Cisco Crosswork Network Controller 7.2.x Installation Guide

Cisco Crosswork Network Controller
Release: 7.2.x | Updated August 14, 2026

© 2026 Cisco Systems, Inc. All rights reserved.

Full Cisco Trademarks with Software License

Full Cisco Trademarks with Software License

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

Topics included

Full Cisco Trademarks with Software License....ii

Part I: Get Started....11

1 Overview....13

About this guide....14
Audience....14
Cisco Crosswork Network Controller....14
Cisco Crosswork Infrastructure....15
Crosswork Data Gateway....15
Cisco integrated components....15
Components that support integration with Crosswork Network Controller....16
Security....16

2 Plan Your Deployment....17

Prepare for installation....18
Determine how many VMs you need....18
Identify the resource footprint....20
VM requirements....21
Resource footprint for VMware....24
Resource footprint for AWS EC2....25
Resource footprint for KVM....26
Resource footprint for single VM deployments....28
Physical server requirements....29
Special considerations....32

3 Choose Installation Workflow....33

Overview....34
Install Cisco Crosswork Network Controller on VMware vCenter....34
Install Cisco Crosswork Network Controller on KVM....36
Install Cisco Crosswork Network Controller on AWS EC2....38
Install Cisco Crosswork Network Controller on a single VM....40

4 Installation Prerequisites....43

Installation requirements overview....44
VM requirements....44
Crosswork cluster VM requirements....44

Crosswork Data Gateway VM requirements....46
Validate storage performance and review diagnostics....49
Crosswork network ports and communication details....52
IP address restrictions....70
Supported network topology models....72

Part II: Install Cisco Crosswork Network Controller on VMware vCenter....79

5 Install Cluster on VMware vCenter....81
Installation overview....82
VMware settings....82
Installation parameters....83
Automate application installation....92
Install Cisco Crosswork on VMware vCenter using the Docker installer tool....94
Manual installation of Cisco Crosswork using vCenter vSphere UI....102
Import the CA certificate chain on VMware....102
Build the OVF template....103
Deploy the template....109
Monitor cluster activation....113
Diagnostic assessment....115
IOPs calculation....118
Log into the Cisco Crosswork UI....120
Import cluster inventory....122
Troubleshoot the cluster....122

6 Install Cisco Crosswork Data Gateway on VMware vCenter....129
Crosswork Data Gateway workflows....130
Crosswork Data Gateway parameters and deployment scenarios....131
Install Crosswork Data Gateway using vCenter vSphere Client....144
Install Crosswork Data Gateway via the OVF Tool....154
Log in and Log out of the Data Gateway VM....156
Access Crosswork Data Gateway VM from SSH....156
Access Crosswork Data Gateway through vCenter....157
Log out of Crosswork Data Gateway VM....157
Crosswork Data Gateway authentication and enrollment....158
Crosswork Data Gateway post-installation tasks....159
Configure the timezone of the Crosswork Data Gateway VM....159
Troubleshoot installation and enrollment....160
Import controller signing certificate file....164
View the controller signing certificate file....164

Part III: Install Cisco Crosswork Network Controller on KVM....167

7 Install Crosswork Cluster on KVM.....169
Installation overview.....170
KVM host bare metal requirements.....170
Installation parameters.....171
Install Crosswork Network Controller cluster using CLI.....178
Preliminary checks.....179
Set up and validate KVM on RHEL.....179
Configure network bridges or SRIOV.....180
Install Crosswork Network Controller cluster on KVM using CLI.....184
Monitor cluster activation.....187
Log into the Cisco Crosswork UI.....188
Troubleshoot the cluster.....190
Installation workflow for deploying SR-PCE (Cisco IOS XRv 9000) on RHEL 8.10 KVM.....193
Prerequisites and setup for Cisco IOS XRv 9000 virtualization.....193
Configure network bridges for SR-PCE.....194
Prepare VM storage.....195
Update the VM configuration.....195
Deploy and verify the VM.....196
Connect to the router console.....196

