Manage the Crosswork Network Controller VMs

You can deploy Cisco Crosswork Network Controller on a single virtual machine or as a cluster of multiple VMs. This section covers essential concepts, tasks, and troubleshooting procedures for managing virtual machine nodes in any deployment scenario.

Virtual machines in Crosswork Network Controller

A virtual machine (VM) is a compute node that

- hosts platform services and applications,
- supports both standalone and clustered deployments, and
- enables administrators to monitor, configure, and scale system resources.

In this documentation, the terms **VM** and **node** refer to the same entity and are used interchangeably. Crosswork Network Controller supports two deployment models:

- **Single VM deployment**: All system functions run on a single virtual machine, providing a streamlined management experience with limited redundancy and device capacity.

- **Cluster deployment**: Multiple VMs form a cluster, distributing workloads for scalability, high availability, and extensibility.

Administrators use the Crosswork Manager interface to:

- monitor the health and status of each VM,
- view resource consumption and operational details,
- add, update, or remove VMs as network demands change, and
- assign administrative roles for VM management tasks.

Additional reference information

- Role assignment controls user access to VM configuration settings.
- Management actions and monitoring features are available for both individual VMs and clusters.
- For advanced operational guidance, see tasks such as deploying new VMs, troubleshooting faults, and performing system recovery.

Management actions in Crosswork Manager

The Crosswork Manager interface enables administrators to monitor and manage cluster health, resources, nodes, and installed applications.

Table 1: Crosswork manager actions

Action	Description
Navigation	Use the **Crosswork Manager** window to check the health of the cluster. To access: from the main menu, choose **Administration** > **Crosswork Manager**.
Crosswork summary tab	Displays summary information about the status of nodes, the Platform Infrastructure, and the applications currently installed.
Cluster Management window	Displays node details and can be viewed only when Crosswork Network Controller is deployed as a cluster. Click on the **System summary** tile to see the node details.
System Summary window	When deployed on a single VM, allows access to details for that VM. Click on the **System summary** tile to see the VM details.

Additional notes

- In a cluster, the Cluster Management window provides summarized details about cluster health, overall resource consumption, and per-node resource utilization.
- The UI shows the IP addresses in use for each node and whether they are hybrid or worker nodes.

- On AWS EC2 deployments, the VM status may show "unknown" initially and then "initializing" after updating the inventory file—this is normal behavior for EC2 clusters.
- To see more visualizations, use the **View more visualizations** link in the top-right corner.
- To inspect node details, click on a node tile and select **View details** for components, microservices, and alarms.
- To request metrics or logs, click  under the **Actions** column and select the desired operation (such as metrics, logs, or restart microservice).
- For additional platform or application health, refer to the **Crosswork health** tab.

Common troubleshooting scenarios in cluster management

These scenarios describe common troubleshooting cases in Crosswork Network Controller cluster management and their expected behaviors:

Table 2: Troubleshooting scenarios

Scenario	Resolution
One of the Hybrid nodes is faulty in a cluster with one or more worker nodes.	Follow the *Clean system reboot* procedure described in System recovery options and requirements, on page 19.
More than one Hybrid node is faulty.	Follow the *Redeploy and recover* procedure described in System recovery options and requirements, on page 19.
Last_updated_time deviation	On the **Cluster Management** window, it is normal to see deviation on the *last_updated_time* across the nodes in the cluster based on when the data was updated. This is an expected behavior.

Additional information

- If multiple node or application faults persist after recommended recovery actions, contact the Cisco Customer Experience team for further assistance.
- When performing recovery actions, always verify backup recency and ensure the operational architecture matches the original deployment (number/type of nodes).
- For further recovery steps, see System recovery options and requirements, on page 19 for detailed actions covering VM replacement, system reboot, and redeployment.

Edit data center credentials

Update and store the current credentials for your data center.

If you changed your password after deploying Crosswork Network Controller, update the stored credentials to ensure the correct password is used when deploying the new VM.

Before you begin

Ensure you have the current credentials for your data center.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager** .

Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.

Step 3 Choose **Actions** > **View/Edit data center** to display the **Edit data center** window.

The **Edit data center** window displays details of the data center.

Step 4 Use the **Edit data center** window to enter values for the **Access** fields: Address, Username, and Password.

Step 5 Click **Save** to save the data center credential changes.

The new credentials are saved for the data center and will be used for subsequent deployments.

Add a VM to the Crosswork Network Controller cluster

Add a new VM in Crosswork Network Controller cluster to expand your cluster and handle increased workload.

As your network grows and you add more Crosswork applications, you may need to expand resources to handle increased workload. You can add a new VM to your Crosswork Network Controller cluster to scale capacity. The deployment steps are similar whether you use the UI or the API; for API details, see Crosswork Network Controller APIs. This guide describes the procedure using the UI.

Important

- If you install your cluster manually, import the cluster inventory file into Crosswork Network Controller before deploying a new VM. The **Deploy VM** option remains disabled until you complete the import. For more information, see Import the inventory file, on page 5.
- When a new Worker (or Hybrid) node is added and an existing node is subsequently deleted, the system can become unstable and many pods may enter a degraded state. This occurs because the system requires a rebalance operation after the new node is added. To avoid instability, users must manually run the **Rebalance** option from the **Actions** tab immediately after adding the Worker/Hybrid node.
- If worker nodes are deployed on an ESXi host with a down Nexus connection, the nodes may appear as successfully added but will not join the cluster. Only nodes that successfully join (for example, hybrid nodes) are shown in the UI, while the backend may still reflect the total expected count. This behavior is expected because a node tile appears in the UI only after the VM boots and joins the cluster.

