



# Install Crosswork Cluster on KVM

---

This chapter contains the following topics:

- [Installation overview, on page 1](#)
- [Installation parameters, on page 1](#)
- [Install Crosswork Network Controller cluster using CLI, on page 7](#)
- [Monitor Cluster Activation, on page 13](#)
- [Log into the Cisco Crosswork UI, on page 15](#)
- [Troubleshoot the cluster, on page 16](#)

## Installation overview

The Crosswork Network Controller cluster is installed on KVM primarily via CLI. This is the recommended installation approach.

Python scripts are used on the bare metal where the VMs will be running prior to installing the Crosswork Network Controller cluster.



---

**Attention** SR-PCE is not currently supported on KVM. Customers using KVM to host Crosswork Network Controller will need to either enable SR-PCE functions on a physical device or deploy the OVA using VMware.

---

## Installation parameters

This section explains the important parameters that must be specified while installing the Crosswork cluster.

Kindly ensure that you have relevant information to provide for each of the parameters mentioned in the table and that your environment meets all the requirements specified under [Installation Prerequisites for KVM](#).

The settings recommended in the table represent the least complex configuration. If you encounter network conflicts or wish to implement more advanced security settings (e.g., self-signed certificates), please work with the Cisco Customer Experience team to ensure you are prepared to make the necessary changes for your cluster.

**Attention**

- Please use the latest template file that comes with the installation.
- Secure ZTP and Secure Syslog require the Crosswork cluster to be deployed with FQDN.

**Table 1: General parameters**

Parameter name	Description
ClusterName	Name of the cluster file.
ClusterIPStack	The IP stack protocol: IPv4, IPv6, or DUALSTACK.
AdminIPv4Address AdminIPv6Address	The Admin IP address of the VM (IPv4 and/or IPv6).
AdminIPv4Netmask AdminIPv6Netmask	The Admin IP subnet in dotted decimal format (IPv4 and/or IPv6).
ClusterCaKey	The CA private key. Use the default value (Empty).
ClusterCaPubKey	The CA public key. Use the default value (Empty).
CwInstaller	Set as "False".
Deployment	Enter the deployment type.
Disclaimer	Enter the disclaimer message.
ManagementIPv4Address ManagementIPv6Address	The Management IP address of the VM (IPv4 and/or IPv6).
ManagementIPv4Netmask ManagementIPv6Netmask	The Management IP subnet in dotted decimal format (IPv4 and/or IPv6).
ManagementIPv4Gateway ManagementIPv6Gateway	The Gateway IP on the Management Network (IPv4 and/or IPv6). The address must be reachable, otherwise the installation will fail.
ManagementVIP	The Management Virtual IP address for the cluster.
ManagementVIPName	Name of the Management Virtual IP for the cluster. This is an optional parameter used to reach Crosswork cluster Management VIP via DNS name. If this parameter is used, the corresponding DNS record must exist in the DNS server.
ManagementPeerIPs	The Management peer IP addresses (IPv4 and/or IPv6) for the cluster.
DataIPv4Address DataIPv6Address	The Data IP address of the VM (IPv4 and/or IPv6).
DataIPv4Netmask DataIPv6Netmask	The Data IP subnet in dotted decimal format (IPv4 and/or IPv6).
DataIPv4Gateway DataIPv6Gateway	The Gateway IP on the Data Network (IPv4 and/or IPv6). The address must be reachable, otherwise the installation will fail.

Parameter name	Description
DataVIP	The Data Virtual IP address for the cluster.
DataVIPName	Name of the Data Virtual IP for the cluster. This is an optional parameter used to reach Crosswork cluster Data VIP via DNS name. If this parameter is used, the corresponding DNS record must exist in the DNS server.
DataPeerIPs	The Data peer IP addresses (IPv4 and/or IPv6) for the cluster.
NBIIIPv4Address NBIIIPv6Address	The NBI IP address of the VM (IPv4 and/or IPv6).
NBIIIPv4Netmask NBIIIPv6Netmask	The NBI IP subnet in dotted decimal format (IPv4 and/or IPv6).
NBIIIPv4Gateway NBIIIPv6Gateway	The Gateway IP on the NBI Network (IPv4 and/or IPv6). The address must be reachable, otherwise the installation will fail.
NBIVIP	The NBI Virtual IP address for the cluster.
DNSv4 DNSv6	The IP address of the DNS server (IPv4 and/or IPv6). The address must be reachable, otherwise the installation will fail.
NTP	NTP server address or name. The address must be reachable, otherwise the installation will fail.
DomainName	The domain name used for the cluster.
CWPassword	<p>Password to log into Cisco Crosswork. When setting up a VM, ensure the password is strong and meets the following criteria:</p> <ul style="list-style-type: none"> <li>• It must be at least 8 characters long and include uppercase and lowercase letters, numbers, and at least one special character.</li> <li>• The following special characters are not allowed: backslash (\), single quote ('), or double quote (").</li> <li>• Avoid using passwords that resemble dictionary words (e.g., "Pa55w0rd!") or relatable words. While such passwords may meet the specified criteria, they are considered weak and will be rejected, resulting in a failure to set up the VM.</li> </ul>
VMSize	<p>Sets the VM size for the cluster. For cluster deployments, the only supported option is "Large".</p> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• If you leave this field blank, the default value ("Large") is automatically selected.</li> <li>• This parameter accepts a string value, so be sure to enclose the value in double quotes.</li> </ul>

