



Deploying the Cisco Application Services Engine in KVM (Fabric External Mode)

- [Prerequisites, on page 1](#)
- [Deploying the Cisco Application Services Engine in KVM, on page 1](#)

Prerequisites

Complete the following one time pre-requisites before you start:

- Deploying Cisco Application Services Engine is supported on Linux operating system such as CentOS, Ubuntu or RedHat.
- Ensure that the following minimum Kernel and Virsh requirements are met:
 - Linux Kernel: 3.10.0-957.el7.x86_64
 - Virsh: libvirt-4.5.0-23.el7_7.1.x86_64

- Each cluster node requires a dedicated disk partition and minimum of 800 GB of disk space.

- The disk must have I/O latency less than 20ms. Assuming, /home is the disk/partition then,

```
# mkdir /home/test_data
```

```
# fio --rw=write --ioengine=sync --fdatasync=1 --directory=test_data_with_se --size=22m --bs=2300 --name=mytest
```

Check for **99.00th**=[<VALUE>] under **fsync/fdatasync/sync_file_range** section. It should be less than 20 ms.

- Memory: 48G for each service node
- vCPUs: 16 for each service node
- You have installed all the required packages for QEMU-KVM support.

Deploying the Cisco Application Services Engine in KVM

This procedure is used for setting up Cisco Application Services Engine cluster in Linux KVM.

Step 1 Choose the Cisco Application Services Engine ISO image image.

- a) Browse to the [Software Download](#) page.
- b) Choose the Cisco Application Services Engine image for KVM (apic-sn-dk9.1.1.2(x).qcow2).

Step 2 Create a directory for service node base qcow2 image and download the **apic-sn-dk9.1.1.2h.qcow2** file.

Note Execute this on all the KVM hosts for all cluster nodes.

Note Each node needs to have its qcow2 path on a unique disk partition.

```
[ node1 ] # mkdir -p /home/sn_base/qcow2
[ node1 ] # cd /home/sn_base/qcow2
[ node1 ] # <wget/scp file from CCO to this location>
[ node1 ] # ls
apic-sn-dk9.1.1.2h.qcow2
[ node1 ] #
```

Step 3 Create a directory for the data path for the service node on each host and create a snapshot of the base image. The service node will always write to this snapshot.

Note Perform this action on all the service nodes in the cluster.

```
[ node1 ] # mkdir -p /home/mso-node1/
[ node1 ] # cd /home/mso-node1
[ node1 ] # qemu-img create -f qcow2 -b /home/sn_base/qcow2/apic-sn-dk9.1.1.2h.qcow2
/home/mso-node1/disk0.qcow2
```

Step 4 Open the KVM console and click **New Virtual Machine**.

Step 5 On the **New VM**, choose **import existing disk image option**. Click **Forward**.

Step 6 In the **provide existing storage path** tab, choose the **/home/mso-node1/disk0.qcow2** file.

Note Each node needs to have its qcow2 path on a unique disk partition.

Step 7 Choose the **Generic** value for the operating system and the version. Click **Forward**.

Step 8 For memory, choose the value 48000. For CPU, choose the value 16. Click **Forward**.

Step 9 Enter the name of the virtual machine **mso-node1**. Select the **customize configuration before install**. Choose the appropriate option from **Network selection** and click **Finish**.

Step 10 In the window **mso-node1 on QEMU/KVM**, choose the appropriate option from **Network selection**.

- a) Select the NIC for the **Virtual Network Interface** and choose the device model as **e1000**.
- b) Leave the default **Mac address**.
- c) Click **Apply**.
- d) Click **Begin Installation**.

The virtual machine should boot from **disk0.qcow2**. The first-boot prompt is displayed.

- a) Specify the mode. To specify that the configuration is not obtained from the Cisco APIC cluster enter **n**.
- b) Enter the serial number and a unique hostname for the service node.
- c) Enter the domain name for the service node. The domain name is equivalent to the name of the cluster or the domain name of the fabric.

```
Setup utility for apic-sn with SerialNumber Not Specified and running version 1.1.2h
Is this running in ACI mode? (y/n) n
Enter node serialnumber: Mynode01
Enter node hostname: mso-node1
```

```
Enter node domain: example.com
Enter the password for rescue-user:
Reenter the password for rescue-user:
```

Step 11 Enter the physical network management IP address and mask.

It is the out-of-band management IPv4/ or Pv6 addresses used to access the Cisco Application Services Engine GUI, CLI, or API.

```
Enter physical network management IP address and mask:192.168.10.100/24
```

Step 12 Enter the physical network gateway IP address.

It is used for communicating to the external networks using out-of-band management.

```
Enter physical network gateway IP address:192.168.10.1
```

Step 13 Enter the number of masters in the cluster.

```
Enter number of Masters in the cluster (recommended is 3) 3
```

Step 14 Enter IP addresses, serial number of other master nodes in the cluster.