Before you begin

- Gather configuration details for Crosswork Network Controller, including the management IP address.
- Collect host information for the new VM, such as the data store and data VM interface IP address.
- Decide which type of VM to add. The cluster supports a minimum of three hybrid VMs and up to two worker VMs.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager**.

Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.

Note
The **Crosswork summary** tab and **Cluster Management** window both display the status of your cluster, but there may be slight differences. The **Crosswork summary** tab shows VM status based on Kubernetes, while the **Cluster Management** window also accounts for the VM status in the data center. For example, if a worker VM deployment fails due to insufficient data center resources, the **Cluster Management** window shows its status as *degraded*, while the **Crosswork summary** window shows the status as *down*.

Step 3 Select **Actions** > **Deploy VM** to display the **Deploy VM node** window.

Step 4 Enter the required VM details and configuration.

Step 5 Click **Deploy** to begin the provisioning process.

A new VM tile appears in Crosswork Manager and displays deployment progress.

Step 6 (Optional) To monitor deployment status, use **Cluster Management** > **Actions** > **View job history**, or check the data center UI.

Step 7 If needed, rebalance cluster resources or restart processes to optimize the load on the new VM. For more information, see Rebalance cluster resources, on page 10.

Import the inventory file

Import the Day0 inventory file to enable Crosswork Network Controller to perform any datacenter-related operations.

If you want to perform any datacenter-related operations, you must first manually import the Day0 inventory file.

Attention Crosswork Network Controller cannot deploy or remove VM nodes in your cluster until you complete this operation.

Before you begin

- Ensure you uncomment the `OP_Status` parameter in your tfvars file. Otherwise, VM status may display incorrectly as **Initializing** even after VMs become functional.
- In KVM or EC2 deployments (single VM or cluster), ensure that the tfvars file includes the required details for each VM in your setup and that the `NonVcenter` flag is set to `true`.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager** .
Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.
Step 3 Choose **Actions** > **Import inventory** to display the **Import Inventory** drawer window.
Step 4 (Optional) Click **Download sample template file** to download and edit the template.
Step 5 Click **Browse** and select the cluster inventory file.
Step 6 Click **Import** to complete the operation.

The cluster inventory imports successfully, allowing Crosswork Network Controller to recognize and manage your VMs.

Export the inventory file

Export the Cisco Crosswork cluster inventory file for monitoring, management, or backup.

Use this process to download the current cluster inventory for external analysis, backup, or compliance.

Before you begin

Ensure you have administrator access to Crosswork Network Controller.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager** .
Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.
Step 3 Choose **Actions** > **Export inventory** .

Crosswork Network Controller downloads the cluster inventory gzip file to your local directory.

What to do next

Save or review the exported file as needed for your workflow.

Retry deployment for failed VMs

Retry deployment of nodes that failed due to incorrect information after correcting the details.

Node deployments with incorrect information can fail. After providing the correct details, you can retry the deployment.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager**.

Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.

Step 3 Click **Retry** on the failed node tile to display the **Deploy VM** window.

Step 4 Provide corrected information in the fields provided.

Step 5 Click **Deploy**.

Node removals

Node removals are cluster maintenance operations that

- allows administrators to delete failed or healthy nodes from a Cisco Crosswork cluster,
- eliminates the node reference from the Crosswork Network Controller cluster, and
- deletes the node from the host VM.

Node removal behaviors and limits

This topic lists the supported limits, expected effects, and actions associated with removing hybrid and worker nodes in the system.

Supported node roles and limits

- **Hybrid nodes**: The system must maintain three operational hybrid nodes at all times to ensure high availability (HA) and system protection. If one of the hybrid nodes stops functioning, Crosswork will attempt to compensate, but performance and resilience against further failures will be severely impacted. In such cases, the faulty node must be erased, and a new hybrid node should be deployed to replace it.
- **Worker nodes**: You can have up to two worker nodes. Both worker nodes can be erased without immediate consequences, but it is recommended to erase and replace them one at a time.

Effects of hybrid node removal

When a hybrid node is removed (either through an erase operation or directly from the backend), the following effects are observed:

- Remaining hybrid nodes display a "degraded" status, indicating high availability (HA) is lost.
- A further node failure could cause operational issues.
- Alarms are generated, and you are expected to restore the down node. Three functioning hybrid nodes should always be present.
- Several pods may enter the "Pending" state. This is expected because some critical infrastructure services, which run as three instances for maximum HA, are pinned to specific hybrid nodes.

- **Examples of services in the "Pending" state**: `cw-ftp`, `cw-sftp`, `nats`, `robot-etc`, `robot-kafka`, and `tyk`.

- Some pods may remain pending due to being configured as **DaemonSet**.
- Once the down hybrid node is restored, the system returns to normal and pending issues are resolved.

Effects of worker node removal

- Up to two worker nodes are supported.
- Both can be erased without immediate system impact.
- It is recommended to erase and replace worker nodes one at a time.

Note When a Worker node is removed while a vCenter alarm on that VM requires user acknowledgement, the node is deleted from the Crosswork Network Controller UI but not from the backend, causing the total count in the UI to remain incorrect and leaving the VM in vCenter. This stale backend entry can also cause new Worker node additions to fail with a duplicate-IP error. To clean up the stale entry, run this command:

```
robotctl remove-node-from-inventory <node-ip>
```

.