Parameter name	Description
VMName	Name of the VM. A unique VM name is required for each node on the cluster (Hybrid or Worker).
VMLocation	Location of the VM.
VMType	Indicates the type of VM. Choose either "Hybrid" or "Worker". This parameter accepts a string value, so be sure to enclose the value in double quotes.  <b>Note</b> The Crosswork cluster requires at least three VMs operating in a hybrid configuration.
IsSeed	Choose "True" if this is the first VM being built in a new cluster. Choose "False" for all other VMs, or when rebuilding a failed VM.  This parameter accepts a string value, so be sure to enclose the value in double quotes.
InitNodeCount	Total number of nodes in the cluster including Hybrid and Worker nodes. The default value is 3. Set this to match the number of VMs (nodes) you are going to deploy.
InitMasterCount	Total number of Hybrid nodes in the cluster. The default value is 3.
BackupMinPercent	Minimum percentage of the data disk space to be used for the size of the backup partition. The default value is 35 (valid range is from 1 to 80).  Please use the default value unless recommended otherwise.  <b>Note</b> The final backup partition size will be calculated dynamically. This parameter defines the minimum.
ddataafs	Refers to the data disk size for the nodes (in Giga Bytes). This is an optional parameter and the default value is 485 (valid range is from 485 to 8000), if not explicitly specified.  Please use the default value unless recommended otherwise.
ssd	Refers to the ssd disk size. This is an optional parameter and the default value is 15.  Please use the default value unless recommended otherwise.
ThinProvisioned	Set to <b>false</b> for production deployments.

Parameter name	Description
EnableHardReservations	<p>Determines the enforcement of VM CPU and Memory profile reservations (see <a href="#">Installation Prerequisites for KVM</a> for more information). This is an optional parameter and the default value is "True", if not explicitly specified.</p> <p><b>Note</b> This parameter accepts a string value, so be sure to enclose the value in double quotes.</p> <p>If set as "True", the VM's resources are provided exclusively. In this state, the installation will fail if there are insufficient CPU cores, memory or CPU cycles.</p> <p>If set as "False" (only set for lab installations), the VM's resources are provided on best efforts. In this state, insufficient CPU cores can impact performance or cause installation failure.</p>
ramdisk	<p>Size of the Ram disk.</p> <p>This parameter is only used for lab installations (value must be at least 2). When a non-zero value is provided for <code>RamDiskSize</code>, the <code>HSDatastore</code> value is not used.</p>
OP_Status	<p>This optional parameter is used (uncommented) to import inventory post manual deployment of Crosswork cluster.</p> <p>The parameter refers to the state for this VM. To indicate a running status, the value must be 2 (<code>#OP_Status = 2</code>).</p>
SchemaVersion	<p>The configuration Manifest schema version. This indicates the version of the installer to use with this template.</p> <p>Schema version should map to the version packaged with the sample template on <a href="#">cisco.com</a>. You should always build a new template from the default template provided with the release you are deploying, as template requirements may change from one release to the next.</p>
logfs	<p>Log partition size (in Giga Bytes). Default value is 20 GB and Maximum value is 1000 GB. You are recommended to use the default value.</p>
corefs	<p>Core partition size (in Giga Bytes). Default value is 18 GB and Maximum value is 1000 GB. You are recommended to use the default value.</p>

