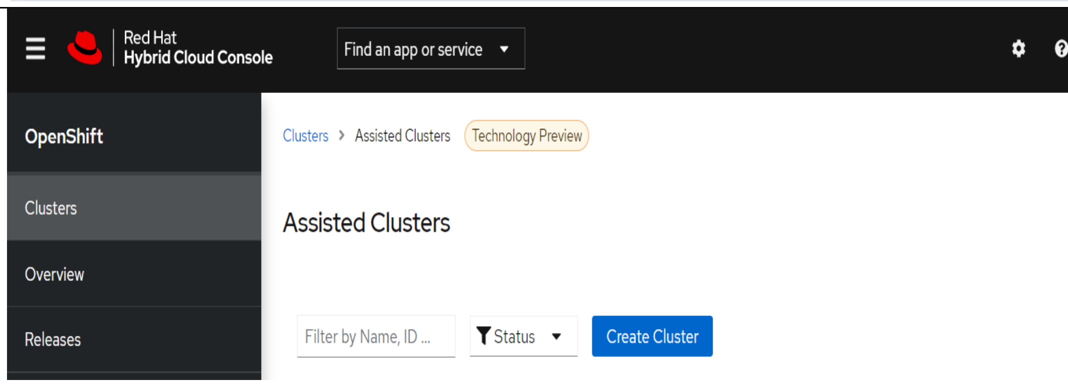

Name	OS	Type	Updated	Version	Description	Action
IPMICFG	DOS, Windows, Linux, UEFI Shell	zip	05-28-21	1.33.0	IPMICFG is an in-band utility for configuring IPMI devices.	Accept EULA
IPMICFG	For VMWare ESXi only	zip	05-28-21	1.33.0	IPMICFG is an in-band utility for configuring IPMI devices.	Accept EULA
IPMIView	Windows	zip	06-04-21	2.19.0	IPMIView is a GUI-based software application that allows administrators to manage multiple target systems through BMC.	Accept EULA
IPMIView	Linux64	tar.gz	06-04-21	2.19.0	IPMIView is a GUI-based software application that allows administrators to manage multiple target systems through BMC.	Accept EULA
SMCIPMITool	Linux64	tar.gz	06-04-21	2.25.0	SMCIPMITool is an out-of-band Supermicro utility that allows a user to interface with SuperBlade® systems and IPMI devices via CLI (Command Line Interface).	Download File per EULA acceptance
SMCIPMITool	Windows	zip	06-04-21	2.25.0	SMCIPMITool is an out-of-band Supermicro utility that allows a user to interface with SuperBlade® systems and	Accept EULA

<table>
<tr>
<td>Enter the desired cluster name and the version of Red Hat OpenShift to be deployed.

This example uses ldc2-sanjose as the cluster name.

Click Next to continue.</td>
<td>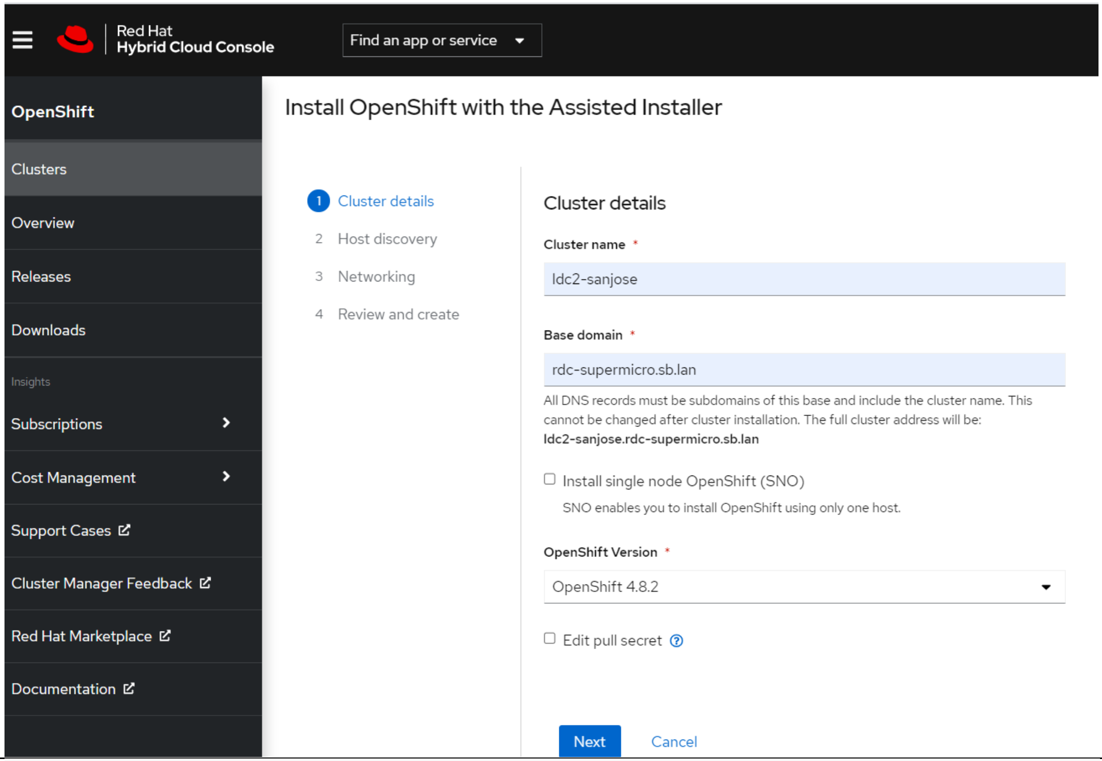
</td>
</tr>
<tr>
<td>Step 9