Manual cluster installation requirements

For manual cluster installations, you must erase the VM from the Crosswork UI and then delete it from the data center (for example, from vCenter).

Troubleshooting and escalation

If you continue to experience issues after performing these steps, contact the Cisco Customer Experience team for assistance.

Remove a node

Remove a node from Crosswork Network Controller.

Use this task to permanently erase a VM node in Crosswork Network Controller. This operation is disruptive and should be performed during a maintenance window.

Before you begin

- Erasing a node can disrupt services and block certain processes until the action completes. Perform this operation during a scheduled maintenance window.
- Removing worker or hybrid nodes increases the load on remaining nodes and may impact system performance. Contact Cisco Customer Experience before removing nodes.

Follow these steps to erase a node:

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager** > **System summary**.

Step 2 On the VM node you want to remove, click ⋮ and select **View details**.

Step 3 Click ··· and select **Erase VM node**.

Step 4 On the dialog prompt, click **Erase** to confirm the action.

Note

- During the removal of a hybrid or worker node, the Crosswork Network Controller UI may become temporarily unreachable for a short duration due to the relocation of the `robot-ui` pod to another node.
- A removed node will continue to be visible in the Grafana dashboard as an entry with only historical data.

The selected node is erased and removed from active management in Crosswork, but remains in Grafana as a historical entry.

What to do next

Review cluster performance and update operational procedures to account for the removed node.

Enable or disable maintenance mode

Use maintenance mode to temporarily suspend Crosswork Network Controller operations for maintenance or restart activities and resume normal service.

Maintenance mode provides a graceful shutdown for system updates and synchronizes application data before suspending services.

Attention It can take several minutes for the system to enter maintenance mode and to restart when maintenance mode is turned off. During these periods, users should not attempt to log in or use the Crosswork applications.

Before you begin

- Back up the Crosswork Network Controller cluster.
- Notify users, and ensure they log out. The operation cannot be canceled once started.

Procedure

Step 1 Navigate to **Administration** > **Settings** > **System Settings** > **Maintenance Mode**.

Step 2 To enable maintenance mode, set the **Maintenance mode** slider to **On**.

- When prompted, confirm the shutdown to proceed.
- Wait for the system to fully enter maintenance mode (this may take several minutes).

Note
If you plan to reboot the cluster, wait at least 5 minutes after entering maintenance mode to allow data synchronization.

Step 3 Perform required maintenance activities.

Step 4 To disable maintenance mode and resume service, set the **Maintenance mode** slider to **Off**.

When prompted, confirm the action. If you do not see a prompt but the system remains in maintenance mode, toggle it on and off again to restore applications.

Crosswork Network Controller enters or exits maintenance mode and synchronizes all application data. Users cannot access Crosswork applications during maintenance mode.

What to do next

Verify system state and notify users when service is restored.

Rebalance cluster resources

- Rebalancing ensures that workloads are evenly distributed, preventing performance bottlenecks caused by uneven resource utilization.
- Efficient resource utilization is critical for maintaining a healthy and well-performing cluster.

You can initiate rebalancing at any time through the user interface. Additionally, Crosswork Network Controller continuously monitors CPU usage across all VMs and will notify you if utilization exceeds predefined thresholds. These alarms serve as prompts to take corrective actions, such as adding more worker VMs and redistributing resources, before performance issues arise.

Rebalancing is required in these scenarios:

1. A new VM is added on day N in the cluster.
2. An existing VM is replaced on day N in the cluster.
3. A VM is down for over 5 minutes in the cluster.
4. The CPU or memory utilization of a VM constantly exceeds 95% in the cluster.

To avoid performance degradation, it is recommended to deploy new worker VMs (see Add a VM to the Crosswork Network Controller cluster, on page 4) before CPU usage exceeds 90%. However, note that when new VMs are added, active workloads are not automatically redistributed, making rebalancing a necessary step. If you already have 5 or 6 VMs in your cluster and still experience resource shortages, please reach out to the Cisco Customer Experience team for assistance.

Caution Rebalancing can take from 15 to 30 minutes during which the Crosswork Applications will be unavailable. Once initiated, a rebalance operation cannot be canceled.

To rebalance resources between the existing VMs in your cluster, follow these steps:

Before you begin

- Crosswork must be in maintenance mode before rebalancing to ensure data integrity.
- Any users logged in during the rebalancing will lose their sessions. Notify other users beforehand that you intend to put the system in maintenance mode for rebalancing, and give them time to log out. You can use the **Active Sessions** window (**Administration** > **Users and Roles** > **Active sessions** tab) to see who is currently logged in (or sessions that were abandoned and have not been cleaned up yet).

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager**.

Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.

Step 3 Click **Actions** > **Rebalance**, and the **Rebalance Requirements** are displayed. Read through the requirements and select the two check boxes once you are ready to start the rebalancing.

