

Deploy Red Hat OpenShift 4.11 3-node compact cluster with Assisted Installer on Ampere Altra platform

This tutorial explains how to deploy Red Hat OpenShift Container Platform 4.11 with Assisted Installer on Ampere Altra Platform. We will install Rook Ceph Operator for block storage, shared filesystems, and Object storage. Estimated time to complete this tutorial: 1 hour.

Overview

Cloud native computing is an approach in software development that utilizes cloud computing to "build and run **scalable** applications in modern, dynamic environments such as public, private, and hybrid clouds".¹ It's growing rapidly on many kinds of workloads and use cases appearing in a large number of areas with all kinds of physical footprints. Developing and deploying applications in cloud native way is increasingly being adopted from small business to enterprises.

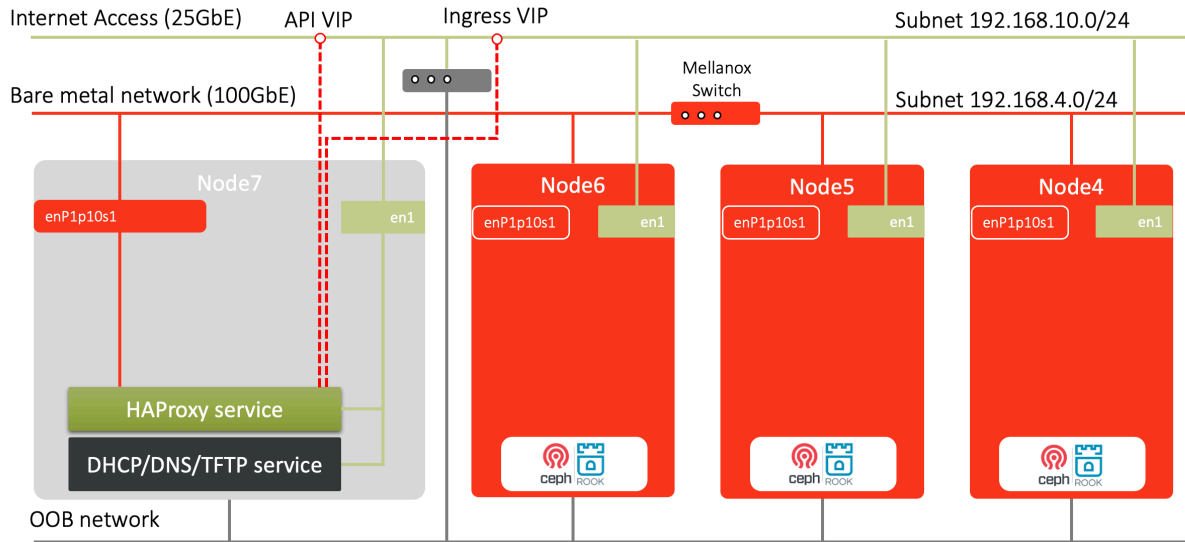
At Ampere, we observe the trend of infrastructure providers and application owners wanting a consistent workload lifecycle and predictable performance across their business. We offer Ampere Altra CPU family with the values for cloud native workloads with High Availability (HA).

- High core counts in single socket processors
- Predictability for less jitter and lower latency, while protecting against noisy neighbor effects in the processor
- Linear Scalability to maximize heavily loaded server performance
- The most sustainable, low power architecture for power sensitive edge locations and more efficient data centers
- Scale-out computational horsepower with optimized video codecs for high performance video transcoding

The combination of Ampere Altra processors and Red Hat OpenShift 3-node Compact Cluster will be a great match for the scenarios below:

- CDN Edge cloud
- Video Service Providers
- Digital Service Providers
- Small & Medium Business who develop cloud native applications
- 5G User Plane Function (UPF)

Network overview for OpenShift 3-node Compact cluster on Ampere Altra platform



Prerequisites

- A DNS service like Bind (named) running on the bastion node
- A HAProxy service as the external load balancer on the bastion node

Instructions

The following is the step-by-step guide to install OpenShift Container Platform 4.11 with Assisted Installer on Ampere Altra Platform:

1. Login to <http://cloud.redhat.com> and click “OpenShift”, then click “Data Center” for “Bare Metal (Arm64)”, and “Create Cluster” for new cluster
2. Click “Assisted Installer”



3. Provide domain name, cluster name and IP addresses for each node (if the static IP address option was chosen) for the cluster details.
 - a. We used a local DNS for managing domain names under “hhii.amp” and use “ocp4” as the cluster name.

- 1 Cluster details
- 2 Static network configurations >
- 3 Host discovery
- 4 Networking
- 5 Review and create

Cluster details

Cluster name *

ocp4



Base domain *

hhii.amp

All DNS records must be subdomains of this base and include the cluster name. This cannot be changed after cluster installation. The full cluster address will be:

ocp4.hhii.amp

OpenShift version *

OpenShift 4.11.0

☐ Install single node OpenShift (SNO)

SNO enables you to install OpenShift using only one host.

☐ Edit pull secret ?

☒ Use arm64 CPU architecture ?

Make sure all the hosts are using arm64 CPU architecture.

Hosts' network configuration

☐ DHCP server ☒ Static network configuration

☐ Enable encryption of installation disks on control plane nodes

☐ Enable encryption of installation disks on workers

Next

Cancel

Figure A-2

b. Enter the static IP address range, default gateway and DNS

Install OpenShift with the Assisted Installer

- 1 Cluster details
- 2 Static network configurations v
- 3 Host discovery
- 4 Networking
- 5 Review and create

Static network configurations

Network configuration can be done using either the form view or YAML view. Configurations done in this step are for discovering hosts.

Configure via: ☒ Form view ☐ YAML view

Network-wide configurations

The following configurations are applicable to all the hosts.

Internet protocol version

IPv4

☐ Use VLAN ?

