



Deploying in VMware ESX

- [Prerequisites and Guidelines, on page 1](#)
- [Deploying Cisco Nexus Dashboard in VMware ESX, on page 4](#)

Prerequisites and Guidelines

Virtual deployments are supported starting with Nexus Dashboard, Release 2.0.2h. Earlier releases support only the physical form factor described in [Deploying as Physical Appliance](#).

Before you proceed with deploying the Nexus Dashboard cluster in VMware ESX, you must:

- Review and complete the general prerequisites described in the [Deployment Overview](#).

Note that this document describes how to initially deploy a three-node Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `worker` or `standby`), see the "Deploying Additional Nodes" section of the *Cisco Nexus Dashboard User Guide* instead.

The guide is available from the Nexus Dashboard UI or online at [Cisco Nexus Dashboard User Guide](#)

- Ensure that the ESX form factor supports your scale and application requirements.

Scale and application co-hosting vary based on the cluster form factor. You can use the [Nexus Dashboard Capacity Planning](#) tool to verify that the virtual form factor satisfies your deployment requirements.

- Ensure you have enough system resources:

Table 1: Deployment Requirements

Nexus Dashboard Version	Requirements
Release 2.0.2h Earlier releases are not supported.	<ul style="list-style-type: none"> • VMware vCenter 6.x • VMware ESXi 6.5 or 6.7 • Each VM requires: <ul style="list-style-type: none"> • 16 vCPUs • 64 GB of RAM • 500 GB disk • We recommend that each Nexus Dashboard node is deployed in a different ESXi server.

- After each node's VM is deployed, ensure that the VMware Tools periodic time synchronization is disabled as described in the deployment procedure in the next section.

ESX Host Network Connectivity


If you plan to install Nexus Dashboard Insights or Fabric Controller service and use the Persistent IPs feature, you must ensure that the ESX host where the cluster nodes are deployed has a single logical uplink. In other words, it is connected via a single link, PC, or vPC and not a dual Active/Active (A/A) or Active/Standby (A/S) link without PC/vPC.

The following diagrams summarize the supported and unsupported network connectivity configurations for the ESX host where the nodes are deployed:

- In case the ESX host is connected directly, the following configurations are supported:
 - A/A uplinks of Port-Group or virtual switch with PC or vPC
 - Single uplink of Port-Group or virtual switch
 - Port-Channel used for the uplink of Port-Group or virtual switch.

A/A or A/S uplinks of Port-Group or virtual switch without PC or vPC are not supported


Figure 1: ESX Host Connectivity (Direct)



- In case the ESX host is connected via a UCS Fabric Interconnect (or equivalent), the following configurations are supported:
 - A/S uplinks of Port-Group or virtual switch at UCS Fabric Interconnect level without PC or vPC
 In this case, the Active/Standby links are based on the server technology, such as Fabric Failover for Cisco UCS and not at the ESXi hypervisor level.
 - Single uplink of Port-Group or virtual switch

A/A or A/S uplinks of Port-Group or virtual switch at the hypervisor level without PC or vPC are not supported

Figure 2: ESX Host Connectivity (with Fabric Interconnect)



Deploying Cisco Nexus Dashboard in VMware ESX

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter.

Before you begin

- Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines, on page 1](#).

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

<https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html>

b) Click the **Downloads** tab.

c) Choose the Nexus Dashboard version you want to download.