Figure 1: Rebalancing requirements

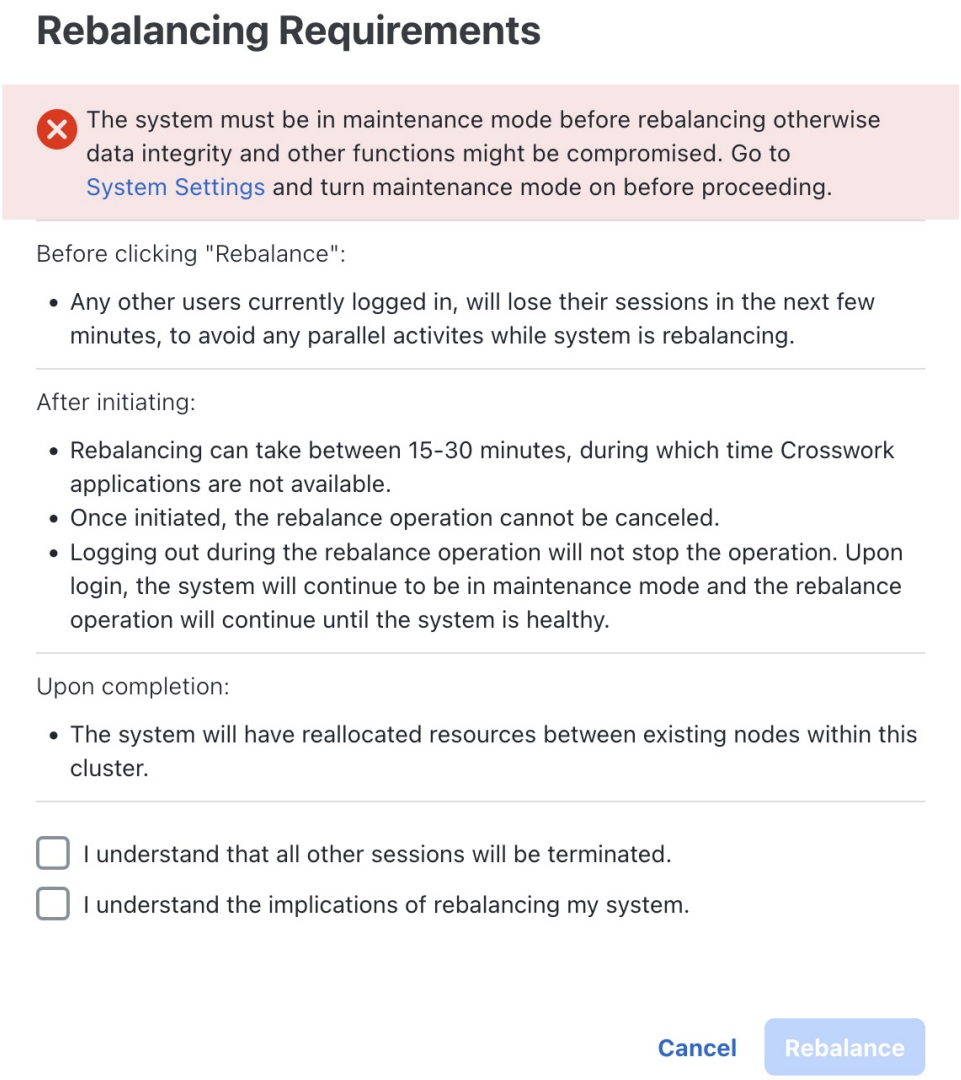

Step 4 Click **Rebalance** to initiate the process. Crosswork begins to reallocate the resources in the over utilized VM to the other VMs in the cluster.

A dialog box indicating the status of rebalancing is displayed. Kindly wait for the process to complete.

Step 5 After the rebalancing process is completed, you may see one of the following result scenarios:

- **Success scenario:** A dialog box indicating successful rebalancing operation. Follow the instructions in the dialog box to proceed further.

Figure 2: Rebalancing result - success

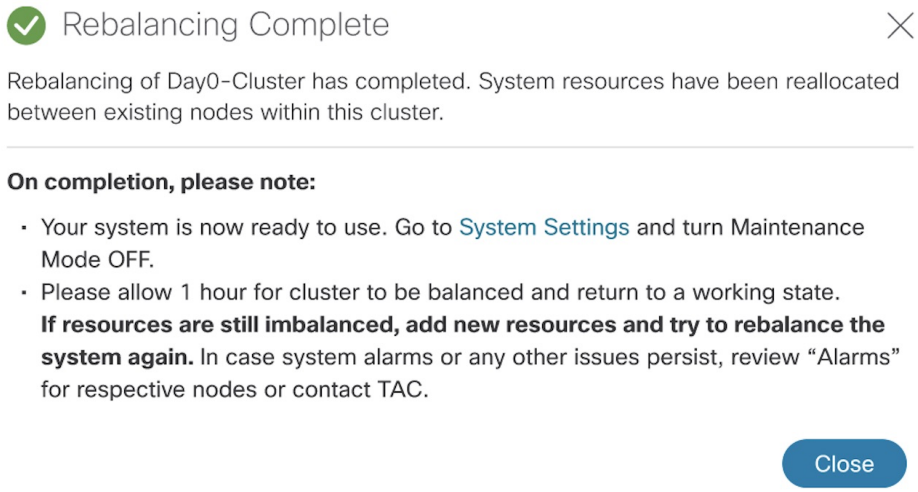

- **Failure scenario - scope available to add new worker nodes:** A dialog box indicating rebalancing failure is displayed. In this case, the system prompts you to add a new worker VM and try the rebalance process again.