8 Install Crosswork Data Gateway on KVM.....199
Crosswork Data Gateway on KVM.....200
Installation workflow of Crosswork Data Gateway.....200
Parameters required for Crosswork Data Gateway installation.....201
Sample configuration file for IPv4.....211
Sample configuration file for IPv6.....213
Install Crosswork Data Gateway using the CLI.....216
Access and manage the Crosswork Data Gateways on KVM.....220

Part IV: Install Cisco Crosswork Network Controller on AWS EC2.....223

9 Install Cisco Crosswork Network Controller on AWS EC2.....225
Installation overview.....226
Amazon EC2 requirements for Crosswork.....226
Extract CF template image.....229
Roles and policy permissions.....231
CloudFormation (CF) template parameters.....232
CF template parameters for installing Cisco Crosswork cluster VMs.....232
CF template parameters for installing Crosswork Data Gateway.....238
CF template parameters for installing a single hybrid cluster or worker node.....241
Module-based deployment method.....244
Cisco Crosswork cluster deployment on Amazon.....244

Crosswork Data Gateway deployment on Amazon....245
Additional Crosswork cluster node deployment....250
CloudFormation template deployment and management....251
Deploy a CF template....251
Monitor the installation....252
Access the Crosswork UI....253
Crosswork Data Gateway post-installation tasks....254
Configure the time zone of the Data Gateway VM....254
Data Gateway VM login and logout....256
Installation and enrollment troubleshooting....257

Part V: Install Crosswork Network Controller Applications....263

10 Install Crosswork Applications....265
Install Crosswork Network Controller applications....266
Add the auto-action file....269

Part VI: Integrate Cisco NSO and SR-PCE with Cisco Crosswork Network Controller....271

11 Integrate Cisco NSO....273
NSO integration workflow....274
Configure standalone NSO....274
Restrict sync-to access for the Crosswork NSO user....276
Add a standalone Cisco NSO provider....277
Configure Cisco NSO for layered service architecture....280
Add Cisco NSO providers for layered service architecture....282
Install Cisco NSO function pack bundles from Crosswork UI....286
View NSO function pack bundles....286
Manage NSO function pack bundles....287
Deploy NSO function pack bundles....289
NSO function pack installation troubleshooting....292
View NSO function pack job history....293
Install NSO function packs manually....294
Configure JSON web token authentication for NSO....294

12 Integrate SR-PCE....297
SR-PCE integration workflow....298
SR-PCE certificate configurations....298
Configure SR-PCE with default certificates....298
Configure SR-PCE with custom certificates....299
Configure SR-PCE....307

Sample SR-PCE configurations....313
Add Cisco SR-PCE providers....322

Part VII: Install Cisco Crosswork Network Controller on a Single VM....329

13 Install Crosswork Network Controller on a Single VM....331
Introduction....332
Installation requirements....332
Installation parameters....335
Install Crosswork Network Controller using the vCenter vSphere UI....342
Install Crosswork Network Controller via the OVF Tool....348
Install Crosswork Network Controller using the Docker installer tool....351
Install Crosswork Network Controller VM using CLI....357
Preliminary checks....358
Set up and validate KVM on RHEL....358
Configure network bridges or SRIOV....359
Install Crosswork Network Controller VM on KVM using CLI....363
Post installation tasks (for Premier tier installation)....365

Part VIII: Upgrade Cisco Crosswork Network Controller....367

14 Upgrade Cisco Crosswork Network Controller....369
Upgrade overview....370
Upgrade requirements....372
Save AWS EC2 configuration....372
Disable network bandwidth management feature packs....372
Upgrade using existing hardware....374
Shut down the Data Gateway VMs....375
Create backup and shut down....376
Install the latest version of the cluster....379
Install the applications....380
Migrate the backup....381
Upgrade Crosswork Data Gateway....382
Post-upgrade checklist....386
Upgrade using parallel hardware....387
Deploy a new cluster....388
Back up the cluster....389
Update DNS server and run migration....392
Add Crosswork Data Gateway to Cisco Crosswork....393
Shut down the old cluster....395
Upgrade Crosswork Network Controller applications (standalone activity)....396
Upgrade Crosswork Network Controller on a single VM....398
Migration prerequisites....398