IPv4

Machine Network *

10.9.27.24

/ 24

(10.9.27.0 - 10.9.27.255)

Default gateway *

10.9.27.1

DNS *

10.9.27.27

Next

Back

Cancel

Figure A-3

- c. Enter all 3 hosts' MAC addresses and assigned IP addresses on the Internet-access NICs.

Install OpenShift with the Assisted Installer

Static network configurations

Network configuration can be done using either the form view or YAML view. Configurations done in this step are for discovering hosts.

Configure via: ☒ Form view ☐ YAML view

▼ Host 1

MAC Address *
a0:b1:c2:d4:e5:f6

IP address (IPv4)
10.9.27.24

▼ Host 2

MAC Address *

IP address (IPv4)

Add another host

Figure A-4

- d. When all 3 nodes' static network configurations are ready, click “Next” button to move to Step 3

Install OpenShift with the Assisted Installer

Static network configurations

Network configuration can be done using either the form view or YAML view. Configurations done in this step are for discovering hosts.

Configure via: ☒ Form view ☐ YAML view

▼ Host 1

MAC Address *
a0:b1:c2:d4:e5:f6

IP address (IPv4)
10.9.27.24

▼ Host 2

MAC Address *

IP address (IPv4)

Add another host

Figure A-5

e. Click “Add host” on Step 3 Host discovery.

1 Cluster details

2 Static network configurations >

3 Host discovery

4 Networking

5 Review and create

Host discovery

View cluster events

Add hosts

☐ Install OpenShift Virtualization ⓘ
Run virtual machines along containers.

☐ Install OpenShift Data Foundation ⓘ
Persistent software-defined storage for hybrid applications.

☐ Integrate with vSphere ⓘ

☒ Run workloads on control plane nodes ⓘ

Information and warnings

ⓘ Minimum hardware requirements ⓘ Hosts not showing up?

ⓘ All bootable disks will be formatted during installation

Host Inventory

☐ 0 selected

Actions

Hostna... ↑ Role ⓘ Status ⓘ Discove... ⓘ CPU Co... ⓘ Memory ⓘ Total st... ⓘ (0)

(:))

Waiting for hosts...

Hosts may take a few minutes to appear here after booting.

Figure A-6

- f. The web console will prompt a diagram for generating Discovery ISO image for provisioning the target node. For provisioning bare metal, we choose “Full image file”.

Add host [X]

Generate a Discovery ISO in order to add hosts to the cluster.

☐ Minimal image file: Provision with virtual media ?

☒ Full image file: Provision with physical media ?

SSH public key ?

Drag a file here or browse to upload Browse... Clear

Paste the content of a public ssh key you want to use to connect to the hosts into this field. [Learn more](#)

☐ Configure cluster-wide proxy settings

If hosts are behind a firewall that requires the use of a proxy, provide additional information about the proxy.

Generate Discovery ISO Cancel

Figure A-7

- g. Drag an SSH public key file or enter the content of an SSH public key from the bastion node to access the target node later. Then click “Generate Discovery ISO” button, the web console will embed the SSH key into the Discovery ISO image.

system

Host

Ad

☐ Ins
Ru

☐ Ins
Pe

Info

☒ Mi

☐ A

Host

Host

Add host ×

i Generate a Discovery ISO in order to add hosts to the cluster.

☐ Minimal image file: Provision with virtual media [?](#)

☒ Full image file: Provision with physical media [?](#)

SSH public key [?](#)

Drag a file here or browse to upload Browse... Clear

ssh-ed25519
AAAAC3N[REDACTED]nlutprZNqPPIMu2p721nld+0z
277RUZGOZhNuL amplab[REDACTED]

Paste the content of a public ssh key you want to use to connect to the hosts into this field. [Learn more](#)

☐ Configure cluster-wide proxy settings

If hosts are behind a firewall that requires the use of a proxy, provide additional information about the proxy.

Generate Discovery ISO Cancel

Figure A-8

- h. The web console will prompt the ISO URL or the wget command for downloading Discovery ISO image.