Figure 3: Rebalancing result - add new worker node

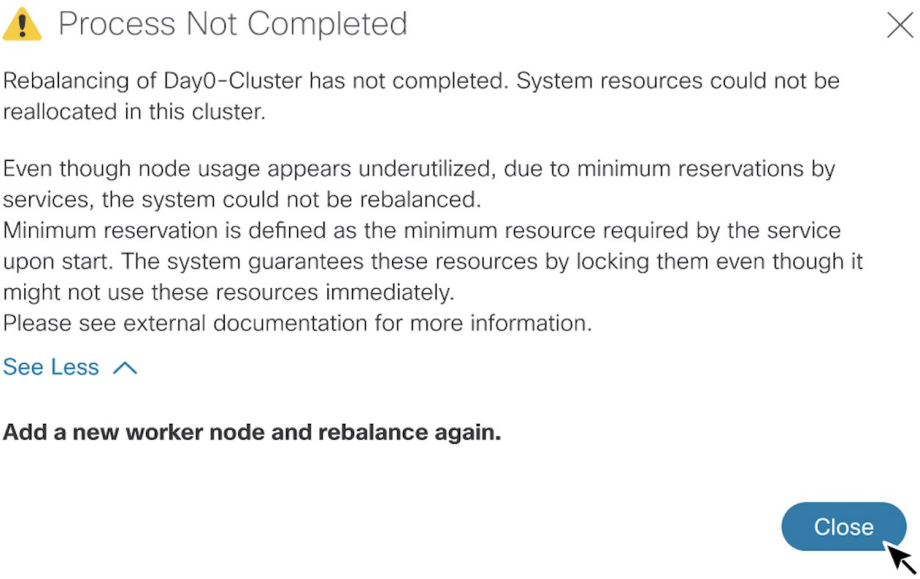

- **Failure scenario - no scope to add new worker nodes:** A dialog box indicating rebalancing failure is displayed. In this case, the system prompts you to contact the TAC as new worker VMs cannot be added.

Figure 4: Rebalancing result - contact TAC

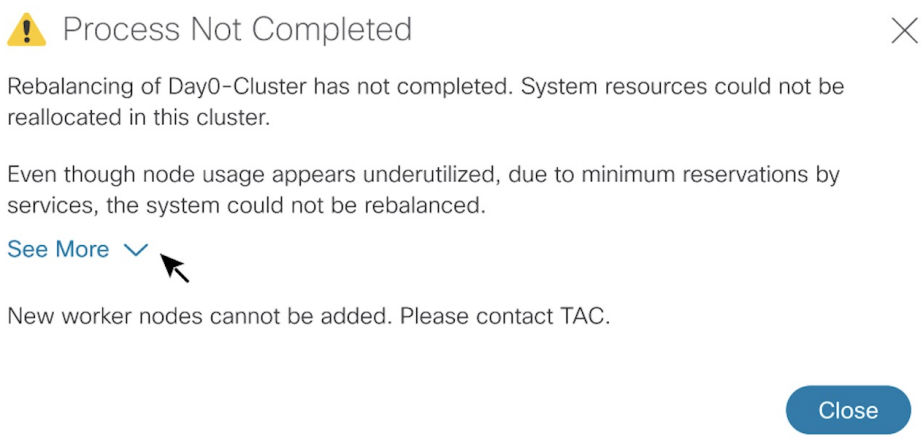

Best practice for moving workloads with placement APIs

Use these guidelines to ensure reliable workload movement in your cluster when using placement APIs, especially if the Crosswork Network Controller UI is unavailable or during VM or database recovery scenarios:

- The API method is preferred if the Crosswork Network Controller UI is not working due to high CPU utilization (>=95%) for a period of time.
- When replacing a VM containing a database, use the placement API to move the database before rebalancing workloads across the VMs.
- During a VM power-down and power-up scenario, typically the database pod recovers automatically within a few hours. If the VM is down for more than 5 minutes, redistribute resources using the placement API and rebalance the cluster.
- When moving non-core service and application workloads, exclude database services when identifying services to be moved.

Capabilities of placement APIs for workload distribution

Understand how placement APIs support manual workload movement between cluster VMs when automated or UI-based placement is unavailable.

You can use APIs to manually move database or application service workloads from one VM to other VMs in the cluster. The API method is preferred if the Crosswork Network Controller UI is not working due to high CPU utilization (>=95%) for a period of time.

Databases refer to `robot-postgres` and `cw-timeseries-db`. If a VM containing a database is replaced, the placement API must be explicitly invoked to instantiate the database on a new VM. In the event of VM replacement, the recommended order is to first use the API to move the database, followed by rebalancing to evenly distribute workloads across the VMs.

On clusters with worker VMs installed, the `robot-postgres` and `cw-timeseries-db` database services are pinned to the worker VMs, while the `local-postgres` pods are pinned to the hybrid VMs.

API example: place services for database pods

Request

```
curl --request POST --location
'https://<Vip>:30603/crosswork/platform/v2/placement/move_services_to_nodes' \
--header 'Content-Type: application/json' \
--header 'Authorization: <your-jwt-token>' \
--data '{
    "service_placements": [
        {
            "service": {
                "name": "robot-postgres",
                 "clean_data_folder": true,
                 "pin_to_node":true
            },
            "nodes": [
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-114-worker.cisco.com"
                },
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-115-worker.cisco.com"
                }
            ]
        },
        {
            "service": {
                "name": "cw-timeseries-db",
                 "clean_data_folder": true ,
                 "pin_to_node":true
             },
            "nodes": [
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-114-worker.cisco.com"
                },
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-115-worker.cisco.com"
                }
            ]
        }
    ]
}'
```

