



Install CWM using OVA

This section contains the following topics:

- [Install CWM using OVA, on page 1](#)

Install CWM using OVA

The Crosswork Workflow Manager 1.1 is installed as a guest virtual machine by deploying an OVA image using the VMware vSphere 6.7 (and higher) virtualization platform.

Prerequisites

- An **ed25519** SSH public and private key pair.

System requirements

| Minimum system requirements | |
|-----------------------------|--|
| Server | VMware vSphere 6.7+ account with an ESXi 6.7+ host |
| CPU | 8 cores |
| Memory | 64 GB |
| Storage | 100 GB |

Download the CWM package

To get the CWM 1.1 software package:

- Step 1** Go to the Cisco Software Download service and in the search bar, type in '**Crosswork Workflow Manager**', then select it from the search list.
- Step 2** From Select a software type, select **Crosswork Workflow Manager Software**.
- Step 3** Download the Crosswork Workflow Manager software package for Linux.

Step 4 In a terminal, use the `sh` command to extract the downloaded **.signed.bin** file and verify the certificate. See example output below for reference:

```
sh cwm-1.1.signed.bin
Unpacking...
Verifying signature...
Retrieving CA certificate from http://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from http://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from tailf.cer.
Successfully verified the signature of cwm-1.1.tar.gz using tailf.cer
```

The `cwm-1.1.tar.gz` file and other files have been extracted and validated against the signature file.

Step 5 To extract the `cwm-1.1.tar.gz` file, double click on it (Mac users) or use `gzip` utility (Linux and Windows users). This will extract the CWM OVA file that will be used for installation.