Parameter name	Description
Timezone	<p>Enter the timezone. Input is a standard IANA time zone (for example, "America/Chicago"). If left blank, the default value (UTC) is selected. This parameter accepts a string value, so be sure to enclose the value in double quotes.</p> <p>This is an optional parameter.</p> <p><b>Note</b> The timestamp in Kafka log messages represents the NSO server time. To avoid any mismatch between the Crosswork server time and the NSO event time, ensure you update the NSO server time before changing the <code>Timezone</code> parameter in Crosswork.</p>
EnableSkipAutoInstallFeature	<p>Pods marked as "skip auto install" will not be brought up unless explicitly requested by a dependent application or pod. By default, the value is set as "False".</p> <p>The recommended value for cluster deployment is "False".</p> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• If left blank, the default value is automatically selected.</li> <li>• This parameter accepts a string value, so be sure to enclose the value in double quotes.</li> </ul>
EnforcePodReservations	<p>Enforces minimum resource reservations for the pod. If left blank, the default value ("True") is selected.</p> <p>This parameter accepts a string value, so be sure to enclose the value in double quotes.</p>
K8Orch	<p>Enforces minimum resource reservations for the pod. If left blank, the default value ("True") is selected.</p> <p>This parameter accepts a string value, so be sure to enclose the value in double quotes.</p>
K8sServiceNetwork	<p>The network address for the kubernetes service network. By default, the CIDR range is fixed to '10.96.0.0/16'. If you wish to change this default value, please work with the Cisco Customer Experience team.</p>
K8sPodNetwork	<p>The network address for the kubernetes pod network. By default, the CIDR range is fixed to '10.224.0.0/16'. If you wish to change this default value, please work with the Cisco Customer Experience team.</p>
bootOptions.efiSecureBootEnabled	<p>Default value is "True".</p> <p>This parameter accepts a string value, so be sure to enclose the value in double quotes.</p>

Parameter name	Description
IgnoreDiagnosticsCheckFailure	<p>Used to set the system response in case of a diagnostic check failure.</p> <p>If set to "False" (default value), the installation will terminate if the diagnostic check reports an error. If set to "True", the diagnostic check will be ignored, and the installation will continue.</p> <p>You are recommended to select the default value. This parameter accepts a string value, so be sure to enclose the value in double quotes.</p> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>The log files (<code>diagnostic_stdout.log</code> and <code>diagnostic_stderr.log</code>) can be found at <code>/var/log</code>. The result from each diagnostic execution is kept in a file at <code>/home/cw-admin/diagnosis_report.txt</code>.</li> <li>Use <b>diagnostic all</b> command to invoke the diagnostic manually on day N.</li> <li>Use <b>diagnostic history</b> command to view previous test report.</li> </ul>

## Install Crosswork Network Controller cluster using CLI

This section provides the high-level workflow for installing Crosswork Network Controller cluster on KVM via CLI.

**Table 2: Installation workflow**

Step	Action
1. Ensure you have performed the preliminary checks.	See <a href="#">Preliminary checks</a> , on page 8 for details.
2. Set up and validate the KVM environment.	See <a href="#">Set up and validate KVM on RHEL</a> , on page 8.
3. Configure network bridges and SRIOV	See <a href="#">Configure network bridges or SRIOV</a> , on page 9.
4. Install Crosswork Network Controller cluster on KVM.	See <a href="#">Install Crosswork Network Controller cluster on KVM using CLI</a> , on page 11.

### Known limitations

- If you are using a non-root user ID for the deployment of nodes on the bare metals, ensure that the particular user ID has been added to the sudoers list (i.e., `/etc/sudoers`).

## Preliminary checks

- 1. Virtualization:** Ensure that your system supports virtualization. This is typically enabled in the BIOS. To check, use these commands:
  - For Intel CPUs: `grep -wo 'vmx' /proc/cpuinfo`
  - For AMD CPUs: `grep -wo 'svm' /proc/cpuinfo`
- 2. KVM modules:** Ensure that the KVM modules are loaded: `lsmod | grep kvm`

## Set up and validate KVM on RHEL

Follow these steps to set up KVM on RHEL.

### Procedure

---

- Step 1** Refresh repositories and install updates. This command updates all the packages on your system to their latest versions.
- ```
sudo dnf update -y
```
- Step 2** Reboot the system after all the updates are installed successfully.
- ```
sudo reboot
```
- Step 3** Install virtualization tools.
- Install `virt-install` and `virt-viewer`.

```
sudo dnf install virt-install virt-viewer -y
```

`virt-install` is a command-line tool for creating virtual machines.  
`virt-viewer` is a Lightweight UI for interacting with VMs.
  - Install `libvirt` virtualization daemon, which is necessary for managing VMs.

```
sudo dnf install -y libvirt
```
  - Install `virt-manager`, a graphical interface for managing VMs.

```
sudo dnf install virt-manager -y
```
  - Install additional virtualization tools for managing VMs.

```
sudo dnf install -y virt-top libguestfs-tools
```
- Step 4** Start and enable `libvirtd` virtualization daemon.
- Start the `libvirtd` daemon.

```
sudo systemctl start libvirtd
```
  - Enable the `libvirtd` daemon.

```
sudo systemctl enable libvirtd
```
  - Verify that the Daemon is running.

```
sudo systemctl status libvirtd
```