Click the Generate Discovery ISO button.</td>
<td>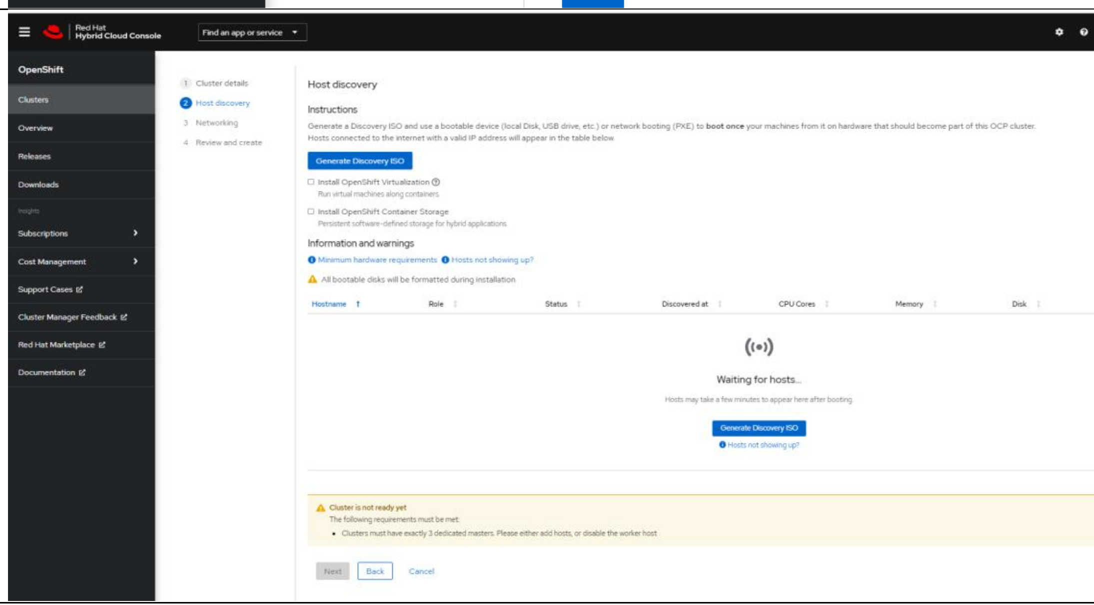
</td>
</tr>
</table>

<table>
<tr><td>Step 10

Provide the SSH public key of the admin/provisioner node from where you can manage the cluster.

Click the Generate Discovery ISO button.</td><td>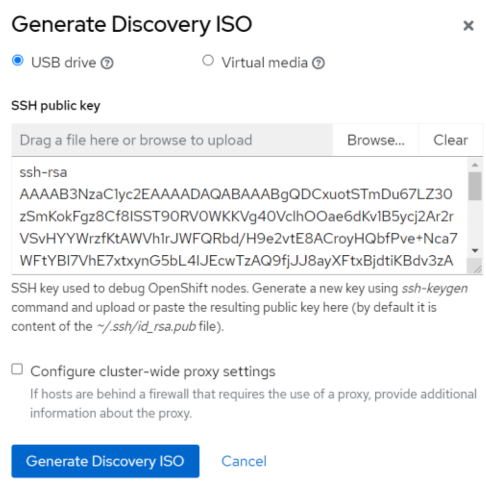
</td></tr>
<tr><td>Step 11

Download the Discovery ISO to a Samba file share or webserver from where you can mount the image using the Supermicro command line tools.</td><td>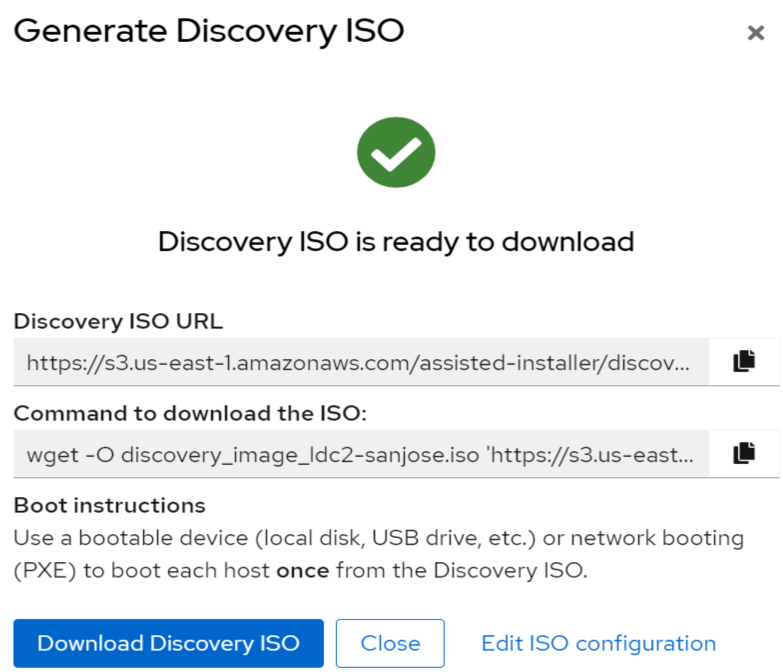
</td></tr>
<tr><td>Step 12

Unmount any existing virtual media from all control and compute nodes</td><td><code>for i in {1..6};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 wsiso umount;
done</code></td></tr>
<tr><td>Step 13

Mount the downloaded Red Hat OpenShift</td><td><code>for i in {1..6};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 wsiso mount 172.17.41.1 /share/discovery_image_ldc2-sanjose.iso;
done</code></td></tr>
</table>

<table>
<tr><td>install virtual media, which can be either an SMB share or http URL for the downloaded ISO.

This example uses an SMB share.</td><td></td></tr>
<tr><td>Step 14

Set the one-time bootable device to the virtual USB media as in Step 7.</td><td>for i in {1..6};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 ipmi power bootoption 13;
done</td></tr>
<tr><td>Step 15

Check the power status of all master and worker nodes.</td><td>for i in {1..6};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 ipmi power status;
done</td></tr>
<tr><td>Reset or turn on the power to mount the one-time virtual USB media which was set in step 13,</td><td>for i in {1..5};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 ipmi power reset;
done
Or
for i in {1..5};
do ./SMCIPMITool 172.17.41.8$i SMCI Super123 ipmi power up;
done</td></tr>
<tr><td>Step 16