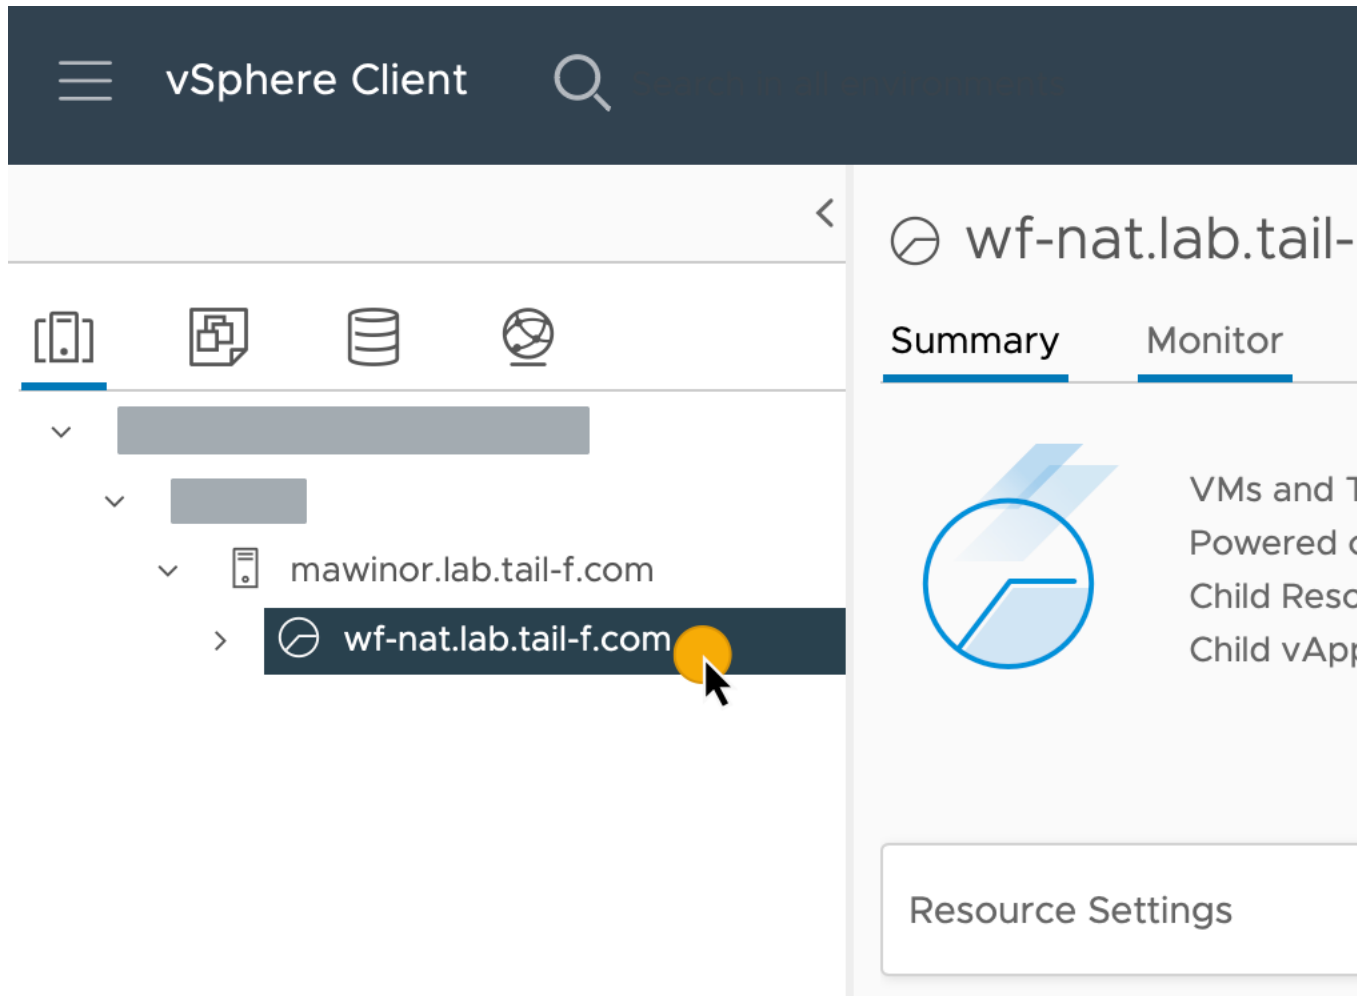
Deploy OVA and start VM

To create a virtual machine using the downloaded OVA image:

Step 1 Log in to your vSphere account.