**Step 5** Add users to the required groups, for example, libvirt and qemu. In the following commands, replace *your\_username* with the actual username.

```
sudo usermod --append --groups libvirt your_username
sudo usermod --append --groups qemu your_username
```

**Step 6** Ensure that IOMMU is enabled. If it is not enabled, run this command to enable it.

```
grubby --update-kernel=ALL --args=intel_iommu=on
dmesg | grep -I IOMMU
```

**Step 7** Check IOMMU and validate the setup. Ensure that all checks show as PASS.

```
virt-host-validate
```

If the IOMMU check is not PASS, then use the following commands to enable it.

```
sudo grubby --update-kernel=ALL --args=intel_iommu=on
sudo reboot
```

## Configure network bridges or SRIOV

Crosswork needs the 10G interface for all the data layer communications to support functionality at a scale. You may choose any networking configuration which can provide 10G throughput.

The following sections explain how to enable bridging and SRIOV network configuration.




---

**Important** For KVM deployment, configure either network bridges or SRIOV, but not both.

---

### Configure network bridges

A network bridge acts like a virtual network switch, allowing multiple network interfaces to communicate as if they are on the same physical network.

Follow these steps to configure network bridges.

#### Procedure

---

**Step 1** Create a new network connection of type "bridge" with the interface name `intMgmt` and assign it the connection name `intMgmt`.

```
nmcli connection add type bridge ifname intMgmt con-name intMgmt
```

**Step 2** Add a new `bridge-slave` connection, associating the physical network interface `<interface1>` with the previously created bridge `intMgmt`.

```
nmcli connection add type bridge-slave ifname <interface1> controller intMgmt
```

Example:

```
nmcli con add type bridge-slave ifname <hostdataIntf> master intData con-name
intData-slave-<hostdataIntf>
```

**Step 3** Assign IP address to the bridge.

```
nmcli connection modify intMgmt ipv4.addresses <IPv4-address>/<subnet-mask>
```

Example:

```
nmcli con modify intMgmt ipv4.addresses <hostmgmtIp/mask> ipv4.gateway
<mgmtgw> ipv4.dns <dnsIp> ipv4.method manual ipv4.route-metric 50
```

**Step 4** Bring up the `intMgmt` network connection.

```
nmcli connection up intMgmt
```

Example:

```
nmcli con up intMgmt
nmcli con up intMgmt-slave-<hostmgmtIntf>
```

**Step 5** Create another network bridge connection with the interface name `intData` and assign it the connection name `intData`.

```
nmcli connection add type bridge ifname intData con-name intData
```

Example:

```
nmcli con add type bridge ifname intData con-name intData
```

**Step 6** Add a bridge-slave connection, associating the physical network interface `<interface2>` with the previously created bridge `intData`.

```
nmcli connection add type bridge-slave ifname <interface2> controller intData
```

Example:

```
nmcli con add type bridge-slave ifname <hostmgmtIntf> master intMgmt con-name
intMgmt-slave-<hostmgmtIntf>
```

**Step 7** Assign IP address to `intData`.

```
nmcli connection modify intData ipv4.addresses <IPv4-address>/<subnet-mask>
```

Example:

```
nmcli con modify intData ipv4.addresses <hostdataIp/mask> ipv4.method manual ipv4.gateway <datagw>
ipv4.route-metric 90
```

**Step 8** Bring up the `intData` network connection.

```
nmcli connection up intData
```

Example:

```
nmcli con up intData
nmcli con up intData-slave-<hostdataIntf>
```

## Configure SRIOV

SRIOV allows a single physical network interface to be shared among multiple VMs by creating multiple Virtual Functions (VFs).

Follow these steps to configure SRIOV.

## Procedure

**Step 1** Open the `rc.local` file in the vi editor.

```
vi /etc/rc.d/rc.local
```

**Step 2** Set the number of VFs for the network interfaces based on your requirement. Two VFs are configured for each interface by default. You may also configure additional VFs for future scalability needs.