The inventory should look like this once the servers boot with the discovery image and the inventory is finished.</td><td>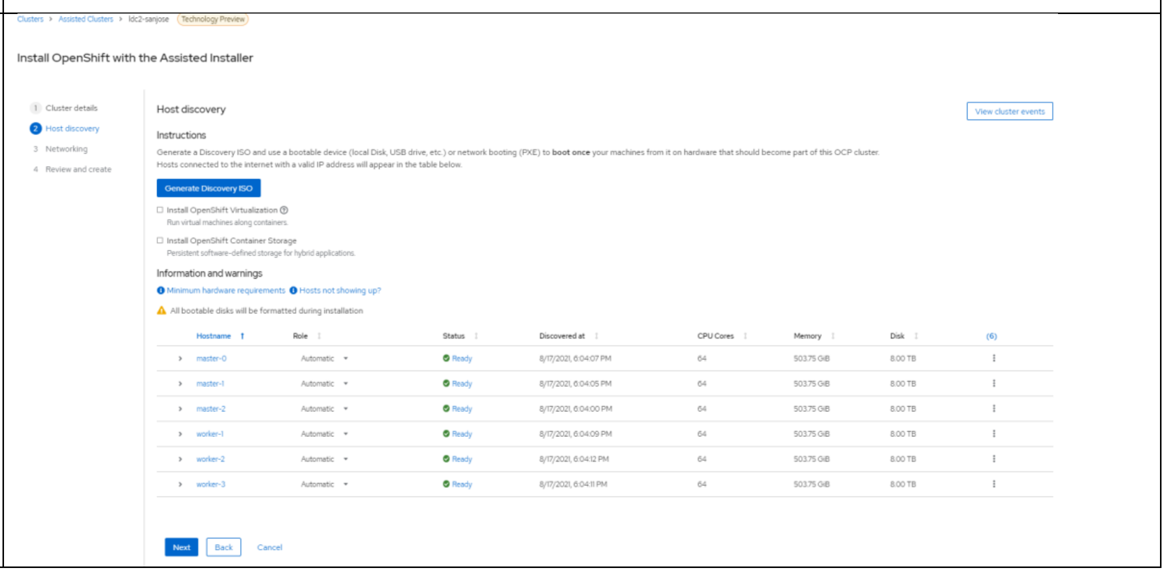
</td></tr>
</table>

Step 17

Assign the appropriate roles and hostnames for the systems.

Check the **Install OpenShift Virtualization** and **Install OpenShift Container Storage** checkboxes for the hyperconverged worker nodes.

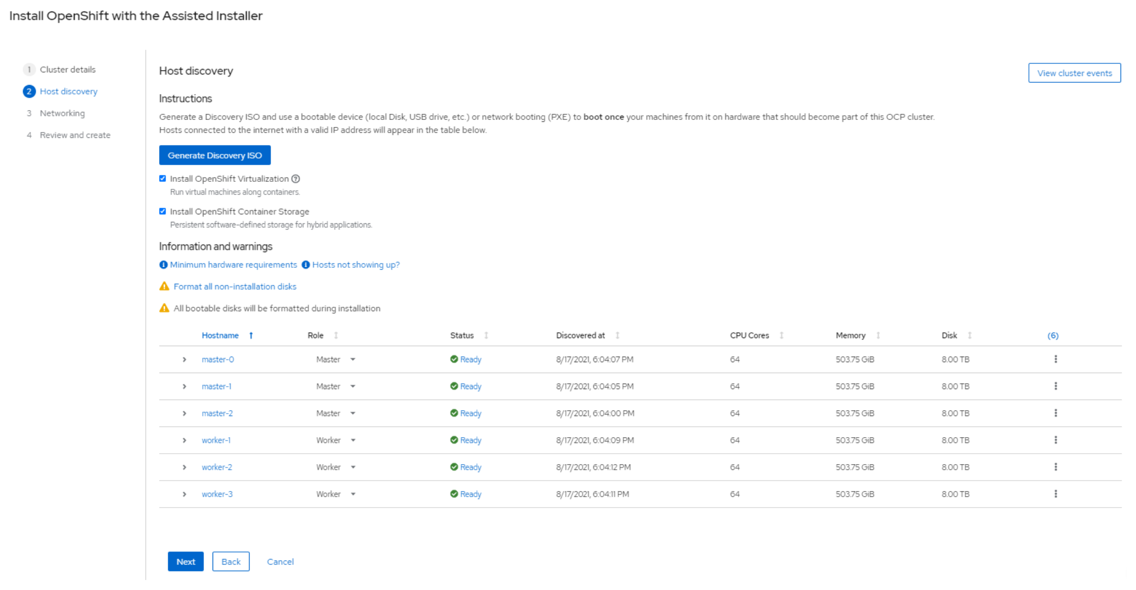

Step 18

Either add the API and ingress VIPs or allow DCHP assignment.

This example uses software-defined networking.