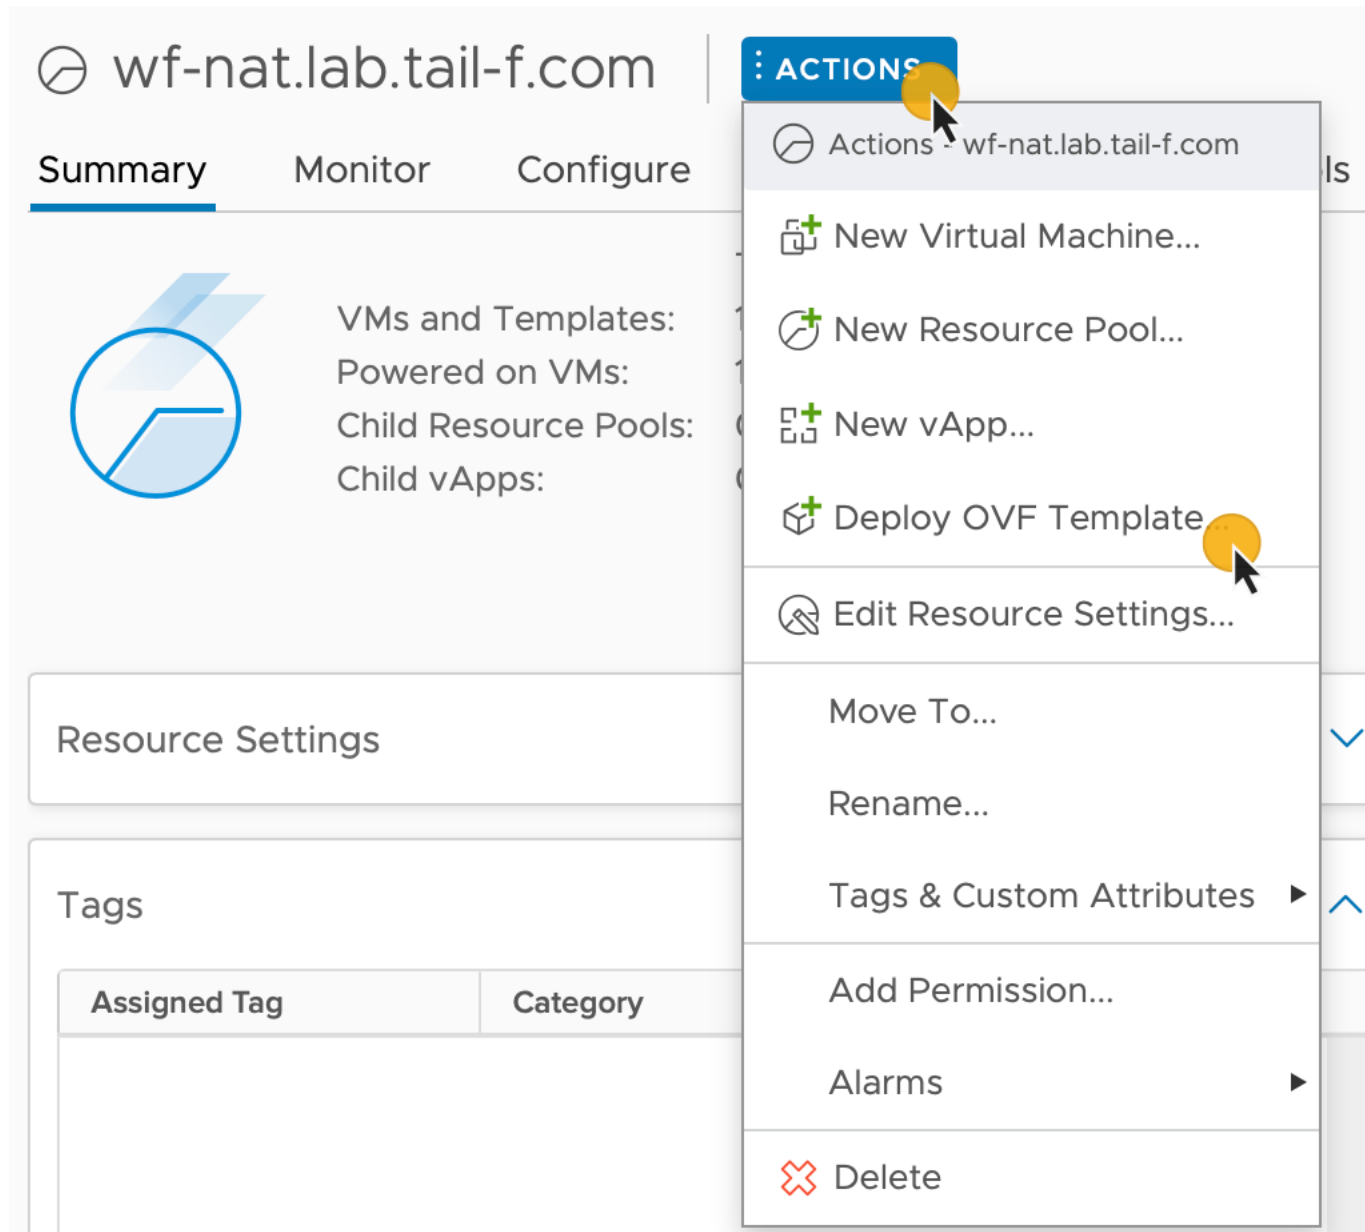
Step 2 In the **Hosts and Clusters** tab, expand your host and select your resource pool.

Figure 1:



Step 3 Click the **Actions** menu and select **Deploy OVF Template**.

Figure 2:



- Step 4** In the **Select an OVF template** step, click **Local file**, **Select files**, and select the CWM OVA image. Click **Next**.
- Step 5** In the **Select a name and folder** step, provide a name for your VM and select its location. Click **Next**.
- Step 6** In the **Select a compute resource** step, select your resource pool. Click **Next**.
- Step 7** In the **Review details** step, click **Next**.
- Step 8** In the **Select storage** step, set **Select virtual disk format** to **Thin provision** and select your storage, then click **Next**.
- Step 9** In the **Select network** step, you need to select destination networks for the **Control Plane** and **Northbound**:
- Control Plane**: select **PrivateNetwork**. If not available, select **VM Network**.