For example, to set the number of VFs to 2 for each `<interface1>` and `<interface2>`, use these commands. In this example, `<interface1>` refers to the management interface and `<interface2>` refers to the data interface.

```
echo 2 > /sys/class/net/<interface1>/device/sriov_numvfs
echo 2 > /sys/class/net/<interface2>/device/sriov_numvfs
```

**Step 3** Change the permissions of the `rc.local` file to make it executable.

```
chmod +x /etc/rc.d/rc.local
```

**Step 4** If any of the interfaces are configured over the VLAN, set the VLAN IDs to the interfaces.

```
ip link set <interface1> vf 0 vlan <vlanid>
ip link set <interface2> vf 1 vlan <vlanid>
```

**Step 5** Save the changes and reboot the system.

**Step 6** List all the PCI devices for all the virtual functions in a tree format. This is useful for verifying the setup and ensuring that the VFs are correctly recognized by the KVM hypervisor.

```
virsh nodedev-list --tree
```

```
|+- pci_0000_17_00_0
|  |  |
|  |+- net_ens1f0_40_a6_b7_ce_04_c8
|  |  |
|  |+- pci_0000_17_00_1
|  |  |
|  |+- net_ens1f1_40_a6_b7_ce_04_c9
|  |  |
|  |+- pci_0000_17_00_2
|  |  |
|  |+- net_ens1f2_40_a6_b7_ce_04_ca
|  |  |
|  |+- pci_0000_17_00_3
|  |  |
|  |+- net_ens1f3_40_a6_b7_ce_04_cb
```

In this procedure, since we set the number of VFs as 2 in Step 2, two VFs for each management interface and data interface are created. As a result, a total of four PCI devices are generated: two for management and two for data.

This PCI device information is used during the installation process with SRIOV (Step 4 of [Install Crosswork Network Controller cluster on KVM using CLI, on page 11](#)).

## Install Crosswork Network Controller cluster on KVM using CLI

Follow these steps to install the Crosswork Network Controller VM on KVM using CLI. These steps must be repeated for each VM in your cluster.




---

**Note** The time taken to create the cluster can vary based on the size of your deployment profile and the performance characteristics of your hardware.

---

### Before you begin

Ensure that

- the KVM host meets all the requirements specified in [KVM host bare metal requirements](#).
- the checks specified in [Preliminary checks, on page 8](#) is completed.
- your KVM environment is set up (see [Set up and validate KVM on RHEL, on page 8](#)), and
- the network bridge or SRIOV is configured (see [Configure network bridges or SRIOV, on page 9](#)).

### Procedure

---

**Step 1** As a first step, prepare the config IOS files for the Crosswork Network Controller cluster. You must create a separate config IOS file (`ovf-env.xml`) for each VM in your cluster. For more information, see [Crosswork Network Controller deployment templates for KVM](#).

#### Warning

Changing the file name from `ovf-env.xml` will cause errors. Use the exact file name.

- a) Update the `ovf-env.xml` file as per your needs. For more information on the parameters, see [Installation parameters, on page 1](#).

```
$ cat ovf-env.xml
```

- b) Generate the IOS file.

```
$ mkisofs -R -relaxed-filenames -joliet-long -iso-level 3 -l -o node-1-Hybrid.iso ovf-env.xml
```

#### Note

It is recommended to use the VM hostname as the name of the `.iso` file to avoid confusion. In the above example command, "node-1-Hybrid" is the VM hostname.

**Step 2** Download the Crosswork Network Controller cluster qcow2 tar file and extract it.

```
tar -xvf cnc-7.1.0-85-release710-250530-qcow2.tar
```

This command creates three qcow2 files:

- `cnc-7.1.0-85-release710-250530_dockerfs.qcow2`
- `cnc-7.1.0-85-release710-250530_extrafs.qcow2`
- `cnc-7.1.0-85-release710-250530_rootfs.qcow2`

**Step 3** Navigate to the required installation folder and create three disks.

```
cd node-1-Hybrid/
qemu-img create -f qcow2 disk3 20G
```

```
qemu-img create -f qcow2 disk4 485G
qemu-img create -f qcow2 disk6 15G

ls -l
cw_dockerfs.vmdk.qcow2
cw_extrafs.vmdk.qcow2
cw_rootfs.vmdk.qcow2
disk3
disk4
disk6
```

#### Step 4 Install the Crosswork Network Controller cluster using network bridge or SRIOV.

In this example, "node-1-Hybrid" is the host name of the Cisco Crosswork VM.

- Using network bridges:

```
virt-install --boot uefi --boot hd,cdrom --connect qemu:///system --virt-type kvm --name
node-1-Hybrid --ram 98304 --vcpus 12 --os-type linux --disk
path=cw-na-cnc-essential-7.1.0-85-release710-250530_rootfs.qcow2,format=qcow2,bus=scsi --disk
path=cw-na-cnc-essential-7.1.0-85-release710-250530_dockerfs.qcow2,format=qcow2,bus=scsi --disk
path=disk3,format=qcow2,bus=scsi --disk path=disk4,format=qcow2,bus=scsi --disk
path=cw-na-cnc-essential-7.1.0-85-release710-250530_extrafs.qcow2,format=qcow2,bus=scsi --disk
path=disk6,format=qcow2,bus=scsi --disk=node-1-Hybrid.iso,device=cdrom,bus=scsi --import --network
bridge=intMgmt,model=virtio --network bridge=intData,model=virtio --noautoconsole --os-variant
ubuntu22.04 --graphics vnc,listen=0.0.0.0
```