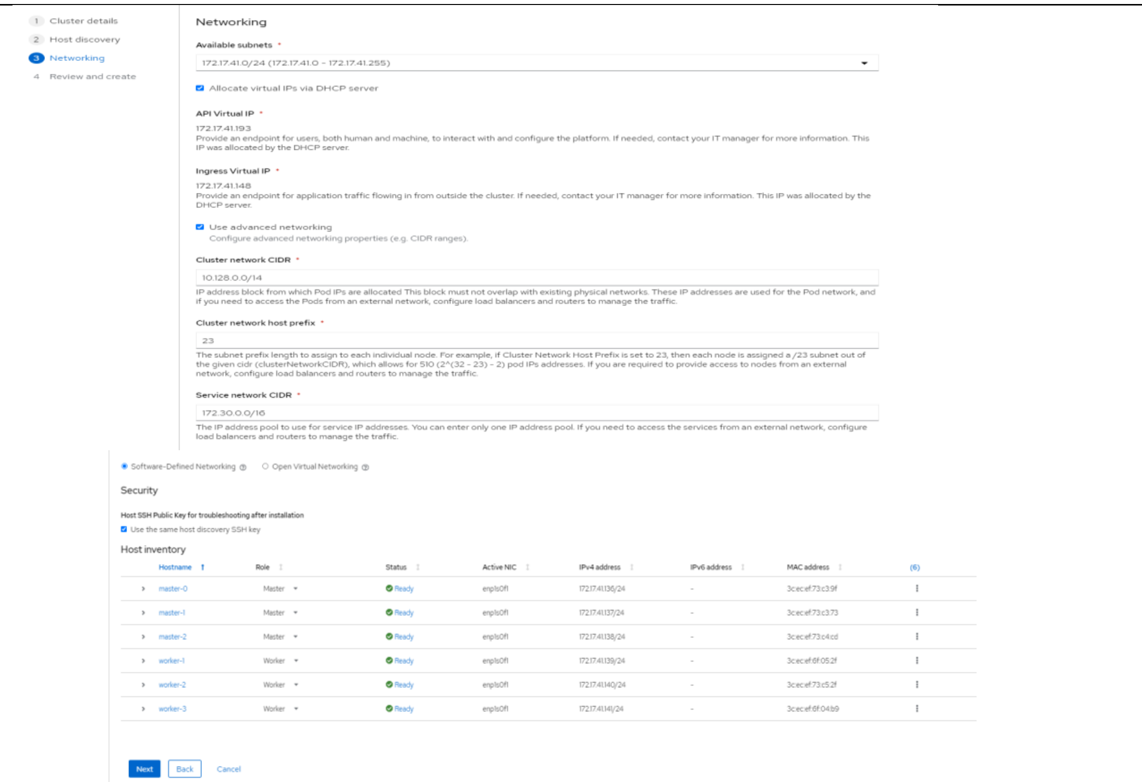

Step 19

Click the **Install Cluster** button to begin the bare metal/VM cluster installation.

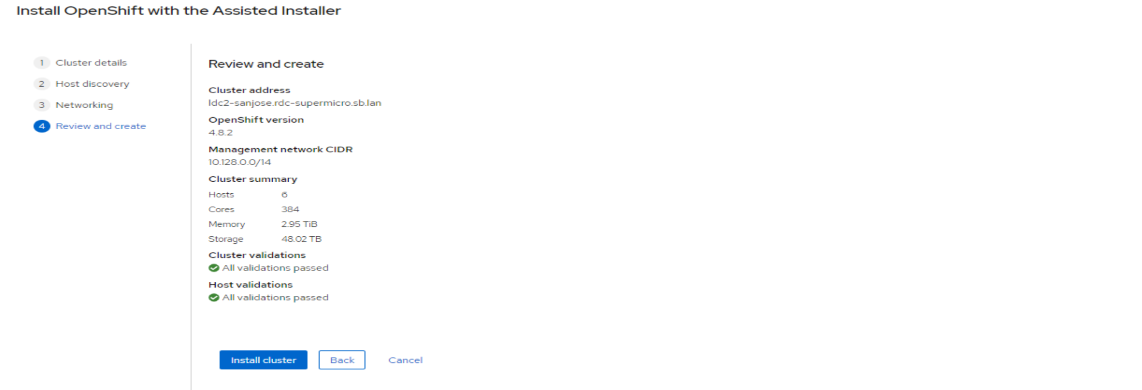

Step 20

One of the masters will be used to bootstrap the cluster then converted into another master node.

The kubeconfig and a link to the cluster console will be available for download once the cluster installation completes.

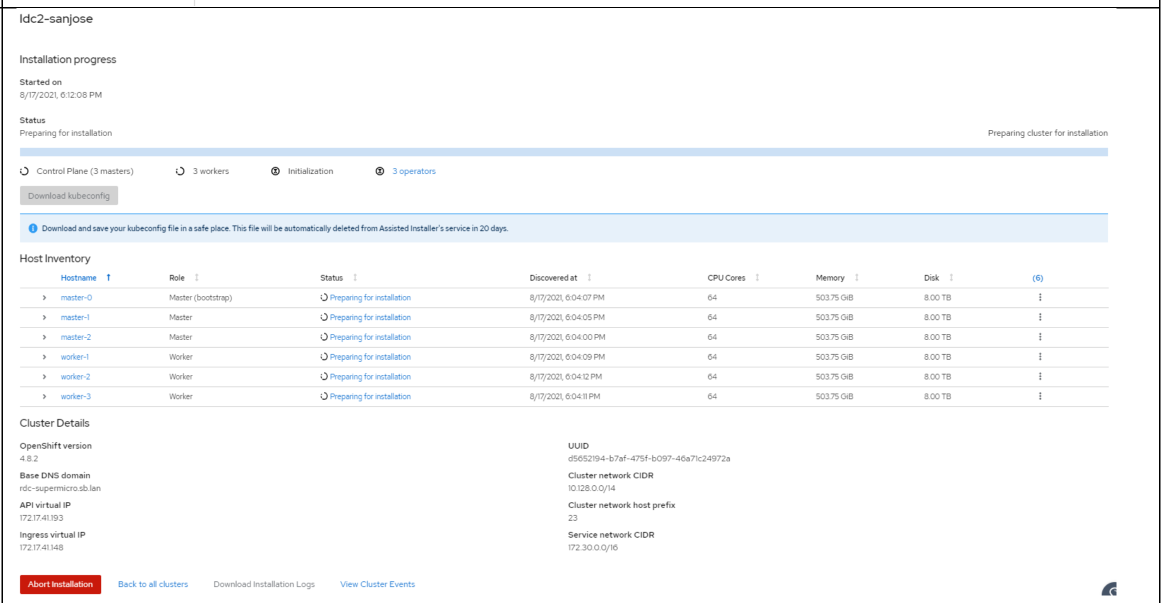

Step 21

Once the installation is complete, you can log in to the console with the provided DNS entries and the kubeadmin password.

Download the kubeconfig file for the OCP CLI tools.