Response

```
{
    "job_id": "PJ5",
    "result": {
        "request_result": "ACCEPTED",
        "error": null
    }
}
```

API example: place services for non-core pods

Request

```
curl --request POST --location
'https://<Vip>:30603/crosswork/platform/v2/placement/move_services_to_nodes' \
--header 'Content-Type: application/json' \
--header 'Authorization: <your-jwt-token>' \
--data '{
    "service_placements": [
        {
            "service": {
                "name": "helios"

            },
            "nodes": [
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-114-worker.cisco.com"
                },
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-115-worker.cisco.com"
                }
            ]
        },
        {
            "service": {
                "name": "dg-manager"
             },
            "nodes": [
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-114-worker.cisco.com"
                },
                {
                    "name": "fded-1bc1-fc3e-96d0-192-168-5-115-worker.cisco.com"
                }
            ]
        }
    ]
}'
```

Response

```
{
    "job_id": "PJ5",
    "result": {
        "request_result": "ACCEPTED",
        "error": null
    }
}
```

Move services between cluster VMs using the placement API

Move database or application service workloads to different VMs in the cluster to address resource imbalances, high CPU utilization, or VM replacement events.

Perform this task when automated placement or the Crosswork Network Controller UI is unavailable, or during planned resource redistributions after VM replacement.

Before you begin

- Ensure you have your authorization token (`<your-jwt-token>`).
- Identify the names of services and target VMs (using Grafana or other cluster tools).
- Confirm access to the Grafana Monitoring Dashboard.

Follow these steps to move services between cluster VMs using the placement API:

Procedure

Step 1 Open the Grafana dashboard for the VM running the service using this link: *[Grafana Monitoring Dashboard](https://clusterendpoint:30603/grafana.monitoring/d/TYiQ9vgWk/platform-summary?orgId=1&refresh=1m\)*

Step 2 Identify the top five services with the highest CPU usage on the VM with the highest CPU utilization. Exclude database services by checking the pod CPU dashboard.

Step 3 Find the top three VMs with the lowest CPU utilization in Grafana.

Step 4 Use the placement API to move the top five services to the underutilized VMs.

For the required API request structure and examples, see Capabilities of placement APIs for workload distribution, on page 14.

Step 5 After moving services, monitor resource utilization in Grafana and follow the cluster rebalancing procedure as needed. For more information, see Rebalance cluster resources, on page 10.

Note
During a VM power-down and power-up, database replica recovery depends on the data size. Typically, the pod recovers on its own within a few hours. If the VM is down for more than 5 minutes in the cluster, redistribute the resources as described above and follow the cluster rebalancing procedure.

View job history

Use the **Job history** window to track the status of jobs, such as deploying a VM or importing cluster inventory.

Procedure

Step 1 From the main menu, choose **Administration** > **Crosswork Manager**.

Step 2 On the **Crosswork summary** tab, click the **System summary** tile to display the **Cluster Management** window.

Step 3 Choose **Actions** > **View job history**.

The **Job history** window displays a list of cluster jobs. You can filter or sort the **Jobs** list using the fields provided: Status, Job ID, VM ID, Action, and Users.

Step 4 Click any job to view it in the **Job details** panel at the right.

Tier upgrades

A tier upgrade is a process that:

- allows users to move from a lower tier to a higher tier in Crosswork Network Controller during the installation lifecycle,
- involves different procedures and requirements depending on whether the deployment is a cluster or a single VM, and
- supports ongoing scaling or feature expansion as business needs evolve.

For detailed information about available product tiers, see the *Release Notes for Crosswork Network Controller, Release 7.2.0*.

Note Ensure all operations are performed with minimal disruption to running workloads.

Upgrade the cluster tier

Follow these steps to upgrade Crosswork Network Controller on a cluster from a lower tier to a higher tier:

Procedure

Step 1 **Add new nodes** : Add new nodes to the cluster to accommodate more applications and resources required for the higher tier. For more information, see Add a VM to the Crosswork Network Controller cluster, on page 4

Step 2 **Move databases** : Move databases to worker nodes to optimize performance. For more information, see Capabilities of placement APIs for workload distribution, on page 14 .

Step 3 **Rebalance pods across nodes** : Use the rebalance feature to redistribute pods across new nodes and restore pod balance after any prolonged node shutdowns or power-ups. For more information, see Rebalance cluster resources, on page 10 .

Step 4 **Redeploy Data Gateway from Standard to Extended for higher tiers (Advantage, Premier)** : Put the Data Gateway in **Maintenance** mode by removing it from the pool and changing its role to **Unassigned** before redeploying. For more information, see Redeploy a Data Gateway VM and Change the administration state of a Data Gateway .

- **For protected pools** :
  a. Start the redeployment with the Data Gateway that has the role **Spare**, if the pool contains one, to minimize downtime for collections.
  b. Add the re-deployed Data Gateway back to the pool.
  c. Initiate a failover so the re-deployed Data Gateway becomes **Assigned** and resumes collections.
  d. Move the other Data Gateway (its role becomes **Spare** after the failover) out of the pool and redeploy it.
- **For unprotected pools** : Move the Data Gateways out of the pool and redeploy them. Collections may stop temporarily until the redeployment completes and the Data Gateways resume processing collection jobs.

Step 5 **Update the number of devices per Data Gateway based on tier** : Reduce the number of devices per Data Gateway as you move to a higher tier to align with the tier's requirements.