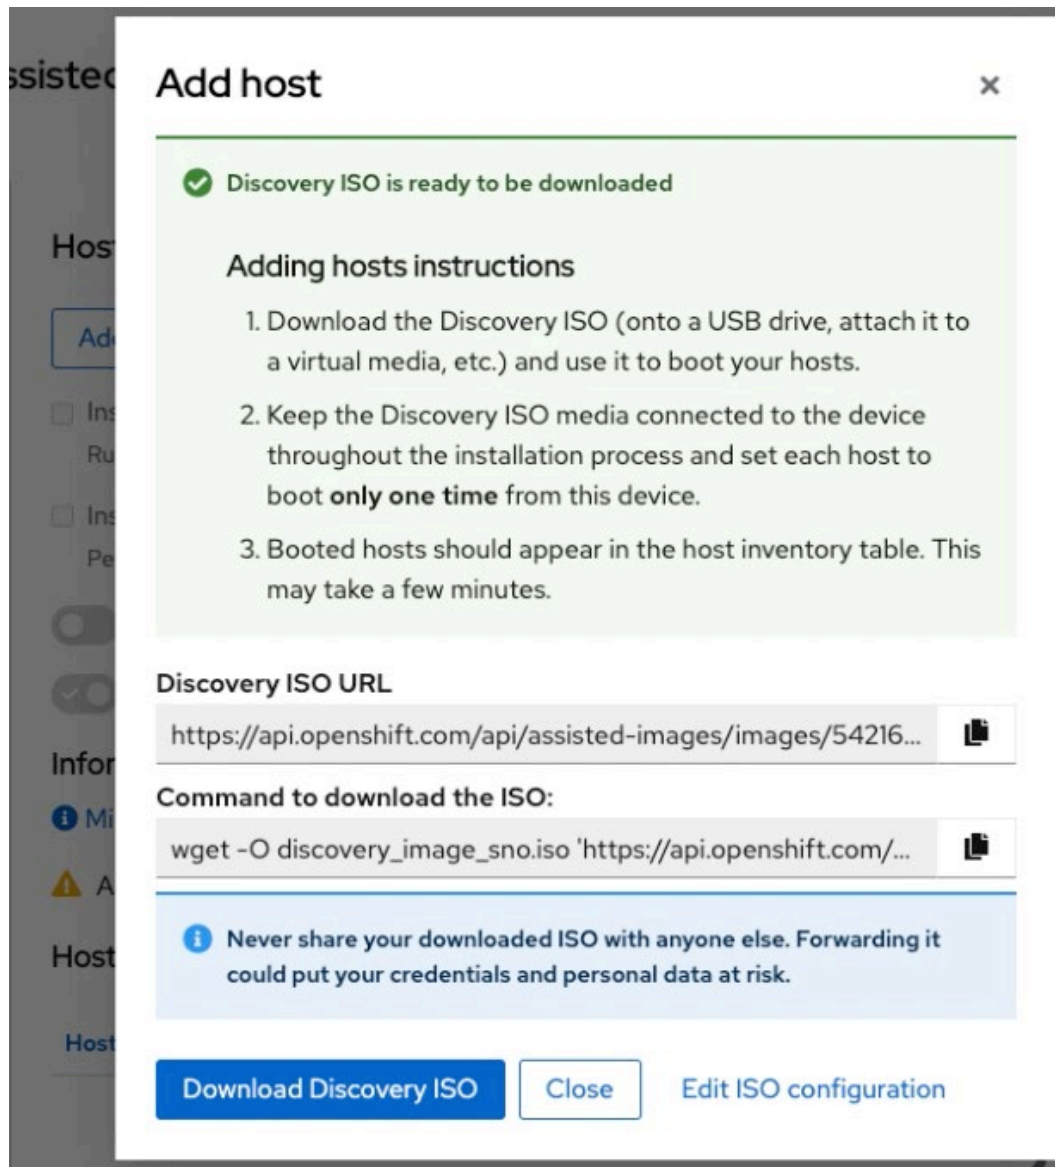


Figure A-9

4. Prepare and download Discovery ISO image for OCP 4.11

```
$ wget -O discovery_image_ocp4.iso 'https://api.openshift.com/api/assisted-images/images/[TOKEN1]?arch=arm64&image_token=[IMAGE_TOKEN]&type=full-iso&version=4.11'
```

5. Clean up the nodes for OCP 4.11 with Red Hat Enterprise Linux CoreOS Live CD. This step is optional but needed if your hardware has been used for other projects or provisioned with OS. Assume there are 9 NVMe drives (1x M.2 for OS and 8x U.2 for data storage) per node.

- a. Download the live CD

```
$ wget "https://mirror.openshift.com/pub/openshift-v4/arm64/dependencies/rhcos/4.11/latest/rhcos-4.11.0-aarch64-live.aarch64.iso"
```

- b. Mount Live CD with KVM on BMC. Once the system is loaded, enter the script below to clean up the drives, then power-off the system.

```
$ for DISK in "/dev/nvme0n1" "/dev/nvme1n1" "/dev/nvme2n1" "/dev/nvme3n1"  
"/dev/nvme4n1" "/dev/nvme5n1" "/dev/nvme6n1" "/dev/nvme7n1" "/dev/nvme8n1" ;  
do echo $DISK && \  
sgdisk --zap-all $DISK && \  
dd if=/dev/zero of=$DISK bs=1M count=100 oflag=direct,dsync && \  
blkdiscard $DISK  
done  
$ poweroff
```

6. Mount the Discovery ISO image for OCP 4.11 with KVM on BMC on the target node.
 - a. Click “Browser File” button on the top-right of KVM browser



Figure A-10. Mount Discovery ISO image

- b. The KVM dialog will prompt a window dialog for the target ISO image in a directory, then choose `discovery_image_sno.iso`



Figure A-11. Choose the ISO image

- c. Click “Start Media” button to mount ISO image to the target node



Figure A-12, Start to mount the ISO image

7. We use ipmitool SOL function to monitor the installation

```
ipmitool -H [BMC IP address] -I lanplus -U [username] -P [password] sol activate
```



Figure A-13, IPMITool SOL function shows the node's states

8. When all 3 nodes are loaded with Red Hat Enterprise Linux CoreOS (RHCOS) and update their status to <http://cloud.redhat.com> , the Assisted Installer web page will also show the status for each phase on the cluster.

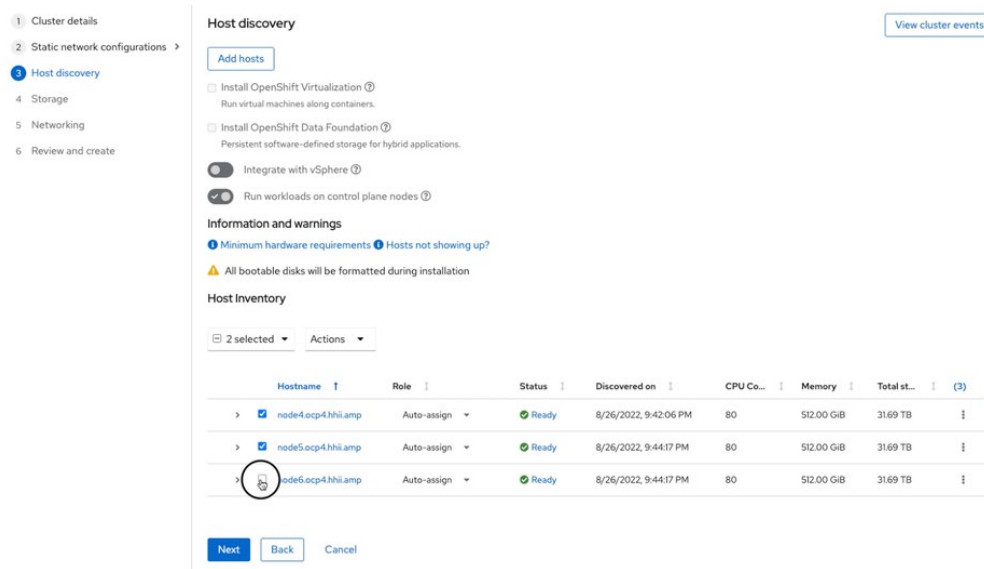


Figure A-14

- 1 Cluster details
- 2 Static network configurations
- 3 Host discovery
- 4 Storage
- 5 Networking
- 6 Review and create