- Using SRIOV:

```
virt-install --boot uefi --boot hd,cdrom --connect qemu:///system --virt-type kvm --name
node-1-Hybrid --ram 98304 --vcpus 12 --cpu host-passthrough --disk
path=cw_rootfs.vmdk.qcow2,format=qcow2,bus=scsi --disk
path=cw_dockerfs.vmdk.qcow2,format=qcow2,bus=scsi --disk path=disk3,format=qcow2,bus=scsi --disk
path=disk4,format=qcow2,bus=scsi --disk path=cw_extrafs.vmdk.qcow2,format=qcow2,bus=scsi --disk
path=disk6,format=qcow2,bus=scsi --disk=node-1-Hybrid.iso,device=cdrom,bus=scsi --import --network
none --host-device=pci_0000_01_10_0 --host-device=pci_0000_01_10_0 --os-variant ubuntu-lts-latest
&
```

## Monitor Cluster Activation

This section explains how to monitor and verify if the installation has completed successfully. As the installer builds and configures the cluster it will report progress. The installer will prompt you to accept the license agreement and then ask if you want to continue the install. After you confirm, the installation will progress and any possible errors will be logged in either `installer.log` or `installer_tf.log`. If the VMs get built and are able to boot, the errors in applying the operator specified configuration will be logged on the VM in the `/var/log/firstboot.log`.



**Note** During installation, Cisco Crosswork will create a special administrative ID (**virtual machine (VM) administrator**, `cw-admin`, with the password that you provided in the manifest template. In case the installer is unable to apply the password, it creates the administrative ID with the default password `cw-admin`). The first time you log in using this administrative ID, you will be prompted to change the password.

The administrative username is reserved and cannot be changed. Data center administrators use this ID to log into and troubleshoot the Crosswork application VM.

The following is a list of critical steps in the process that you can watch for to be certain that things are progressing as expected:

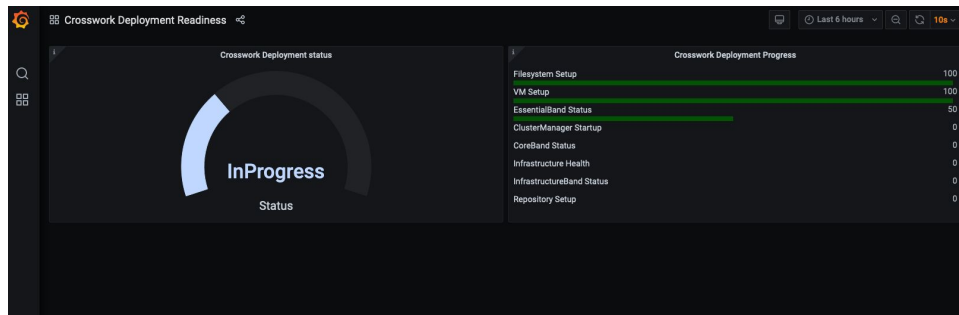
1. The installer uploads the crosswork image file (.ova file) to the data center.
2. The installer creates the VMs, and displays a success message (e.g. "Creation Complete") after each VM is created.
3. After each VM is created, it is powered on (either automatically when the installer completes, or after you power on the VMs during the manual installation). The parameters specified in the template are applied to the VM, and it is rebooted. The VMs are then registered by Kubernetes to form the cluster.
4. Once the cluster is created and becomes accessible, a success message (e.g. "Crosswork Installer operation complete") will be displayed and the installer script will exit and return you to a prompt on the screen.

You can monitor startup progress using the following methods:

- **Using browser accessible dashboard:**

1. While the cluster is being created, monitor the setup process from a browser accessible dashboard.
2. The URL for this grafana dashboard (in the format `http://{VIP}:30602/d/NK1bwXGk/crosswork-deployment-readiness?orgId=1&refresh=10s&theme=dark`) is displayed once the installer completes. This URL is temporary and will be available only for a limited time (around 30 minutes).
3. At the end of the deployment, the grafana dashboard will report a "Ready" status. If the URL is inaccessible, use the SSH console described in this section to monitor the installation process.