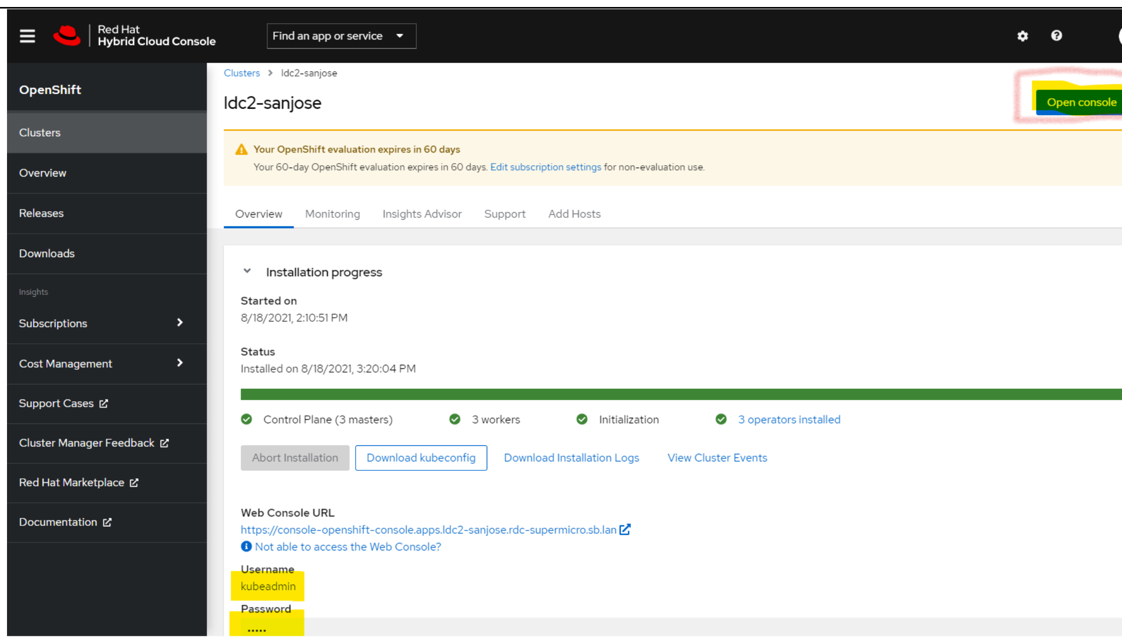

<table>
<tr>
<td>Step 22

Add the DNS entries to your external DNS (preferred) or localhost file.</td>
<td>
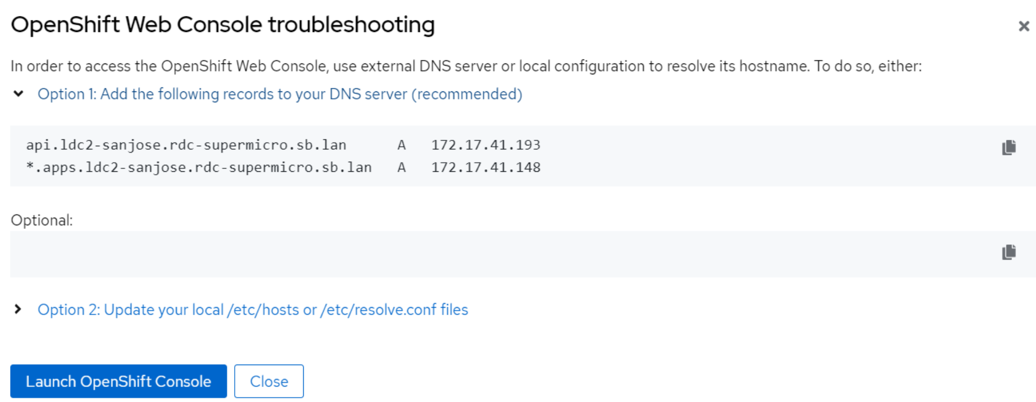

(Local /etc/hosts)
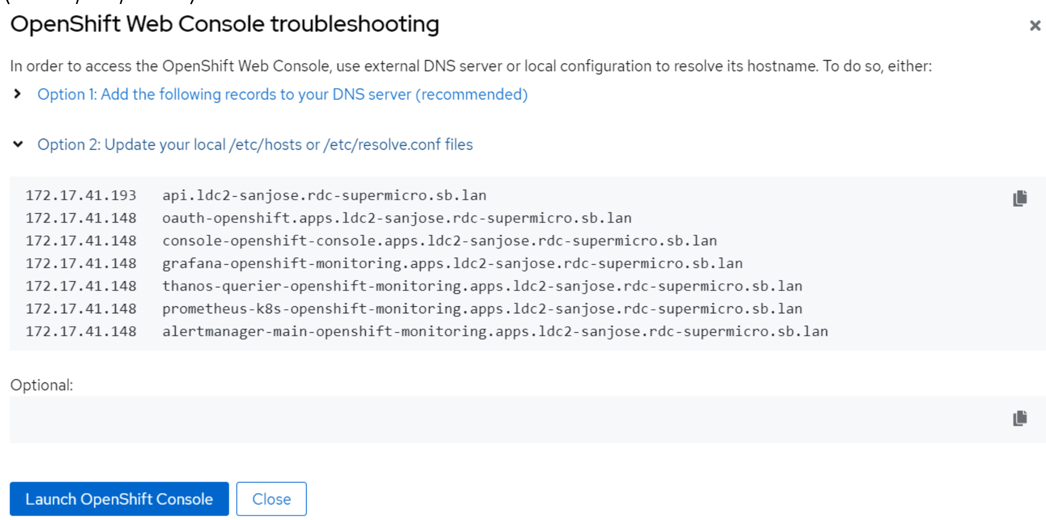

</td>
</tr>
<tr>
<td>Step 23

Launch the OCP console by clicking the “Launch OpenShift Console” in the above step 21 and login with the kubeadmin and its passwords recorded in Step 20.</td>
<td>
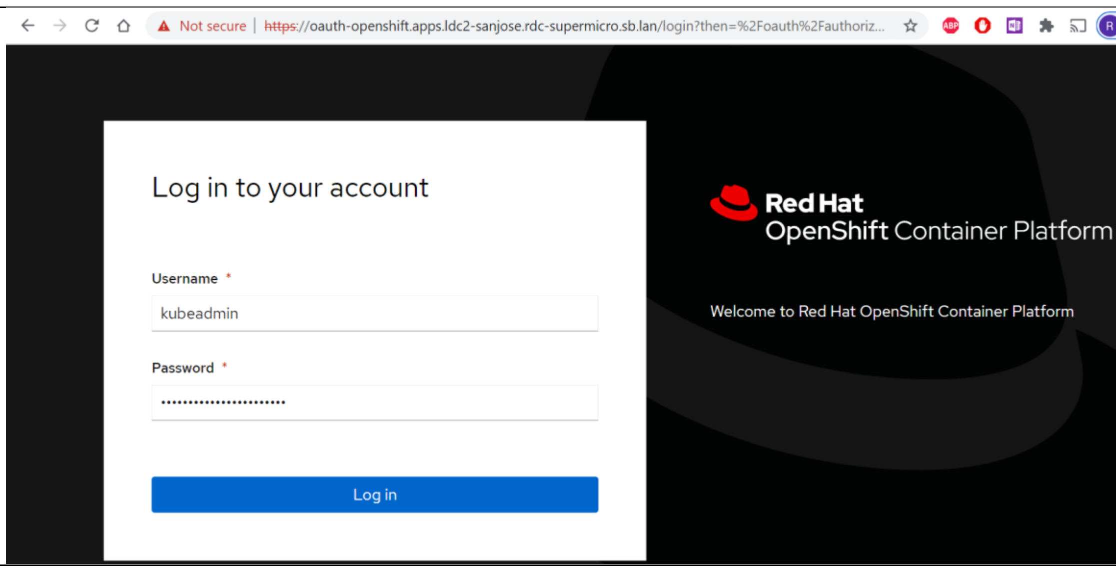