Figure A-15

- 1 Cluster details
- 2 Static network configurations
- 3 Host discovery
- 4 Storage
- 5 **Networking**
- 6 Review and create

Figure A-16

11. Review the OpenShift cluster configuration and then click “Install cluster” to start the Assisted Installation.

Install OpenShift with the Assisted Installer

Review and create

Cluster address
ocp4.hhii.amp

OpenShift version
4.11.0

Stack type
IPv4

CPU architecture
arm64

0

Cluster summary

Hosts	3
Cores	240
Memory	1.50 TiB
Storage	95.06 TB

Cluster validations
✔ All validations passed

Host validations
✔ All validations passed

Cluster support level: Full
✔ Your installed cluster will be fully supported

Buttons: Install cluster, Back, Cancel

Figure A-17

12. The Assisted Installer page shows the nodes are preparing for installation as the role of “Control Plane node, Worker”, node5 also labels the role of “bootstrap”

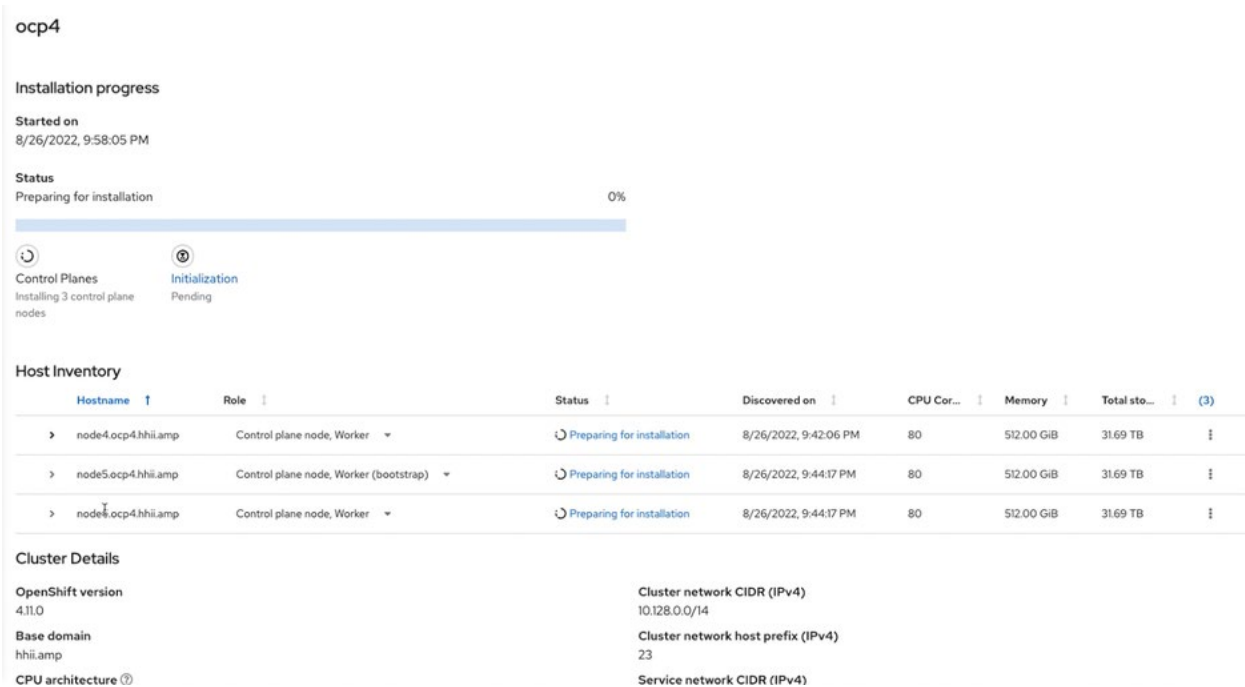


Figure A-18

13. You can click the link of “Status” to find out which step the Assisted Install is working on

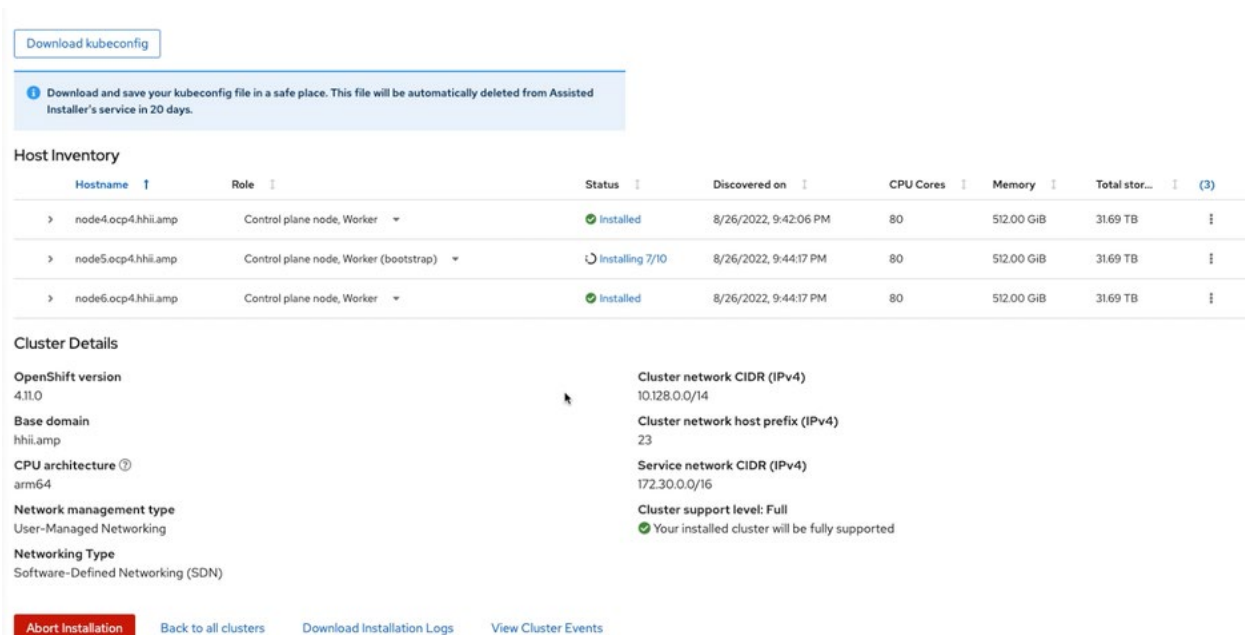


Figure A-19

14. After 35 ~ 65 minutes (depending on the Internet bandwidth), the OpenShift 4.11 3-node compact cluster is ready.

kubeadmin

Password

Download kubeconfig

Download and save your kubeconfig file in a safe place. This file will be automatically deleted from Assisted Installer's service in 20 days.

Host Inventory

Hostname	Role	Status	Discovered on	CPU Cores	Memory	Total storage	(3)
node4.ocp4.hhii.amp	Control plane node, Worker	Installed	8/26/2022, 9:42:06 PM	80	512.00 GiB	31.69 TB	
node5.ocp4.hhii.amp	Control plane node, Worker (bootstrap)	Installed	8/26/2022, 9:44:17 PM	80	512.00 GiB	31.69 TB	
node6.ocp4.hhii.amp	Control plane node, Worker	Installed	8/26/2022, 9:44:17 PM	80	512.00 GiB	31.69 TB	

Cluster Details

OpenShift version
4.11.0

Base domain
hhii.amp

CPU architecture
arm64

Network management type
User-Managed Networking

Networking Type
Software-Defined Networking (SDN)

Cluster network CIDR (IPv4)
10.128.0.0/14

Cluster network host prefix (IPv4)
23

Service network CIDR (IPv4)
172.30.0.0/16

Cluster support level: Full
Your installed cluster will be fully supported

Abort Installation Launch OpenShift Console Back to all clusters Download Installation Logs View Cluster Events

Figure A-20

15. You can access the OpenShift Console via the predefined URL: <https://console-openshift-console.ocp4.hhii.amp>

Not secure | https://console-openshift-console.apps.ocp4.hhii.amp/dashboards

Red Hat OpenShift

Administrator

Home

Overview

Projects

Search

API Explorer

Events

Operators

Workloads

Networking

Storage

Builds

Observe

Compute

User Management

Administration

Overview

Cluster

Getting started resources