Upgrade the single VM tier

Follow these steps to upgrade Crosswork Network Controller on a single VM from a lower tier to a higher tier:

Procedure

Step 1 Create a backup of the current VM to secure all data. For more information, see Manage Backup and Restore .

Step 2 Deploy the higher tier build on a new VM. For installation instructions, see the *Install Cisco Crosswork Network Controller on a Single VM* chapter in *Cisco Crosswork Network Controller 7.2 Installation Guide* .

Step 3 Restore the data from the backup to the newly deployed VM.

Cluster system recovery

A cluster system recovery is a disaster recovery strategy that

- restores critical cluster services and data after failures or disruptions,
- addresses platform-specific considerations to ensure compatibility and resilience, and
- minimizes overall downtime to maintain business continuity.

A robust cluster system recovery approach helps ensure that Cisco Crosswork clusters can be restored quickly and reliably after failures, disruptions, or disasters. Understanding your recovery options and platform-specific requirements is essential to maintaining service continuity and minimizing downtime.

System recovery options and requirements

Successful cluster recovery depends on understanding platform requirements, backup practices, and the nature of the failure. This reference summarizes prerequisites, platform limitations, and actions for common recovery scenarios.

Before you begin

- For cluster recovery, it is essential to have a recent backup.
- The cluster you are restoring should have the same operational architecture, including the same number of hybrid and worker nodes.

Recovery conditions and system behavior

- At some time during normal operations of your Cisco Crosswork cluster, you may need to recover the entire system. This can result from malfunctioning nodes, services, applications, or a disaster destroying hosts for the cluster.
- A functional cluster requires a minimum of three hybrid nodes. These nodes share processing and traffic loads for management, orchestration, and infrastructure services.

- The hybrid nodes are highly available and can redistribute processing among themselves and to worker nodes automatically.
- The cluster can tolerate one hybrid node reboot (graceful or ungraceful); the system remains functional but with degraded availability.
- The system can tolerate any number of failed worker nodes (with degraded availability until restored).
- If two or more hybrid nodes are lost ("double fault"), recovery cannot be guaranteed – in such cases, redeploy a new cluster and restore from a recent backup.

Alarms and troubleshooting

- Cisco Crosswork generates alarms when nodes, applications, or services malfunction.
- Examine alarms and check health of the affected component(s). Use Crosswork features to drill down and, for service faults, attempt to restart the problem service.
- If alarms show a single hybrid node, or a hybrid plus worker node(s) failure, start by rebooting or replacing (erasing, then readding) failed nodes; if unsuccessful, attempt a clean system reboot.
- If the system remains unstable or degraded (loss of two or more hybrid nodes), deploy a new cluster and recover using a backup.

Platform limitations

- Unintentional VM shutdown is not supported on a 3 VM cluster running Crosswork Network Controller. If a VM fails, the remaining two VMs cannot support migrating all pods from the failed VM. Add worker nodes to enable VM shutdown.
- A reboot of one VM is supported in a 3 VM cluster. Restore may take 5 minutes (if the `orch pod` is not on the rebooted VM) up to 25 minutes (if it is).

Perform a clean system reboot (VMware)

Perform a coordinated reboot of all cluster VMs to restore operations or after failure.

A clean system reboot is sometimes required to restore cluster health following multiple node or service issues, or after system maintenance. This process ensures all VMs are properly powered down and brought back online in a specific order, supporting the stability and recovery of both hybrid and worker nodes in VMware deployments.

Follow these steps to perform a clean system reboot:

Procedure

Step 1 Place Crosswork Network Controller in Maintenance mode. See Enable or disable maintenance mode, on page 9 for details.

a) (Optional) Shut down Crosswork Data Gateways and other non-essential components, such as NSO and SR-PCE, that communicate with Crosswork.

Step 2 Power down all VMs:

a) Log in to the VMware vSphere Web Client.
b) In the **Navigator** pane, right-click the VM you want to shut down.
c) Choose **Power** > **Power Off**.
d) Wait for the VM's status to change to **Off**.
e) Repeat for each VM in the cluster.

Step 3 Power up the VM hosting the first hybrid node:

a) In the **Navigator** pane, right-click the VM to power up.
b) Choose **Power** > **Power On**.
c) Wait for the VM's status to change to **On**, then wait 30 seconds before continuing.

Step 4 Repeat the previous step for each remaining hybrid node, staggering reboots by 30 seconds. Continue with each worker node using the same staggered interval.

Step 5 After all VMs are powered on, wait a few minutes and login to Crosswork Network Controller.

Step 6 Move Crosswork Network Controller out of maintenance mode. See Enable or disable maintenance mode, on page 9 for details.

- If your cluster is not healthy, maintenance mode attempts may fail. Alarms may indicate failed services and reasons.
- If issues persist, follow the "redeploy and restore" method. For more details, see Redeploy and restore a Crosswork cluster from backup (VMware), on page 21.

Step 7 Restart Crosswork Data Gateways and any other components in your ecosystem that communicate with Crosswork Network Controller.

The Crosswork Network Controller cluster completes a clean system reboot. If cluster health does not return, proceed with the redeploy and restore procedure.

Redeploy and restore a Crosswork cluster from backup (VMware)

Rebuild and restore a failed Crosswork cluster using a previously taken backup.

Redeployment and restoration from backup is required when a cluster is severely degraded (such as after double faults or catastrophic failures), and cannot be recovered through standard node replacement or reboot procedures. The procedure involves powering down and deleting existing VMs, deploying a new cluster, and then restoring system state from a backup to recover services and data.