</td>
</tr>
</table>

Step 24

Access the OCP **Nodes** screen.

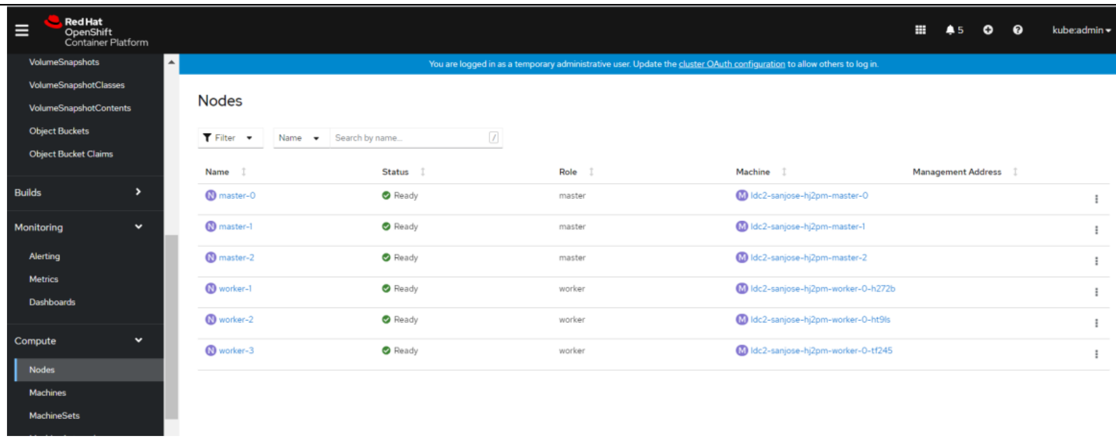

Step 25

Export the downloaded kubeconfig file to run the OCP commands.

```
export KUBECONFIG=/home/kni/kubeconfig
```

Step 26

Get the Red Hat OpenShift nodes status.

```
[kni@provisioner ~]$ oc get nodes
NAME        STATUS   ROLES    AGE   VERSION
master-0    Ready    master   50m   v1.21.1+051ac4f
master-1    Ready    master   52m   v1.21.1+051ac4f
master-2    Ready    master   33m   v1.21.1+051ac4f
worker-1    Ready    worker   37m   v1.21.1+051ac4f
worker-2    Ready    worker   37m   v1.21.1+051ac4f
worker-3    Ready    worker   38m   v1.21.1+051ac4f
```

Step 27

Get the installed Red Hat OpenShift version.

```
[kni@provisioner ~]$ oc version
Client Version: 4.8.4
Server Version: 4.8.2
```