Set up your cluster

Finish setting up your cluster with recommended configurations.

Configure alert receivers

View all steps in documentation

Build with guided documentation

Follow guided documentation to build applications and familiarize yourself with key features.

Monitor your sample application

Get started with Quarkus using a Helm Chart

View all quick starts

Explore new admin features

Explore new features and resources within the admin perspective.

API Explorer

OperatorHub

See what's new in OpenShift 4.11

Details

Cluster API address
https://api.ocp4.hhii.amp:6443

Cluster ID
774c94e-a7be-4f6d-9494-38fbc09bac2f

OpenShift Cluster Manager

Infrastructure provider
None

OpenShift version
4.11.0

Service Level Agreement (SLA)
Self-support, 60 day trial

20 days remaining

Manage subscription settings

Update channel
stable-4.11

Status

Cluster

Control Plane

Operators

Insights 2 issues found

Dynamic Plugins

Aug 10, 2022, 6:45 PM

Simple content access (SCA) is not enabled. Once enabled, Insights Operator can automatically import the SCA certificates from Red Hat OpenShift Cluster Manager making it easier to use the content provided by your Red Hat subscriptions when creating container images. See https://docs.openshift.com/container-platform/latest/cicd/builds/running-entitled-builds.html for more information.

Aug 10, 2022, 6:44 PM

System memory usage of 1294G on node5.ocp4.hhii.amp exceeds 95% of the reservation. Reserved memory ensures system processes can function even when the node is fully allocated and protects against workload out of memory events impacting the proper functioning of the node. The default reservation is expected to be sufficient for most configurations and should be increased (https://docs.openshift.com/container-

Activity

Ongoing

There are no ongoing activities.

Recent events

11:03 AM Stopping container registry-se...

11:02 AM Created container registry-ser...

11:02 AM Started container registry-serv...

11:02 AM Successfully pulled image 'regi...

11:02 AM Pulling image 'registry.redhat.L...

11:02 AM Add eth0 [10.130.0.223] fro...

11:02 AM Successfully assigned openshif...

11:02 AM Back-off pulling image "image-...

Figure A-21

Deploy Rook Ceph Operator

1. On the bastion node, use git to pull Rook Ceph Operator v.1.10.0

```
$ git clone --single-branch --branch v1.10.0 https://github.com/rook/rook.git
$ mv rook rook-v1.10.0
$ cd rook-1.10.0/deploy/examples/
```

2. Deploy CRDs, common, and operator yaml files

```
$ oc create -f crds.yaml -f common.yaml
$ oc create -f operator-openshift.yaml
$ oc create -f cluster.yaml
$ oc create -f toolbox.yaml
$ oc create -f csi/rbd/storageclass.yaml
$ oc create -f csi/rbd/snapshotclass.yaml
```