Before you begin

- Ensure you have a recent and valid backup file.
- This method assumes you have taken periodic backups before recovery is required. (For details on backup, see Back up data.

Follow these steps to redeploy and restore the cluster:

Procedure

Step 1 Power down all VMs:

a) Log in to the VMware vSphere Web Client.

b) In the **Navigator** pane, right-click the VM you want to shut down.
c) Choose **Power** > **Power Off**.
d) Wait for the VM's status to change to **Off**.
e) Repeat for each VM in the cluster.

Step 2 Delete all VMs:

a) In the VMware vSphere Web Client **Navigator** pane, right-click the VM you want to delete.
b) Choose **Delete from Disk**.
c) Wait for the VM's status to show **Deleted**.
d) Repeat for each VM in the cluster.

Step 3 Deploy a new Cisco Crosswork cluster as explained in the *Cisco Crosswork Network Controller 7.2 Installation Guide*.

Step 4 Recover the system state to the newly deployed cluster. For more information, see Restore data after a disaster.

A new Crosswork Network Controller cluster is deployed, and system state is restored using the most recent backup.

Shut down and restart the standby cluster safely

Safely shut down the standby cluster without Maintenance Mode and bring it back online with data consistency.

In geo HA deployments, the standby cluster can be shut down while data continues syncing from the active cluster, ensuring consistency without the need for Maintenance Mode.

Before you begin

- Ensure you do not place the standby cluster in Maintenance Mode.
- Verify that data is syncing from the active cluster.

Follow these steps to shut down and restart the standby cluster safely:

Procedure

Step 1 Shut down the standby cluster without placing it in Maintenance Mode.

Step 2 Keep the active cluster running so it continues syncing data during the shutdown.

Step 3 Power on the standby cluster when needed to start its automatic recovery.

Step 4 Wait for the standby cluster to become fully healthy, which may take about 20–40 minutes.

Step 5 Trigger an on-demand sync from the active cluster or wait for the next periodic sync.

The standby cluster returns to a fully healthy state with all data synchronized from the active cluster.

Collect cluster logs and metrics

Monitor or audit Cisco Crosswork cluster components by collecting and managing periodic logs and metrics for each cluster component.

Collecting logs and metrics helps administrators track the health and performance of the Cisco Crosswork cluster, including its nodes and microservices. Use this task for troubleshooting or routine audits.

Note Showtech logs must be collected separately for each application.

Before you begin

- Ensure you have administrator access to Cisco Crosswork Manager.
- Know which components (cluster, node, or microservice) you want to collect logs or metrics from.

Procedure

Step 1 From the main menu, select **Administration** > **Crosswork Manager**.

Step 2 On the **Crosswork summary** tab, select **System summary** to open cluster management.

Step 3 To collect logs and metrics for the entire cluster, click **Actions** and choose a showtech option:

- **Request all**: Collect both logs and metrics.
- **Request metrics**: Collect only metrics.
- **Collect logs**: Collect only logs.

Step 4 To collect logs or metrics for a specific node:

a) Select the node.
b) Click **Showtech options** and choose a showtech operation.

Step 5 To collect logs or metrics for an individual microservice on a node:

a) Under **Actions** for the desired microservice, click  and select a showtech operation.

Step 6 To view the status of showtech jobs, select **Actions** > **View showtech jobs**. Under the **Action** column, use  menu to:

- **Publish** a completed showtech log.
- **Delete** a showtech log.
- **View details** of a job.

The system collects and displays the requested logs and metrics for your chosen cluster component. Collected showtech logs are available for audit, troubleshooting, or compliance verification.

What to do next

Review collected logs and metrics as needed, and publish or delete showtech logs to manage storage or share information with other stakeholders.

Crosswork Network Controller containers

To give users a single reference that explains what each container does at a basic level, helping them identify components quickly when collecting logs or investigating issues.

Attention This topic includes information only for the containers that were available at the time of publication. It does not represent a complete list of all system containers.

Table 3: Crosswork Network Controller containers

Container name	Container role
robot-ui	This container provides the user interface. In a clustered environment, multiple instances run for resiliency. It typically starts after the core services are up, ensuring that all required processes are available before users can log in.
robot-dlminvmgr	The device lifecycle manager (dlm) tracks devices as they are onboarded to Crosswork Network Controller and monitors their health through basic reach checks.
robot-kafka	Kafka is an open-source messaging system used by Crosswork Network Controller services to process large volumes of streaming data.
nats	The nats is a lightweight, open-source messaging system used by Crosswork Network Controller services.
robot-etcd	The etcd is an open source key value database used by services on Crosswork Network Controller.
descheduler	The descheduler runs on demand and is responsible for moving services to help balance container placement and optimize resource usage across the nodes.
robot-orch	The orchestrator manages infrastructure services, including application lifecycle operations, backup and restore, geo HA functions, node management, and cluster management.
docker-registry	This component is part of the Crosswork Network Controller application lifecycle and handles installing and uninstalling Crosswork Network Controller applications within the cluster, servicing Kubernetes requests for Docker images.
cw-sftp	This component provides an SFTP server for applications that need to download device-related files during the application lifecycle.

Container name	Container role
cw-ftp	This component provides an FTP server for applications that need to download device-related files during the application lifecycle.
cw-ipsec	This component encrypts pod-to-pod communication across nodes.