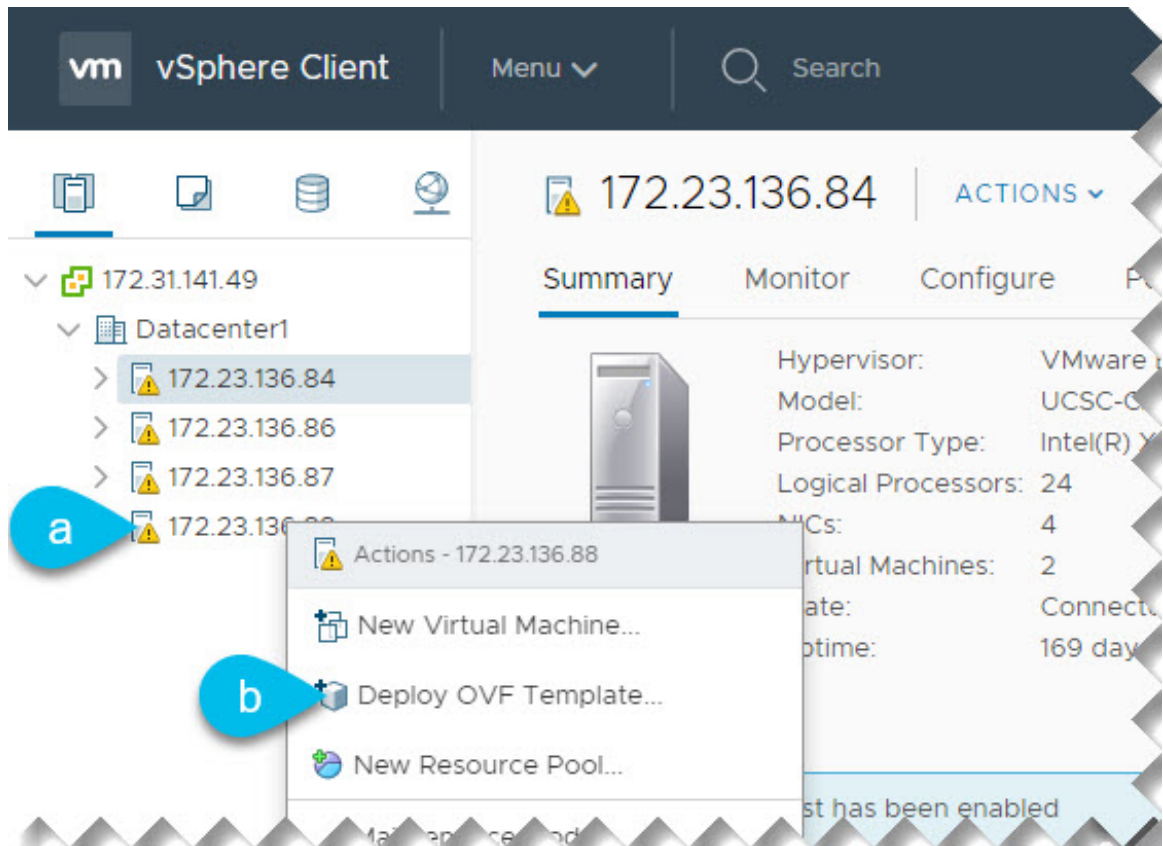
d) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.ova`).

Step 2 Log in to your VMware vCenter.

You cannot deploy the OVA directly in the ESX host, you must deploy it using the vCenter.

Note Depending on the version of your vSphere client, the location and order of configuration screens may differ slightly. The following steps provide deployment details using VMware vSphere Client 6.7.

Step 3 Start the new VM deployment.



- a) Right-click the ESX host where you want to deploy.
- b) Then select **Deploy OVF Template...**

The **Deploy OVF Template** wizard appears.

Step 4

In the **Select an OVF template** screen, provide the OVA image location.

Deploy OVF Template

1 Select an OVF template | Select an OVF template

2 Select a name and folder | Select an OVF template from remote URL or local file system

3 Select a compute resource |

4 Review details | Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

5 Select storage | URL

6 Ready to complete | Local file

Local file

nd-2.0.1.2a.ova

- a) Select **Local file** and click **Choose Files** to select the OVA file you downloaded..
- b) Click **Next** to continue.

Step 5

In the **Select a name and folder** screen, provide a name and location for the VM.

Deploy OVF Template

✓ 1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and select location

Virtual machine name:


Select a location for the virtual machine.

- > [Folder]
- > [Folder]
- > Fabric 2

CANCEL BACK NEXT

- Provide the name for your virtual machine.
- Select the location for the virtual machine.
- Click **Next** to continue


Step 6 In the **Select a compute resource** screen, select the ESX host.



- Select the vCenter datacenter and the ESX host for the virtual machine.
- Click **Next** to continue

Step 7 In the **Review details** screen, click **Next** to continue.

Step 8 In the **Select storage** screen, provide the storage information.



a) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

b) From the **Select virtual disk format** dropdown, select **Thick Provision Lazy Zeroed**.

c) Click **Next** to continue

Step 9

In the **Select networks** screen, accept default values and click **Next** to continue.

There are two networks, **fabric0** is used for the data network and **mgmt0** is used for the management network.