3. Set rook-ceph-block as the default StorageClass

```
$ oc patch storageclass rook-ceph-block -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

4. After a few minutes, the Ceph cluster is ready on rook-ceph project (namespace).

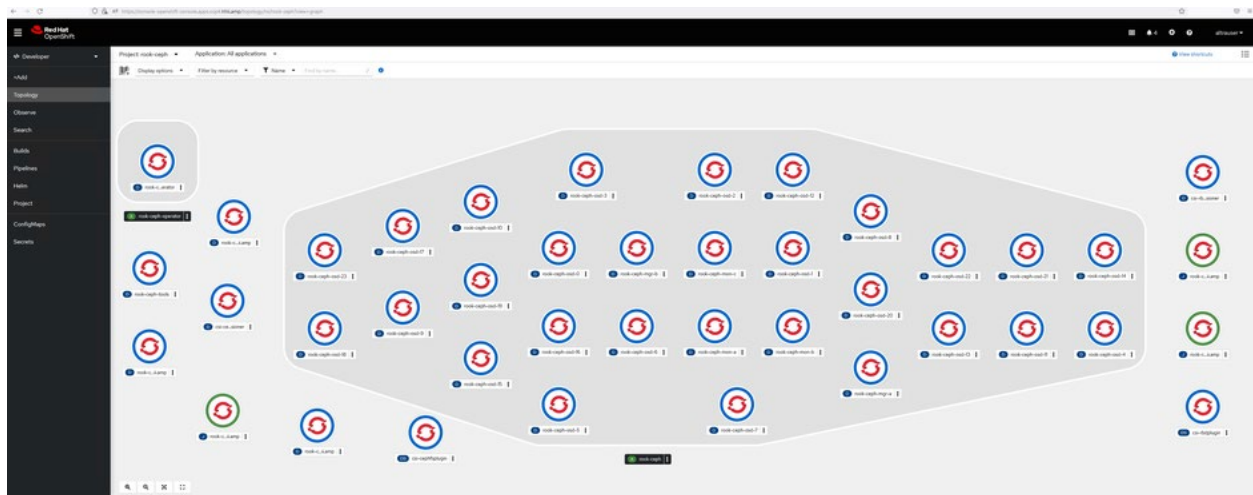


Figure A-22

5. When all OSDs are in ready status, create a YAML file (test-pvc.yaml) for a Persistent Volume Claim (PVC) under rook-ceph project (namespace) for testing Ceph block storage

```
kind: PersistentVolumeClaim
apiVersion: v1
```

```
metadata:
  name: test-pvc
  namespace: rook-ceph
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
  storageClassName: rook-ceph-block
  volumeMode: Filesystem
```

6. Create a test PVC by CLI

```
$ oc create -f test-pvc.yaml
```

7. Or create a PVC on OpenShift Console

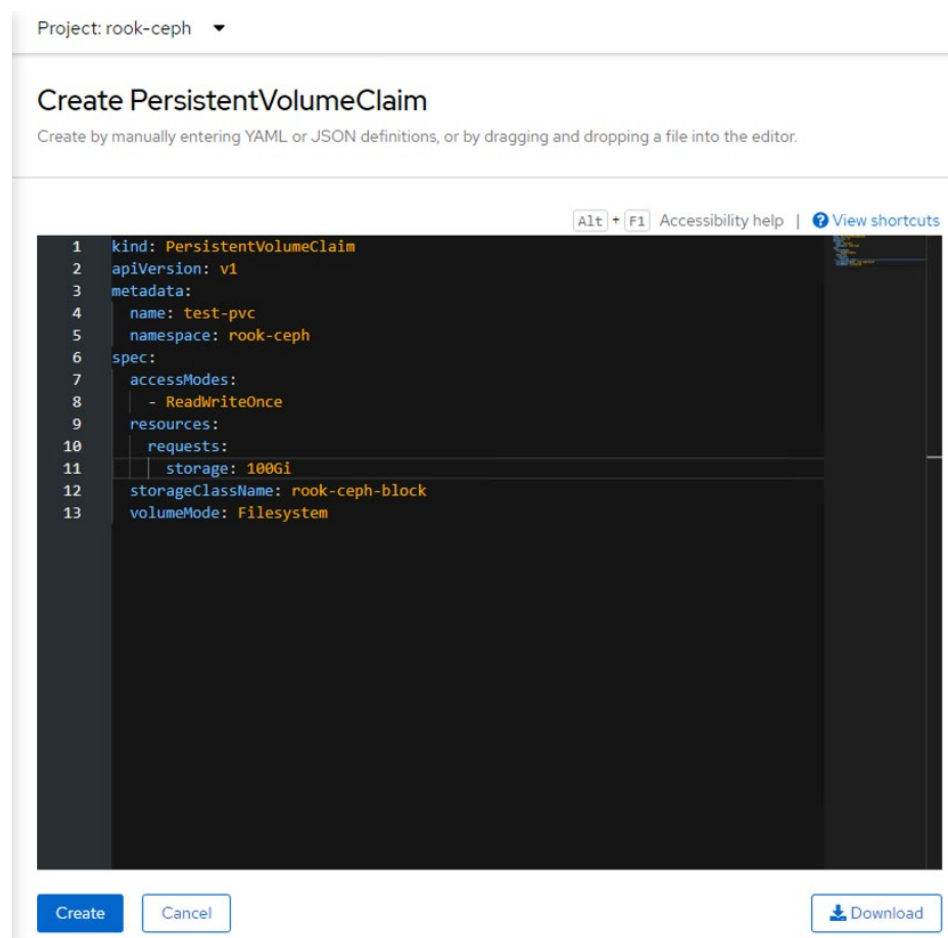


Figure A-23

8. Then verify its status and PVC details on OpenShift Console (Project → PersistentVolumeClaim → test-pvc)

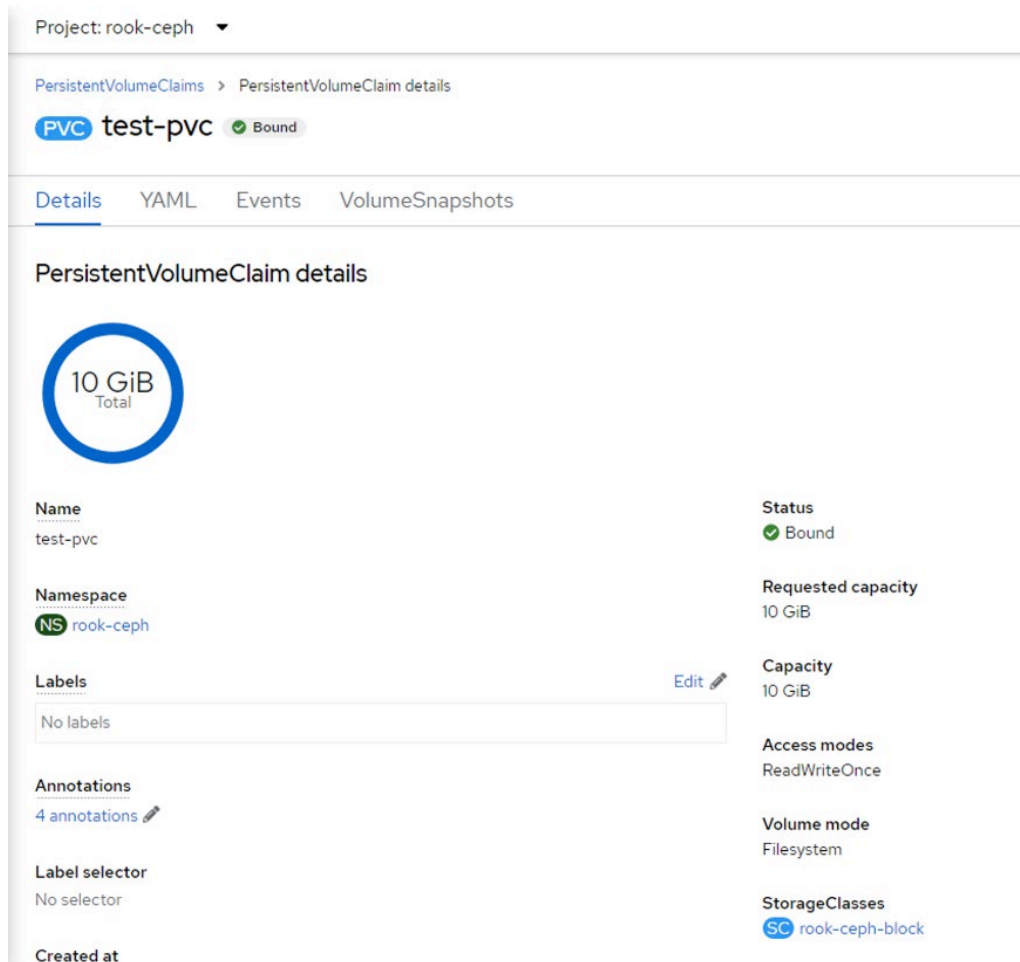


Figure A-24

9. Use ceph-tool pod to verify Ceph cluster's healthiness on bastion node