Note Control plane settings are essential only in case of an HA cluster setup. For single-node setups, control plane settings need to be provided, but are not essential and should not conflict with any other devices connected to the control network.

- b) **Northbound**: select **VM Network**.
- c) Click **Next**.

Step 10


In the **Customize template** step, provide the following selected properties:

- a) **Instance Hostname**: type a name for your instance.
- b) **SSH Public Key**: provide an **ed25519** SSH public key that will be used for command-line access to the VM.
- c) **Node Name**: provide a name for installation node.

Note For single-node setups, it's not recommended to modify the node name. If you modify it, remember that it must match the **Zone-A Node Name** below.

- d) **Control Plane Node Count**: change to more than 1 only in case of HA cluster setup. Not supported for CWM 1.1.
- e) **Control Plane IP (ip subnet)**: provide a network address for the control plane. This address cannot conflict with any other devices in the control network, but is otherwise inessential in a single-node setup. Note that the default subnet mask is /24. You can add your custom subnet mask value if applicable for your network settings.
- f) **Initiator IP**: set the initiator IP for the starter node. In a single-node setup, it is the same address as *Control Plane IP**

Figure 3:



- g) **IP (ip subnet) - if not using DHCP:** provide the network address for the node. Note that the default subnet mask is /24. You can add your custom subnet mask value if applicable for your network settings.
- h) **Gateway - if not using DHCP:** provide the gateway address. By default, it is 192.168.1.1.
- i) **DNS:** provide the address for the DNS. By default, it is 8.8.8.8, or you can use your local DNS.
- j) **Northbound Virtual IP:** provide the network address for the active cluster node. In a single-node setup this address is also required, as this is where the HTTP service is working.
- k) **Zone-A Node Name:** provide the name of the Zone-A node. Note that it must match the **Node Name** above.
- l) **Zone-B Node Name:** provide the name of the Zone-B node. For single-node setups, this is not essential and must not be modified.

- m) **Zone-C Node Name (Arbitrator)**: provide the name of the Zone-C Arbitrator node. For single-node setups, this is not essential and must not be modified.
- n) Click **Next**.

Figure 4:

The screenshot displays the 'Deploy OVF Template' wizard. The left pane shows a list of steps from 1 to 8. Step 7, 'Customize template', is highlighted in a dark blue bar. The right pane shows configuration options for the selected step. The 'Northbound Interface' section is expanded, showing options for Protocol, IP (ip[/subnet]) - if n, Gateway - if not usin, and DNS. The 'Initiator Config' section is also expanded, showing options for Initiator Node, Northbound Virtual I, Zone-A Node Name, Zone-B Node Name, and Zone-C Node Name.

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 Select networks
- 7 Customize template**
- 8 Ready to complete