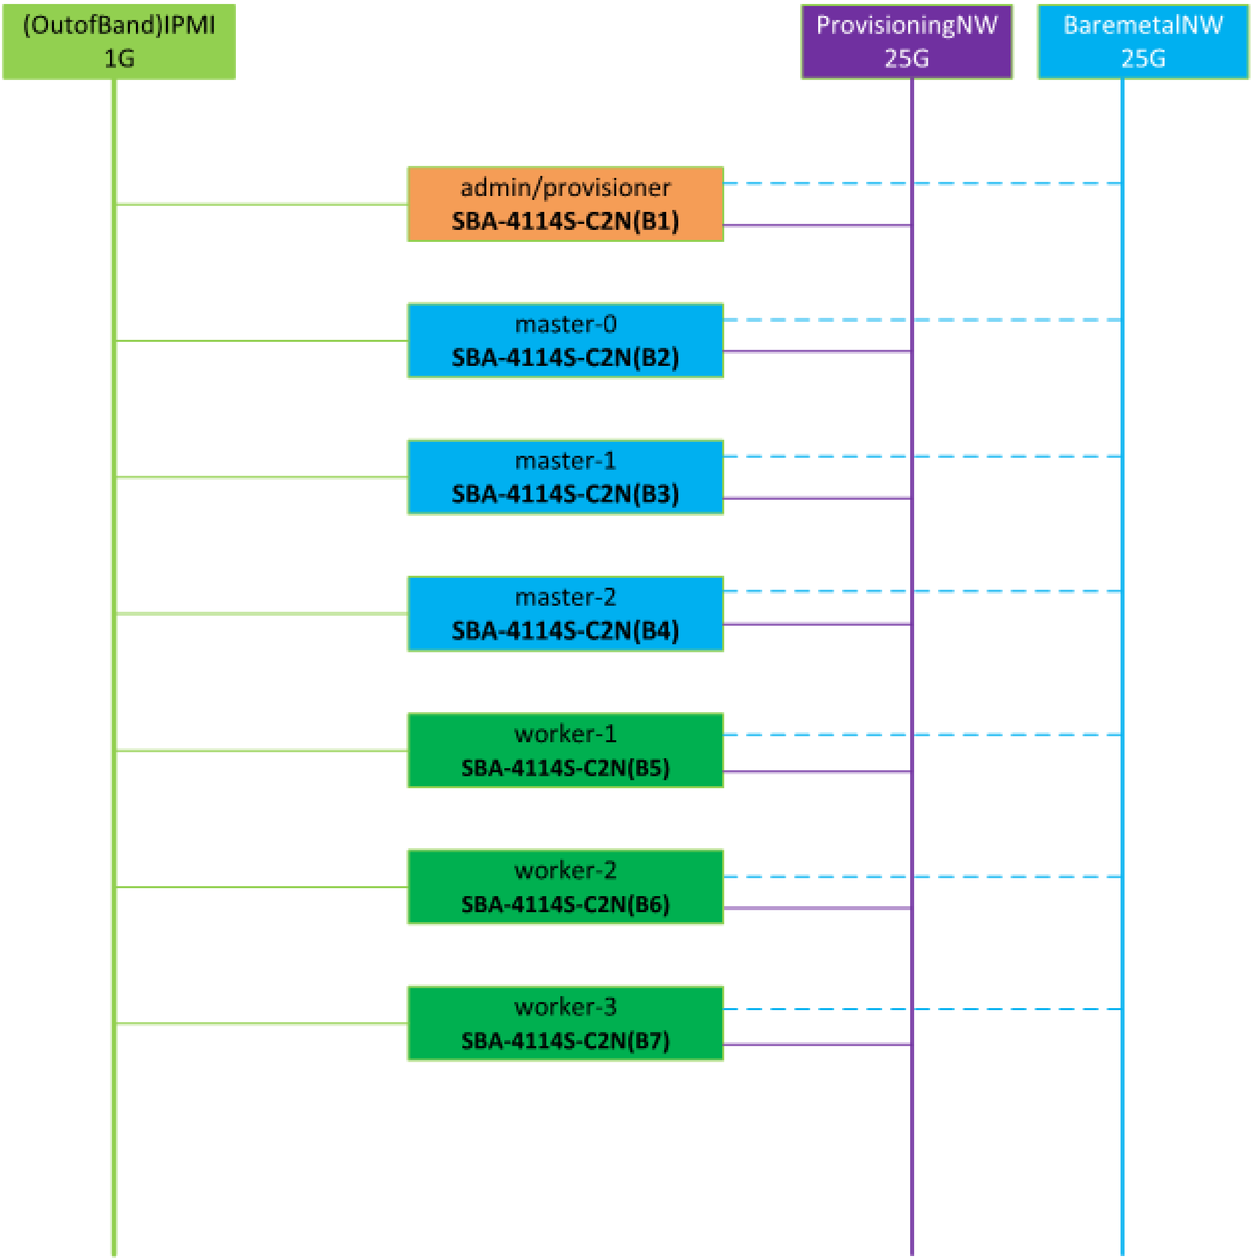

Figure 3: Topology diagram – Red Hat OpenShift admin, master, and worker node deployment on Supermicro SuperBlade

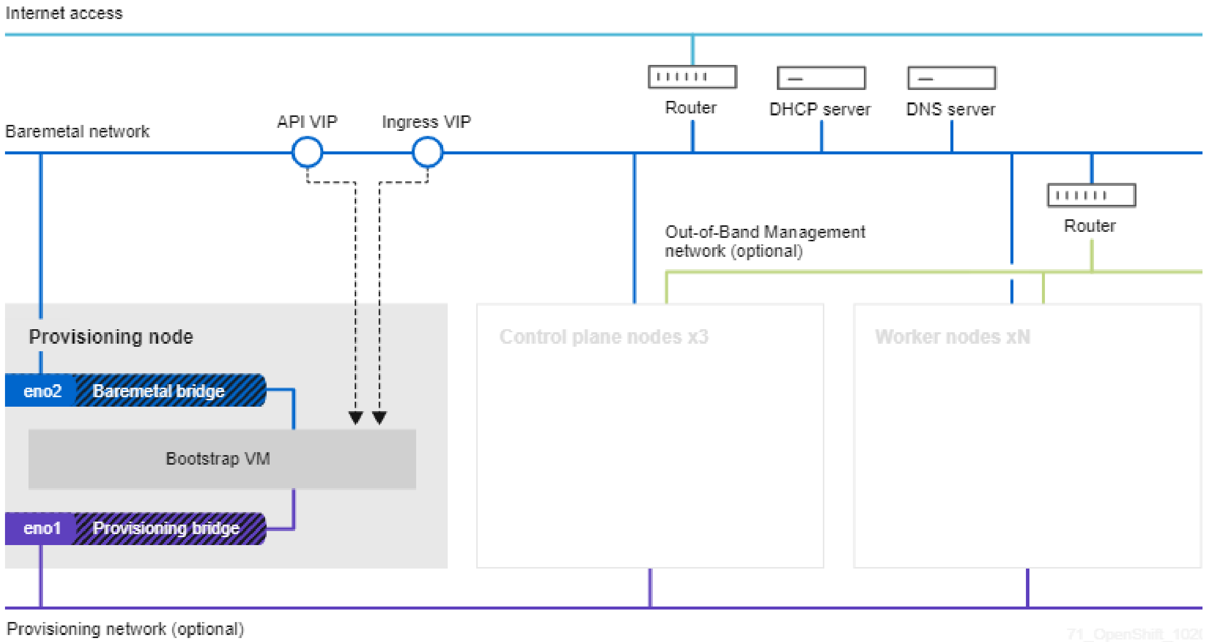

Figure 4: Red Hat OpenShift control plane nodes

The installer automatically destroys the bootstrap VM and moves the virtual IP addresses (VIPs) to the control plane nodes when the Red Hat OpenShift control plane node installation is complete and fully operational.