```
$ oc -n rook-ceph exec -it $(oc -n rook-ceph get pod -o name | egrep rook-ceph-tools ) -- /bin/bash
bash-4.4$ ceph status
cluster:
  id:          9b94ee02-a275-432d-bfef-fae1cf02a1aa
  health: HEALTH_OK

services:
  mon: 3 daemons, quorum a,b,c (age 40h)
  mgr: a(active, since 40h), standbys: b
  osd: 24 osds: 24 up (since 40h), 24 in (since 40h)
  rgw: 1 daemon active (1 hosts, 1 zones)

data:
  pools:   9 pools, 257 pgs
  objects: 411 objects, 928 KiB
  usage:   668 MiB used, 84 TiB / 84 TiB avail
  pgs:    257 active+clean
bash-4.4$ ceph health detail
HEALTH_OK
```

```
bash-4.4$ exit
```

10. Find the 3 YAML files, object-openshift.yaml, rgw-external.yaml, storageclass-bucket-delete.yaml, and change “my-store” to "object-store" in the 3 YAML files
11. Edit rgw-external.yaml for changing port number from 80 to 443 for OpenShift
12. Then deploy object storage related YAML files

```
$ oc create -f object-openshift.yaml
$ oc create -f storageclass-bucket-delete.yaml
$ oc create -f object-bucket-claim-delete.yaml
$ oc create -f rgw-external.yaml
```

13. The external URL Object storage service will be created like “<http://rook-ceph-rgw-object-store-rook-ceph.apps.ocp4.hhii.amp/>”
14. [Add the section for testing the access credential later]
15. Deploy service for Ceph dashboard