Step 10

In the **Customize template** screen, provide the required information.

Note The following few steps may be listed in different order depending on the version of the vSphere client you are using. The provided order and examples are using VMware vSphere 6.7.

In the **Resource Configuration** and **Node Configuration** categories, provide the following details:

Deploy OVF Template

Customize template
Customize the deployment properties of this software solution.

Category	Number of settings
Resource Configuration	1 settings
1. Data Disk Size (GB)	Data disk size (min 300GB, max 1536GB (1.5TB)) 300
Node Configuration	3 settings
1. Node Name	Host name of the node nd-node1
2. Password	Local "rescue-user" password Password: Confirm Password:
3. Role	Node role Master

- a) Provide the sizes for the node's data disks.

We recommend using the default values for the required data disks.

- b) Provide the **Node Name**.

This will be the hostname for node, do not use the fully qualified domain name (FQDN).

For example, `nd-node1`

- c) Provide and confirm the **Password**.

We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node. If you provide different passwords, the first node's password will be used as the initial password of the `admin` user in the GUI.

- d) From the **Role** dropdown, select `Master`.

When first deploying the cluster, all 3 nodes must be `Master`. Adding `Worker` and `Standby` nodes is described in the *Cisco Nexus Dashboard User Guide*.

In the **Network Configuration** category, provide the following details:

Deploy OVF Template

Customize template
Customize the deployment properties of this software solution.

Category	Number of settings
Network Configuration	5 settings
1. Management Network Address and subnet	Management network address. Enter IP/subnet 192.168.10.11/24
2. Management Gateway IP	Management network gateway IP address. Enter IP only 192.168.10.1
3. Data Network Address and subnet	Data network address. Enter IP/subnet 172.10.10.11/24
4. Data Network Gateway IP	Data network gateway IP address. Enter IP only 172.10.10.1
5. Data Network Vlan	Data Network Vlan ID (Optional), leave it empty or set to 0 if no vlan

- a) Provide the **Management Address and Subnet** for the node.

The management IP address can be in the same or different subnet as the data network IP address.

For example, 192.168.10.11/24.

- b) Provide the **Management Gateway IP**.

For example, 192.168.10.1.

- c) Provide the **Data Network Address and subnet**.

The data network IP address can be in the same or different subnet as the management IP address.

For example, 172.10.10.11/24.

- d) Provide the **Data Network Gateway**.

For example, 172.10.10.1.

- e) (Optional) If the data traffic is on a VLAN, provide the **Data Network Vlan**.

For most deployments, you can leave this field blank. If you do want to provide a VLAN ID for your data network, you can enter it in this field, for example 100.

In the **Cluster Configuration Mandatory** and **Cluster Configuration Optional** categories, provide the following details:

Deploy OVF Template

<ul style="list-style-type: none"> ✓ 1 Select an OVF template ✓ 2 Select a name and folder ✓ 3 Select a compute resource ✓ 4 Review details ✓ 5 Select storage ✓ 6 Select networks 7 Customize template 8 Ready to complete 	<ul style="list-style-type: none"> Cluster Configuration Mandatory 4 settings 	
a	1. Cluster Name	Name of the Cluster nd-cluster
b	2. Master List	List the Data Network IPs of _the other_ master nodes in the cluster, separated by spaces. (Ex: 192.192.100.102 192.192.100.103) 172.10.10.12 172.10.10.13
c	3. Enter the latest dbgtoken from the master node in the cluster	Enter the latest dbgtoken from the master node in the cluster. For master node enter some string of at-least length 11 (ignored internally) abcdefg1234
d	4. Download Config From Peers	Download Config From Peers and skip Optional Config Below? <input type="checkbox"/>
	Cluster Configuration Optional	5 settings
e	1. App Subnet	Application Network IP subnet. Enter IP/subnet 172.17.0.1/16
f	2. Service Subnet	Service Network IP subnet. Enter IP/subnet 100.80.0.0/16
g	3. NTP Servers	List of IPs of NTP servers, separated by space 10.197.145.2 10.197.146.2
h	4. Name servers	List of IPs of Name servers, separated by space 10.197.145.3
i	5. Search Domains	List of DNS domains to search, separated by space company.com

CANCEL BACK NEXT

- a) Provide the **Cluster Name** for the Nexus Dashboard cluster.