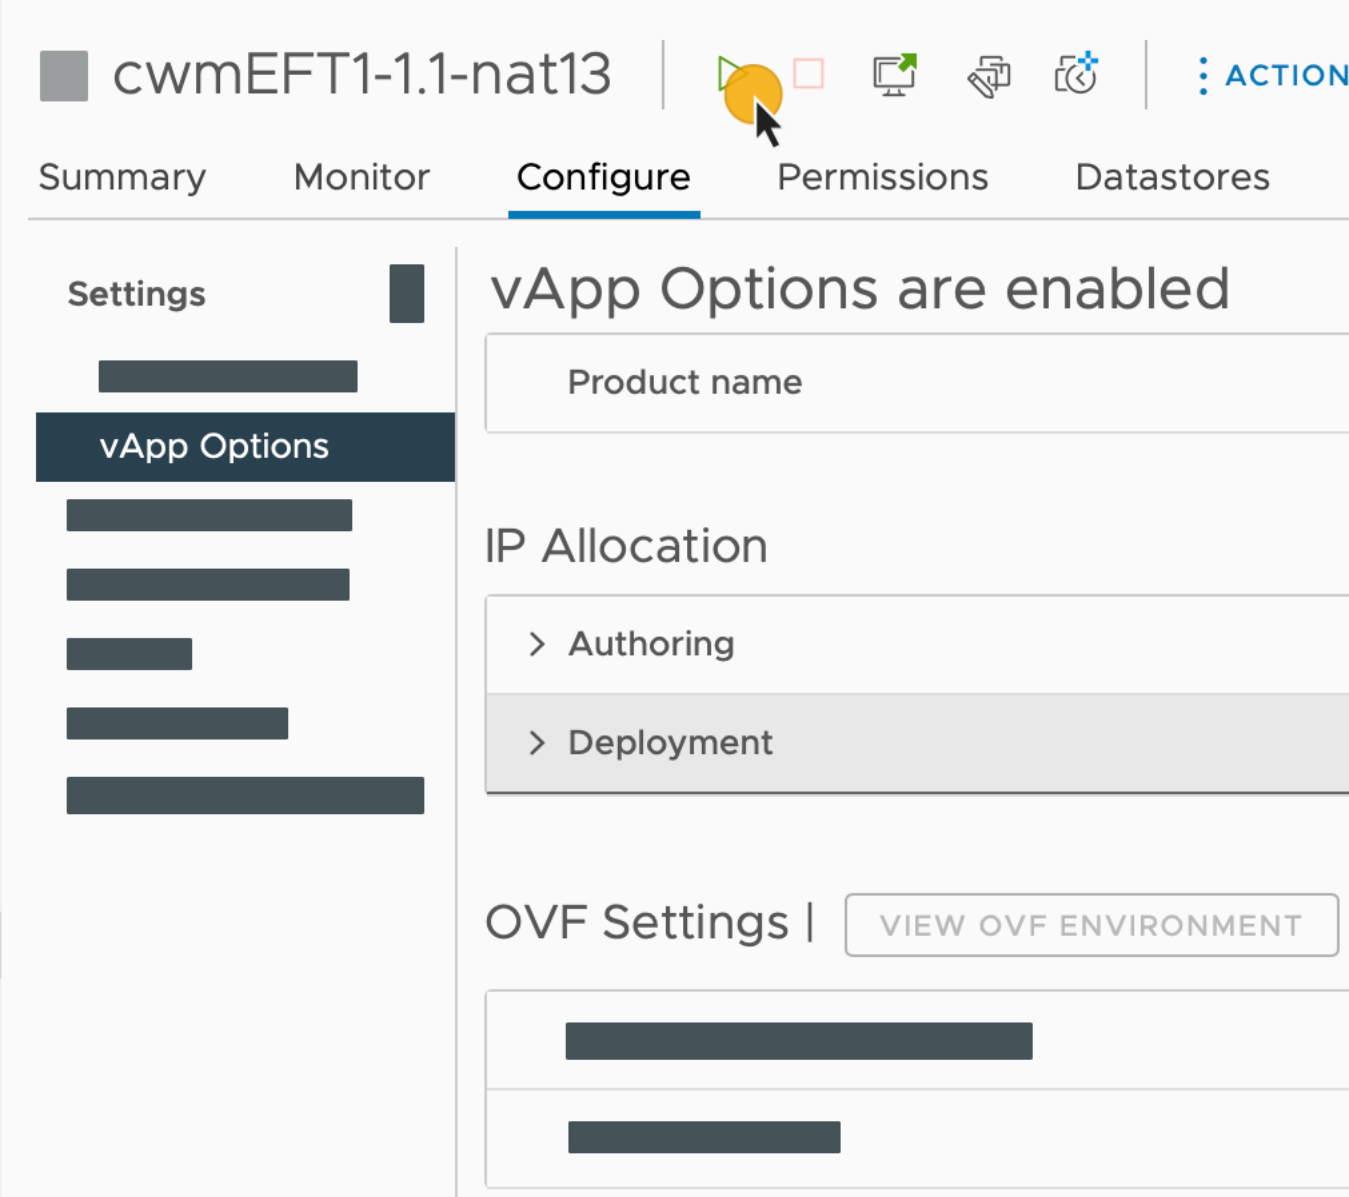
Customize template

- Northbound Interface
 - Protocol
 - IP (ip[/subnet]) - if n
 - Gateway - if not usin
 - DNS
- Initiator Config
 - Initiator Node
 - Northbound Virtual I
 - Zone-A Node Name
 - Zone-B Node Name
 - Zone-C Node Name

Step 11 In the **Ready to complete**, click **Finish**. The deployment may take a few minutes.