```
$ oc create -f dashboard-loadbalancer.yaml
```

16. Create a route for the external access to Ceph Dashboard web UI under “Routes” section on OpenShift Console

Project: rook-ceph ▾

Create Route

Routing is a way to make your application publicly visible.

Configure via: ☒ Form view ☐ YAML view

Name *
rook-ceph-mgr-dashboard-route
A unique name for the Route within the project.

Hostname
dashboard-rook-ceph.apps.ocp4.hhii.amp
Public hostname for the Route. If not specified, a hostname is generated.

Path
/
Path that the router watches to route traffic to the service.

Service *
rook-ceph-mgr-dashboard ▾
Service to route to.

[+ Add alternate Service](#)

Target port *
8443 → 8443 (TCP) ▾
Target port for traffic.

[Create](#) [Cancel](#)

Figure A-25

17. Check the Topology page again to review the external URLs for the endpoints of object-store pod and Ceph Cluster Dashboard (on rook-ceph-mgr-a pod)

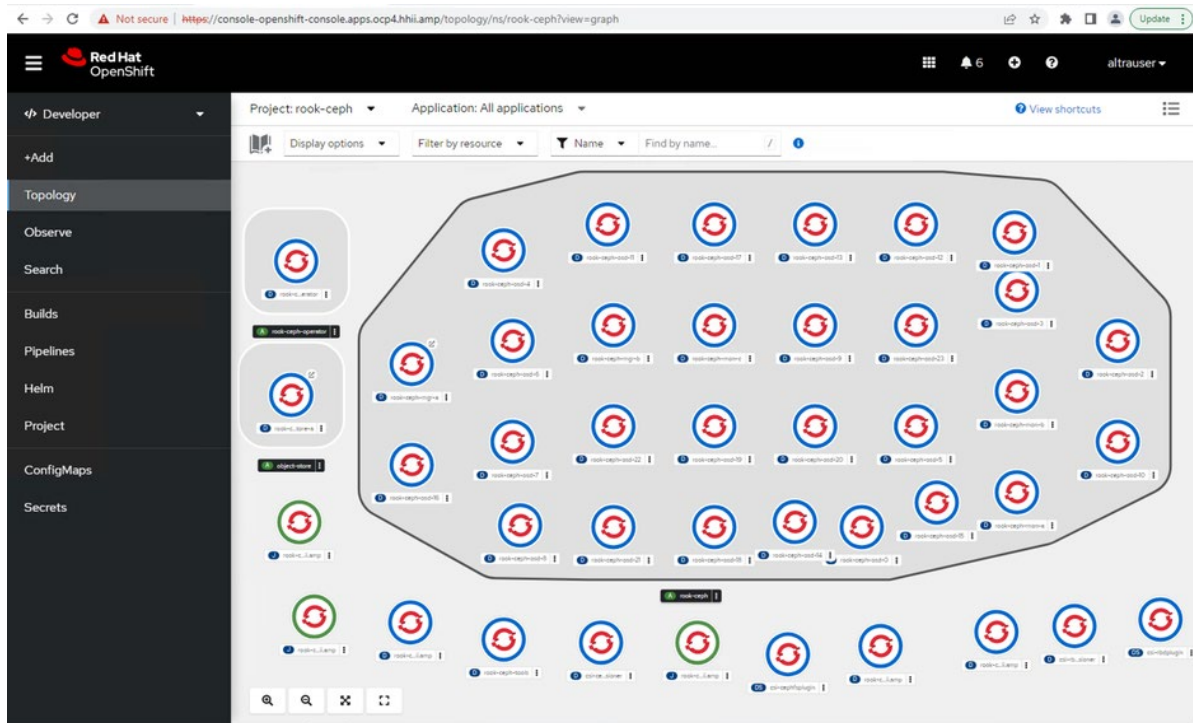


Figure A-26

18. Obtain the generated password for Ceph Dashboard via the command below

```
$ oc -n rook-ceph get secret rook-ceph-dashboard-password -o  
jsonpath="{['data']['password']}" | base64 --decode && echo  
aaabbb123xxxxyy
```

19. Click the link to access Ceph Dashboard with the credential for “admin”, the web UI will look like the one below

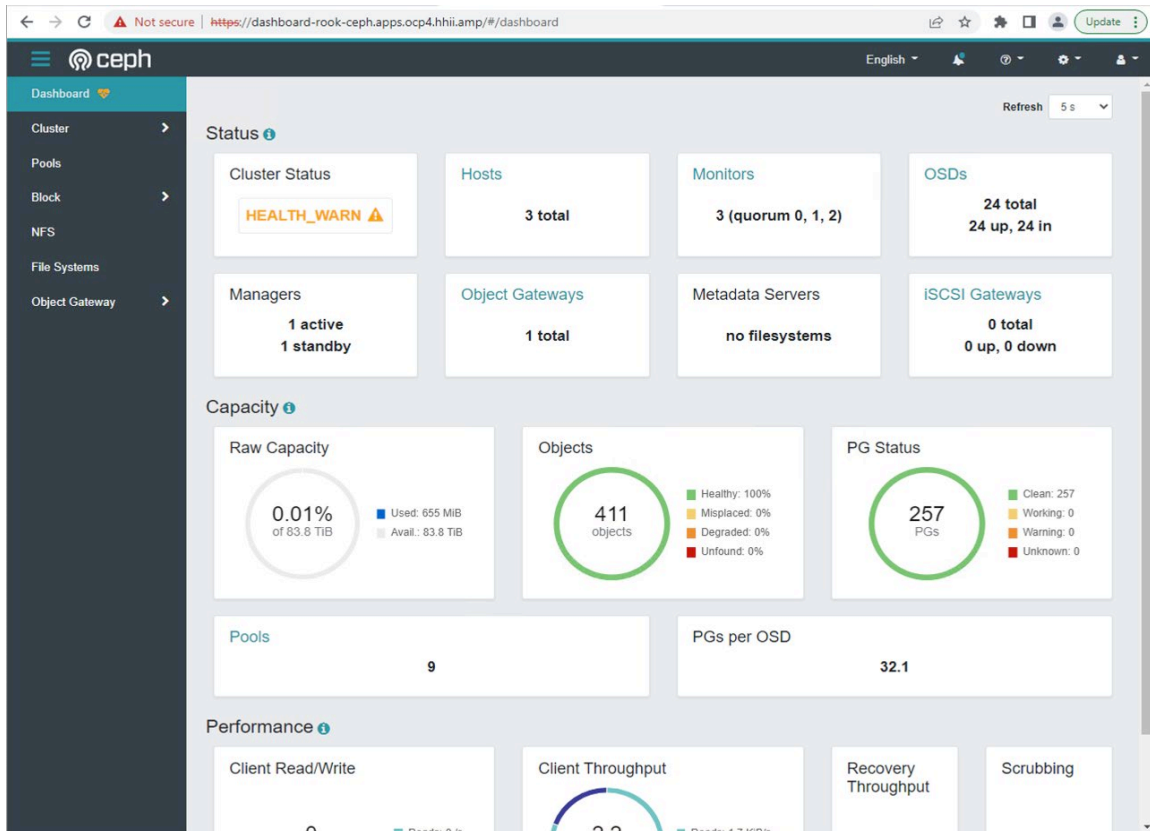


Figure A-27

20. The 3-node compact OpenShift Cluster with Ceph Block storage and Object storage is ready for your next project!