**Figure 1: Crosswork Deployment Readiness**



- **Using the console:**

1. Check the progress from the console of one of the hybrid VMs or by using SSH to the Virtual IP address.
2. In the latter case, login using the `cw-admin` user name and the password you assigned to that account in the install template.
3. Switch to super user using `sudo su -` command.
4. Run `kubectl get nodes` (to see if the nodes are ready) and `kubectl get pods` (to see the list of active running pods) commands.
5. Repeat the `kubectl get pods` command until you see `robot-ui` in the list of active pods.
6. At this point, you can try to access the Cisco Crosswork UI.

### Failure Scenario

In the event of a failure scenario (listed below), contact the Cisco Customer Experience team and provide the `installer.log`, `installer_tf.log`, and `firstBoot.log` files (there will be one per VM) for review:

- Installation is incomplete
- Installation is completed, but the VMs are not functional
- Installation is completed, but you are directed to check `/var/log/firstBoot.log` or `/opt/robot/bin/firstBoot.log` file.

## Log into the Cisco Crosswork UI

Once the cluster activation and startup have been completed, you can check if all the nodes are up and running in the cluster from the Cisco Crosswork UI.



---

**Note** For the supported browser versions, see the *Compatibility Information* section in the [Release Notes for Crosswork Network Controller 7.1.0](#).

---

Perform the following steps to log into the Cisco Crosswork UI and check the cluster health:

### Procedure

---

**Step 1** Launch one of the supported browsers.

**Step 2** In the browser's address bar, enter:

```
https://<Crosswork Management Network Virtual IP (IPv4)>:30603/
```

or

```
https://[<Crosswork Management Network Virtual IP (IPv6)>]:30603/
```

**Note**

Please note that the IPv6 address in the URL must be enclosed with brackets.

**Note**

You can also log into the Crosswork UI using the Crosswork FQDN name.

The **Log In** window opens.

**Note**

When you access the Cisco Crosswork for the first time, some browsers display a warning that the site is untrusted. When this happens, follow the prompts to add a security exception and download the self-signed certificate from the Cisco Crosswork server. After you add a security exception, the browser accepts the server as a trusted site in all future login attempts. If you want to use a CA signed certificate, see the *Manage Certificates* topic in the [Cisco Crosswork Network Controller 7.1 Administration Guide](#).

**Step 3** Log into the Cisco Crosswork as follows:

- a) Enter the Cisco Crosswork administrator username **admin** and the default password **admin**.
- b) Click **Log In**.
- c) When prompted to change the administrator's default password, enter the new password in the fields provided and then click **OK**.

**Note**

Use a strong VM Password (minimum 8 characters long, including upper & lower case letters, numbers, and one special character). Avoid using passwords similar to dictionary words (for example, "Pa55w0rd!") or relatable words.

The **Crosswork Manager** window is displayed.

- Step 4** Click on the **Crosswork Health** tab, and click the **Crosswork Platform Infrastructure** tab to view the health status of the microservices running on Cisco Crosswork.
- Step 5** (Optional) Change the name assigned to the admin account (by default, it is "John Smith") to something more relevant.
- Step 6** After logging into the Crosswork UI, ensure the cluster is healthy. Download the cluster inventory sample (.tfvars file) from the Crosswork UI and update it with information about the VMs in your cluster, along with the data center parameters. For sample .tfvars templates, see [Crosswork Network Controller deployment templates for KVM](#). Then, import the file back into the Crosswork UI.

## Troubleshoot the cluster

By default, the installer displays progress data on the command line. The install log is fundamental in identifying the problems, and it is written into the `/data` directory.

### General scenarios

*Table 3: General scenarios*

Scenario	Possible Resolution
Certificate Error	The RHEL hosts that will run the Crosswork application and the data gateway VM must have NTP configured, or the initial handshake may fail with "certificate not valid" errors.
Floating VIP address is not reachable	The VRRP protocol requires unique router_id advertisements to be present on the network segment. By default, Crosswork uses the ID 169 on the management and ID 170 on the data network segments. A symptom of conflict, if it arises, is that the VIP address is not reachable. Remove the conflicting VRRP router machines or use a different network.