If the cluster size is 1, leave it blank.

```
Enter details of other Masters in the cluster, one at a time?
Select 'n' for a space-separated list (y/n) y
1) Enter IP Address: 192.168.10.101
Enter SerialNumber: Mynode02
2) Enter IP Address: 192.168.11.102
Enter SerialNumber: Mynode03
```

Step 15 You must assign one node in the cluster as the first master. If the cluster already exists, enter **n**.

```
Is this the first node in a new cluster? (y/n) y
```

Step 16 Enter the application overlay network IP address and mask.

It is the private IP address block, /16 network is required for container or pod network.

```
Enter application overlay network IP address and mask: 1.1.0.0/16
```

Step 17 Enter the service network IP address and mask.

It is the private IP address block, /16 network that is required for container or pod network.

```
Enter service network IP address and mask: 2.2.0.0/16
```

Step 18 Enter the search domain.

```
Enter the search domain as a space-separated list: mydomain.com
```

Step 19 Enter the addresses of the DNS name servers.

It is the IP address list required for resolving DNS names outside the cluster.

```
Enter nameserver addresses as a space-separated list: 192.168.12.100 192.168.12.101
```

Step 20 Enter the IP addresses of the NTP servers.

Enter the IP address of the NTP servers. It is required to sync the clock between all the master nodes in the cluster.

```
Enter the ntp servers as a space-separated list: 192.168.13.101
```

Step 21 Review the configuration.

```

Please review the config:
Number Masters cluster: 3
application overlay network: 1.1.0.0/16
first Master: true
management IP: 192.168.10.100/24
nameservers list: [192.168.12.100 192.168.12.101]
node domain: example.com
node hostname: mso-node1
node serialnumber: Mynode01
ntp servers list: [192.168.13.101]
physical gateway IP: 192.168.10.1
rescue-user password: <hidden>

search list: [mydomain.com]
seed list:
- {ipAddress: 192.168.10.101, name: mso-node02, serialNumber: Mynode02}
- {ipAddress: 192.168.11.102, name: mso-node03, serialNumber: Mynode03}
service network: 2.2.0.0/16
Do you wish to reenter the bootstrap config? (y/N) N

mso-node1 login:

```

Step 22 Generate the dbgtoken

a) Log in to SSH

```

$ssh rescue-user@192.168.10.100
password:
bash-4.2$ acidiag dbgtoken
0M080NDSGPRH
bash-4.2$

```

Step 23 Configure the second node similarly.

```

Please review the config:
Number Masters cluster: 3
application overlay network: 1.1.0.0/16
first Master: false
management IP: 192.168.10.101/24
nameservers list: [192.168.12.100 192.168.12.101]
node domain: example.com
node hostname: mso-node2
node serialnumber: Mynode02
ntp servers list: [192.168.13.101]
physical gateway IP: 192.168.10.1
rescue-user password: <hidden>
search list: [mydomain.com]
seed list:
- {ipAddress: 192.168.10.100, name: mso-node01, serialNumber: Mynode01}
- {ipAddress: 192.168.11.102, name: mso-node03, serialNumber: Mynode03}
service network: 2.2.0.0/16
Do you wish to reenter the bootstrap config? (y/N) N
Enter the latest dbgtoken from other active node in the cluster: 0M080NDSGPRH
mso-node2 login:

```

Step 24 For the **Enter the latest dbgtoken from other active node in the cluster**, go to Step 22. Obtain the dbgtoken by logging into the first node using the SSH. Enter the value for node 2.

Note: Always use the latest dbgtoken from the SSH to log in to the nodes.

Step 25 Configure the third node.

```

Please review the config:
Number Masters cluster: 3
application overlay network: 1.1.0.0/16

```

```
first Master: false
management IP: 192.168.11.102/24
nameservers list: [192.168.12.100 192.168.12.101]
node domain: example.com

node hostname: mso-node3
node serialnumber: Mynode03
ntp servers list: [192.168.13.101]
physical gateway IP: 192.168.11.1
rescue-user password: <hidden>
search list: [mydomain.com]
seed list:
- {ipAddress: 192.168.10.100, name: mso-node01, serialNumber: Mynode01}
- {ipAddress: 192.168.10.101, name: mso-node02, serialNumber: Mynode02}
service network: 2.2.0.0/16
Do you wish to reenter the bootstrap config? (y/N) N
Enter the latest dbgtoken from other active node in the cluster: 0M080NDSGPRH
mso-node3 login:
```

Step 26 For the **Enter the latest dbgtoken from other active node in the cluster**, go to Step 22. Obtain the dbgtoken by logging into the first node using the SSH. Enter the value for node 3.

Note: Always use the latest dbgtoken from the SSH to log in to the nodes.

Step 27 After all three nodes are bootstrapped, wait for 15-30 mins and execute the following command using SSH:

```
Server # acidiag health
cluster is healthy
```

Verify that a “healthy” status is displayed to indicate that the installation was performed successfully.

Step 28 Cisco Application Services Engine is available to deploy the apps that can be hosted on the Cisco Application Services Engine.

Note Cisco Application Services Engine, Release 1.1.2 supports the deployment of only the Cisco ACI Multi-Site Orchestrator application (starting with Release 2.2(3)). Refer to the [ACI Multi-Site Orchestrator Cisco ACI Multi-Site Orchestrator Installation and Upgrade Guide](#) for more information