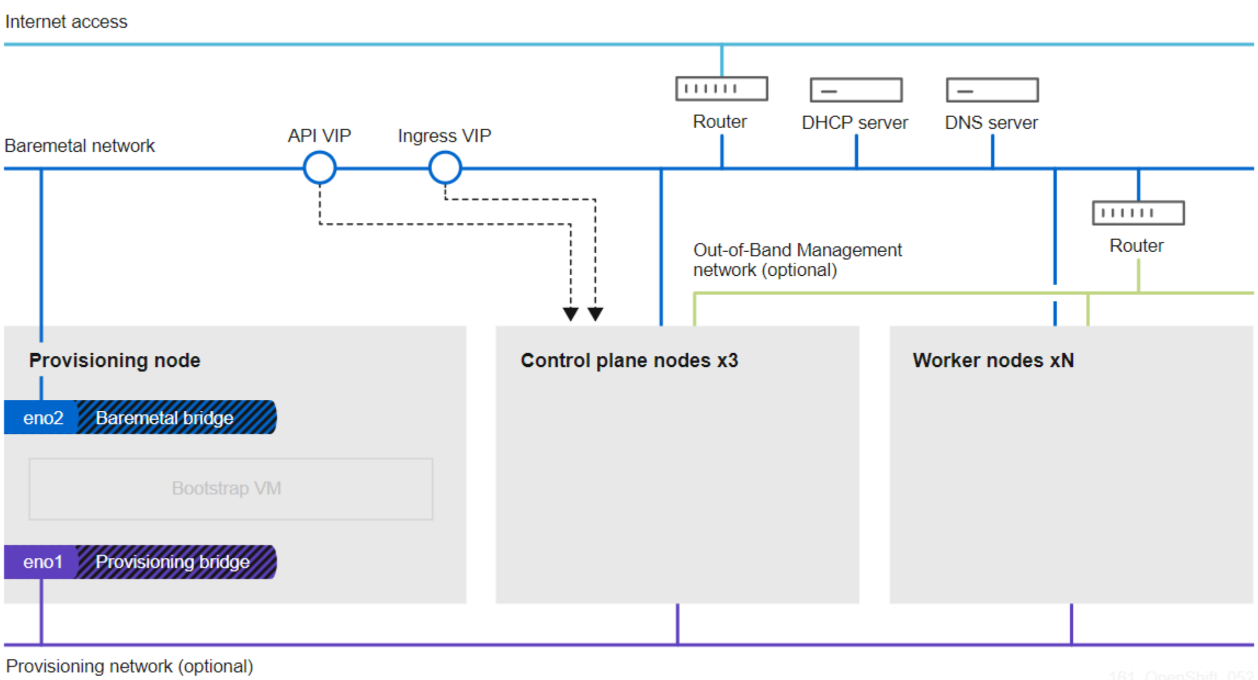

Figure 5: Red Hat OpenShift control plane nodes fully operational

Reference SuperBlade BOM

The Bill of Materials (BOM) used for the Red Hat OpenShift Container Platform installation described in this white paper includes:

ADMIN NODE(x1)			
Part Type	**Part #**	**Part Description**	**Qty**
System	SBA-4114S-C2N	Supermicro SuperBlade (AMD Powered Single-Socket SAS model)	1
Processor	PSE-MLN7543P-001	3rd Gen AMD EPYC™ 7003 Series Processor (7543P - 32 Cores, 2.8GHz, 225W TDP)	1
Memory	MEM-DR416L-SL02-ER32	16GB DDR4-3200	8
Internal Storage (for Operating System)	HDS-IMN0-SSDPELKX020T8	2 TB M.2 NVMe SSD	2
Internal Storage (for Data)	HDS-IUN2-SSDPE2KX020T8	2 TB U.2 NVMe SSD	2
MASTER NODES(x3)			
Part Type	**Part #**	**Part Description**	**Qty**
System	SBA-4114S-C2N	Supermicro SuperBlade (AMD Powered Single-Socket SAS model)	1
Processor	PSE-MLN7543P-001	3rd Gen AMD EPYC™ 7003 Series Processor (7543P - 32 Cores, 2.8GHz, 225W TDP)	1
Memory	MEM-DR432L-SL02-ER32	32GB DDR4-3200	8
Internal Storage (for Operating System)	HDS-IMN0-SSDPELKX020T8	2 TB M.2 NVMe SSD	2
Internal Storage (for Data)	HDS-IUN2-SSDPE2KX020T8	2 TB U.2 NVMe SSD	2
WORKER NODES(x3)			
Part Type	**Part #**	**Part Description**	**Qty**
System	SBA-4114S-C2N	Supermicro SuperBlade (AMD Powered Single-Socket SAS model)	1
Processor	PSE-MLN7763-0312	3rd Gen AMD EPYC™ 7003 Series Processor (7763 – 64 Cores, 2.45GHz, 280W TDP)	1
Memory	MEM-DR464L-SL01-ER32	64GB DDR4-3200	8
Internal Storage (for Operating System)	HDS-IMN0-SSDPELKX020T8	2 TB M.2 NVMe SSD	2
Internal Storage (for Operating Data)	HDS-IUN2-SSDPE2KX020T8	2 TB U.2 NVMe SSD	2

Key Takeaways and Business Values

Red Hat® OpenShift® Container Platform provides enterprise Kubernetes bundled CI/CD pipelines, automated builds, and deployments that allow developers to focus on application logic while leveraging all best-of-class enterprise infrastructure.

Deploying Red Hat OpenShift on Supermicro® SuperBlade® gives organizations a consistent application development and deployment platform. Operations teams benefit from the scalable, secure, enterprise-grade application platform with unified container and cloud management capabilities. This validated reference architecture allows customers to efficiently design, deploy, and operate a containerized DevOps/PaaS platform.

Supermicro SuperBlade powered by 3rd Gen AMD EPYC™ 7003 Series Processors allows Kubernetes to be run in production environments with the operational efficiency and consistency needed to meet various SLAs and IT initiatives. It enables increased productivity, reduced total cost of ownership, and scalability into your datacenter.

SUPERMICRO SUPERBLADE PROVIDES BUSINESS VALUES TO ENTERPRISE CUSTOMERS

Better performance per watt and per dollar
SuperBlade offers:

- Maximum performance and highest density at lowest cost.
- Cost benefits from reduced cabling, shared power and cooling efficiencies, and front serviceability.

Faster, better performance on cloud workloads
SuperBlade offers better performance with:

- High-core-count 3rd Gen AMD EPYC™ processors.
- Unmatched network performance with 10/25/100 GbE network switch connectivity.
- High bandwidth and low latency with 100/200 G InfiniBand.

Greener environmental impact and lower TCO

- Maximum power optimization using titanium-level power supplies with 96% efficiency.
- Modular upgrades that reduce e-waste.