Step 12 From the **Resource pool** list, select you newly created virtual machine and click the **Power on** icon.

Figure 5:



The screenshot shows the configuration page for a VM named `cwmEFT1-1.1-nat13`. The `Configure` tab is selected, and the `vApp Options` section is expanded. The `vApp Options` are enabled. The `IP Allocation` section is also expanded, showing `Authoring` and `Deployment` sub-sections. The `OVF Settings` section is visible with a `VIEW OVF ENVIRONMENT` button. A mouse cursor is pointing at the `vApp Options` icon in the top navigation bar.

Note If the VM doesn't power on successfully, this might be due to an intermittent infrastructure error caused by NxF. As a workaround, remove the existing VM and redeploy the OVA on a new one.

Check installation and create user

Before you create a platform user account for first login to the CWM UI, check if the installation is completed successfully and the system is up:

Step 1 Using a command-line terminal, log in to the NxF in your guest OS with SSH:

```
ssh -o UserKnownHostsFile=/dev/null -p 22 nxf@<virtual_IP_address>
```

Note By default, the virtual IP address is the one you set in **IP (ip subnet) - if not using DHCP**. Depending on how vCenter is set up, this can be the resource pool address along with a specific port. Check this with your network administrator in case of doubt

Optional: If you are logging in for the first time, provide the path name for your private key:

```
ssh -i <ed25519_ssh_private_key_name_and_location> nxf@<virtual_IP_address>
```

Note The default port for SSH is 22, change it to your custom port if applicable.

Step 2 Check NxF boot logs:

```
sudo journalctl -u nxf-boot
```

Note Note that it may take a few minutes for the installation to complete. At the bottom of the NxF logs that appear, look for the `NXF: Done setting up machine` message. If the logs report an issue, you might consider reinstalling CWM.

Step 3 Check if all the Kubernetes pods are up and running:

```
kubectl get pods -A
```

This will display a list of pods accompanied by their status, which will resemble the following:

| NAMESPACE | NAME | READY | STATUS | RESTARTS |
|--------------------|--|-------|---------|----------|
| AGE | | | | |
| kube-flannel | kube-flannel-ds-vh4js | 1/1 | Running | 0 |
| 7m35s | | | | |
| kube-system | coredns-9mnzv | 1/1 | Running | 0 |
| 7m35s | | | | |
| kube-system | etcd-node1 | 1/1 | Running | 0 |
| 7m44s | | | | |
| kube-system | kube-apiserver-node1 | 1/1 | Running | 0 |
| 7m50s | | | | |
| kube-system | kube-controller-manager-node1 | 1/1 | Running | 0 |
| 7m50s | | | | |
| kube-system | kube-proxy-6hwg9 | 1/1 | Running | 0 |
| 7m35s | | | | |
| kube-system | kube-scheduler-node1 | 1/1 | Running | 0 |
| 7m42s | | | | |
| local-path-storage | local-path-provisioner-54c455f95-mbhc9 | 1/1 | Running | 0 |
| 7m34s | | | | |
| nxf-system | authenticator-f74c7c87f-m8p4x | 2/2 | Running | 0 |
| 6m25s | | | | |
| nxf-system | controller-76686f8f5f-gpqc | 2/2 | Running | 0 |
| 6m27s | | | | |
| nxf-system | ingress-ports-node1-zchwz | 1/1 | Running | 0 |
| 4m17s | | | | |
| nxf-system | ingress-proxy-bcb8c9fff-lzm9p | 1/1 | Running | 0 |
| 6m23s | | | | |
| nxf-system | kafka-0 | 1/1 | Running | 0 |
| 7m34s | | | | |
| nxf-system | loki-0 | 3/3 | Running | 0 |
| 6m33s | | | | |
| nxf-system | metrics-5qnzb | 2/2 | Running | 0 |
| 6m30s | | | | |
| nxf-system | minio-0 | 2/2 | Running | 0 |