This name must be the same for all nodes.

For example, `nd-cluster`.

- b) In the **Master List** field, provide the data network IP addresses of the other 2 nodes you will configure for your cluster.

Each IP address in the list must be separated by a space.

For example, if the data network IP addresses of all 3 nodes are `172.10.10.11`, `172.10.10.12`, and `172.10.10.13`, the value of this field for the first node would be `172.10.10.12 172.10.10.13`

- c) Provide a value for the **dbgtoken** field.

Since this is the first node you are deploying, provide any 11-character value for this field (for example, `abcdef12345`). When you deploy the other two nodes, you will use this field to provide a token from the first node to simplify configuration.

- d) Leave the **Download Config From Peers** checkbox unchecked.

You will use this option when configuring the other two nodes.

- e) Provide the **App Subnet**.

The application overlay network defines the address space used by the application's services running in the Nexus Dashboard.

The field is pre-populated with the default `172.17.0.1/16` value.

- f) Provide the **Services Subnet**.

The services network is an internal network used by the Nexus Dashboard and its processes.

The field is pre-populated with the default `100.80.0.0/16` value.

- g) Provide the **NTP Servers** information.

For example, `10.197.145.2 10.197.146.2`.

- h) Provide the **Name servers** information.

For example, `10.197.145.3`.


- i) (Optional) Provide the **Search Domains** information.

For example, `company.com`.

Step 11

Verify that all information is valid and click **Next** to continue.

After you complete the **Customize template** screen, a verification banner is shown at the top.



- Step 12** In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the first node.
- Step 13** Wait for the VM deployment to complete, ensure that the VMware Tools periodic time synchronization is disabled, then start the VM.
- To disable time synchronization:
- Right-click the node's VM and select **Edit Settings**.
 - In the **Edit Settings** window, select the **VM Options** tab.
 - Expand the **VMware Tools** category and uncheck the **Synchronize guest time with host** option.
- Step 14** Log in to the first node's console as the `rescue-user`.
- Use the password you specified in the OVF template when deploying the VM.
- Step 15** Retrieve the `dbgtoken`.
- Run the following command:
- ```
$ acs debug-token
09GZ1PMB8CML
```
- Make a note of this token, you will use it to deploy the other two nodes.
- Keep in mind, the token expires and is refreshed every 30 minutes, so ensure to retrieve it when ready to deploy the second and third nodes.
- Step 16** Deploy the second node.
- The steps to deploy the second and third nodes are similar, with the exception that you can now use the `dbgtoken` from the first node to skip some of the configuration.
- Repeat Steps 2 through 9 to start deploying the 2nd node.  
We recommend using a different ESX host for each node.
  - In the **Cluster Configuration** screen, provide the following information:
    - **Node Name**  
Do not use the fully qualified domain name (FQDN).
    - **Password**  
We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node. If you provide different passwords, the first node's password will be used as the initial password of the `admin` user in the GUI.
    - **Role**  
When first deploying the cluster, all 3 nodes must be `Master`.
    - **Management Network Address and subnet**
    - **Management Gateway IP**
    - **Data Network Address and subnet**
    - **Data Network Gateway**
    - (Optional) If the data traffic is on a VLAN, provide the **Data Network Vlan**.

- **Cluster Name**

This name must be the same for all nodes. For example, `nd-cluster`.

- **Master List**

Provide the data network IP addresses of the other 2 nodes in your cluster separated by a space.

For example, if the data network IP addresses of all 3 nodes are `172.10.10.11`, `172.10.10.12`, and `172.10.10.13`, the value of this field for the second node would be `172.10.10.11 172.10.10.13`

- Provide the **dbgtoken** you obtained from the first node.

The token expires and is refreshed every 30 minutes, ensure to obtain the latest valid token from the first node before continuing. For example, `09GZ1PMB8CML`.

- Check the **Download Config From Peers**

The second and third nodes will download common configuration parameters from the first node using the `dbgtoken`.

- Skip **Cluster Configuration Optional** fields and click **Next** to continue.
- In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the second node.

**Step 17** Repeat the previous step to deploy the third node.

**Step 18** Wait for the second and third node VMs deployment to complete, then start the VMs.

**Step 19** Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

- Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress

$ acs health
k8s services not in desired state - [...]

$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

- Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

---