Scenario	Possible Resolution																
Crosswork VM is not allowing the admin user to log in	<p>This happens when the password is not complex enough. Create a strong password, update the configuration manifest and redeploy.</p> <p>Use a strong VM Password (8 characters long, including upper &amp; lower case letters, numbers, and at least one special character). Avoid using passwords similar to dictionary words (for example, "Pa55w0rd!") or relatable words. While they satisfy the criteria, such passwords are weak and will be rejected resulting in failure to setup the VM.</p>																
Deployment fails with: <i>Failed to validate Crosswork cluster initialization.</i>	<p>The clusters' seed VM is either unreachable or one or more of the cluster VMs have failed to get properly configured.</p> <ol style="list-style-type: none"> <li>1. Check whether the VM is reachable, and collect logs from <code>/var/log/firstBoot.log</code> and <code>/var/log/vm_setup.log</code></li> <li>2. Check the status of the other cluster nodes.</li> </ol>																
The VMs are deployed but the Crosswork cluster is not being formed.	<p>A successful deployment allows the operator logging in to the VIP or any cluster IP address to run the following command to get the status of the cluster:</p> <pre>sudo kubectl get nodes</pre> <p>A healthy output for a 3-node cluster is:</p> <table border="1"> <thead> <tr> <th>NAME</th> <th>STATUS</th> <th>ROLES</th> <th>AGE</th> </tr> </thead> <tbody> <tr> <td>172-25-87-2-hybrid.cisco.com</td> <td>Ready</td> <td>master</td> <td>41d</td> </tr> <tr> <td>172-25-87-3-hybrid.cisco.com</td> <td>Ready</td> <td>master</td> <td>41d</td> </tr> <tr> <td>172-25-87-4-hybrid.cisco.com</td> <td>Ready</td> <td>master</td> <td>41d</td> </tr> </tbody> </table> <p>In case of a different output, collect the following logs: <code>/var/log/firstBoot.log</code> and <code>/var/log/vm_setup.log</code></p> <p>In addition, for any cluster nodes not displaying the Ready state, collect:</p> <pre>sudo kubectl describe node &lt;name of node&gt;</pre>	NAME	STATUS	ROLES	AGE	172-25-87-2-hybrid.cisco.com	Ready	master	41d	172-25-87-3-hybrid.cisco.com	Ready	master	41d	172-25-87-4-hybrid.cisco.com	Ready	master	41d
NAME	STATUS	ROLES	AGE														
172-25-87-2-hybrid.cisco.com	Ready	master	41d														
172-25-87-3-hybrid.cisco.com	Ready	master	41d														
172-25-87-4-hybrid.cisco.com	Ready	master	41d														
The following error is displayed while uploading the image:  <i>govc: The provided network mapping between OVF networks and the system network is not supported by any host.</i>	<p>The Dswitch on the datacenter is misconfigured. Please check whether it is operational and mapped to the RHEL hosts.</p>																

Scenario	Possible Resolution
VMs deploy but install fails with <i>Error: timeout waiting for an available IP address</i>	Most likely cause would be an issue in the VM parameters provided or network reachability. Enter the VM host, review and collect the following logs: <code>/var/log/firstBoot.log</code> and <code>/var/log/vm_setup.log</code>
When deploying on KVM, the following error is displayed towards the end of the VM bringup:  Error processing disk changes post-clone: <i>disk.0: ServerFaultCode: NoPermission: RESOURCE (vm-14501:2000), ACTION (queryAssociatedProfile): RESOURCE (vm-14501), ACTION (PolicyIDByVirtualDisk)</i>	Enable Profile-driven storage. Query permissions for the user at the root level (i.e. for all resources) of the KVM.

### Dual stack scenarios

**Table 4: Dual stack scenarios**

Scenario	Possible Resolution
During deployment, the following error message is displayed:  <i>ERROR: No valid IPv6 address detected for IPv6 deployment.</i>	If you intend to use a dual stack configuration for your deployment, make sure that the host machine running the Docker installer meets the following requirements: <ul style="list-style-type: none"> <li>• It must have an IPv6 address from the same prefix as the Crosswork Management IPv6 network, or be able to route to that network. To verify this, try pinging the Gateway IP of the Management IPv6 network from the host. To utilize the host's IPv6 network, use the parameter <code>--network host</code> when running the Docker installer.</li> <li>• Confirm that the provided IPv6 network CIDR and gateway are valid and reachable.</li> </ul>
During deployment, the following error message is displayed:  <i>ERROR: seed v4 host empty</i>	Ensure you use the approved version of Docker installer (19 or higher) to run the deployment.

Scenario	Possible Resolution
<p>During deployment, the following error message is displayed:</p> <p><i>ERROR: Installation failed. Check installer and the VMs' log by accessing via console and viewing /var/log/firstBoot.log</i></p>	<p>Common reasons for failed installation are:</p> <ul style="list-style-type: none"><li>• Incorrect IPv4 or IPv6 Gateway IP for either Management or Data interfaces.</li><li>• Unreachable IPv4 or IPv6 Gateway IP for either Management or Data interfaces.</li><li>• Errors in mapping the datacenter networks in MgmtNetworkName and DataNetworkName parameters in .tfvars file</li></ul> <p>Check the <code>firstBoot.log</code> file for more information, and contact Cisco Customer Experience team for any assistance.</p>