| | | | | | |
|------------|--------------------------------------|-----|---------|--|--------------|
| 7m34s | | | | | |
| nxf-system | postgres-0 | 2/2 | Running | | 0 |
| 6m59s | | | | | |
| nxf-system | promtail-t7dp4 | 1/1 | Running | | 0 |
| 6m33s | | | | | |
| nxf-system | registry-5486f46b54-c6tf9 | 2/2 | Running | | 0 |
| 7m2s | | | | | |
| nxf-system | vip-nodel | 1/1 | Running | | 0 |
| 6m12s | | | | | |
| zone-a | cwm-api-service-67bd9db5c7-vfszs | 2/2 | Running | | 2 (3m37s |
| ago) | | | | | |
| 4m16s | | | | | |
| zone-a | cwm-dsl-service-7ffd6975ff-wlrwt | 2/2 | Running | | 4 (3m21s |
| ago) | | | | | |
| 4m15s | | | | | |
| zone-a | cwm-engine-frontend-6754445fc-67t5h | 2/2 | Running | | 2 (3m52s |
| ago) | | | | | |
| 4m15s | | | | | |
| zone-a | cwm-engine-history-c4dfffd-d-t2fgv | 2/2 | Running | | 1 (2m35s |
| ago) | | | | | |
| 4m14s | | | | | |
| zone-a | cwm-engine-history-c4dfffd-d-wr5v2 | 2/2 | Running | | 2 (3m51s |
| ago) | | | | | |
| 4m14s | | | | | |
| zone-a | cwm-engine-history-c4dfffd-d-zz74q | 2/2 | Running | | 4 (48s ago) |
| 4m14s | | | | | |
| zone-a | cwm-engine-matching-78dfdf858f-q8wg2 | 2/2 | Running | | 2 (3m46s |
| ago) | | | | | |
| 4m14s | | | | | |
| zone-a | cwm-engine-ui-6b74755499-jwbld | 2/2 | Running | | 0 |
| 4m13s | | | | | |
| zone-a | cwm-engine-worker-589b6bc88b-hs2ch | 2/2 | Running | | 0 |
| 4m13s | | | | | |
| zone-a | cwm-event-manager-5b95bb49db-gw6g5 | 2/2 | Running | | 0 |
| 4m12s | | | | | |
| zone-a | cwm-plugin-manager-76f798446c-qgx27 | 2/2 | Running | | 1 (2m29s |
| ago) | | | | | |
| 4m12s | | | | | |
| zone-a | cwm-ui-779bdb44-98d5v | 2/2 | Running | | 0 |
| 4m11s | | | | | |
| zone-a | cwm-worker-manager-7bd8795b56-f4czp | 2/2 | Running | | 1 (112s ago) |
| 4m10s | | | | | |
| zone-a | logcli-5f8cc8c585-fq7wm | 2/2 | Running | | 0 |
| 4m10s | | | | | |

Note Note that it may take a few minutes for the system to get all the pods running. If any of the pods stays in a status other than Running, consider using the `kubectl delete pod <pod_name> -n <namespace>` command to restart it.

Create user for UI login

You can create CWM platform user accounts using the command-line access to the VM. Here's how to do it:

Step 1 Using a command-line terminal, log in to the NxF in your guest OS with SSH:

```
ssh -o UserKnownHostsFile=/dev/null -p 22 nxf@<virtual_IP_address>
```

Optional: If you are logging in for the first time, provide the path name for your private key:

```
ssh -i <ed25519_ssh_private_key_name_and_location> nxf@<virtual_IP_address>
```

Note The default port for SSH is 22, change it to your custom port if applicable.

Step 2 To create a user with a password, run the following commands:

a) First, set minimum password complexity (default is 3, 0 is complexity disabled):

```
sedo security password-policy set --min-complexity-score 1
```

b) Then create user account and a password:

```
echo -en 'Password123!' | sedo security user add --password-stdin \  
--access permission/admin --access permission/super-admin \  
--access permission/user --display-name Tester test
```

c) Optionally, disable the password change requirement for the test user:

```
sedo security user set test --must-change-password=false
```

Step 3 To see the CWM UI, go to the address that you selected for Northbound IP and default port 8443. For example, <https://192.168.1.233:8443/>.

Step 4 Log in using the `test` username and password.