Back up the current release data....399
Install the newer Crosswork Network Controller version....399
Run data migration....400

15 Security Patches for Crosswork Network Controller....403
Security patches for OS and core components....404
Add a security OS patch....405
Crosswork Data Gateway patches and patch management....406
Precautions during Data Gateway upgrade....406
Apply patch upgrades to Crosswork Data Gateway....406

Part IX: Uninstall Cisco Crosswork Network Controller....411

16 Uninstall Crosswork Network Controller....413
Uninstall the Crosswork cluster....414
Delete the VM using the Docker installer tool....414
Delete the VM using the vSphere UI....415
Uninstall Crosswork Data Gateway....415
Delete the Data Gateway VM from Cisco Crosswork....415
Uninstall Crosswork applications....416

Part X: Enable Geo Redundancy....419

17 Geo Redundancy Overview....421
Geo redundancy....422

18 Geo Redundancy Requirements....425
Crosswork cluster requirements....426
Crosswork Data Gateway requirements....427
Geo redundancy port requirements....427
Unified endpoint requirements....433
Data store roles....433

19 Deploy the Arbiter VM....435
Auto-arbitration in Crosswork Network Controller....436
Arbiter VM requirements....437
Deploy an arbiter VM on VMware....437
Deploy an arbiter VM on KVM....439

20 Enable Geo Redundancy Solution....441

Geo redundancy inventory file guidelines....442
Geo redundancy workflow (Day 0)....443
Perform connectivity checks....450
Enable geo redundancy....452
DNS traffic redirect binary requirements....456
Configure cross cluster settings....458
Configure cross cluster notification settings....463
View cross cluster status....463
Geo redundancy workflow (Day N)....466
Update or add an application after enabling geo redundancy....476
Geo redundancy scenarios....477
Standby cluster behavior....478
Arbiter VM behavior....480

21 Upgrade to Geo Redundancy Solution....483
Cluster: upgrade from non-geo redundant to geo redundant deployment....484
Convert single instance NSO to NSO HA....485
Update providers....486
Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy....486
Configure controller IP for Crosswork Data Gateway....487
Complete geo redundancy enablement....488
Cluster: upgrade an existing geo redundant deployment....489
Create backup of the Cisco Crosswork cluster....493
Install the Crosswork Network Controller cluster and applications....496
Multi cluster upgrade....497
Upgrade Crosswork Data Gateway to 7.2.x on the standby site....498
Update providers....500
Enable geo redundancy on the Site 2 cluster....500
Shut down Crosswork Data Gateway in the active site....501
Perform switchover between sites....502
Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy....503
Enable geo redundancy on the Site 1 cluster....505
Arbiter VM migration....506
Single VM: upgrade from non-geo redundant to geo redundant deployment....509

22 Geo Redundancy Switchover....511
Geo redundancy switchovers....512
Auto-switchover triggers....512
Data Gateway failure triggers considerations, conditions, and alarms....513
How auto-arbitration works....514
Perform switchover manually....515
Crosswork Optimization Engine license count after a switchover....518

23 Geo Redundant Cluster Disaster Recovery.....521
Disaster recovery for geo redundant clusters.....522
Recover the active cluster.....522
Recover the standby cluster.....522
Recover after an active cluster disaster restore.....523

Appendix A: Sample Deployment Templates.....525
Crosswork Network Controller cluster deployment templates for VMware vCenter.....526
Crosswork Network Controller deployment templates for KVM.....535
Crosswork Network Controller VM templates for single VM deployments.....546
Crosswork Data Gateway deployment templates and scripts.....548
Sample script for Crosswork Data Gateway dual-stack deployment with geo-redundancy.....548
Sample script for Crosswork Data Gateway dual-stack deployment.....549
Sample script for deploying Crosswork Data Gateway with IPv6 address.....551
Sample script for deploying Crosswork Data Gateway with IPv4 address.....553
Sample auto action templates.....555
Geo redundancy templates.....558
Sample cross cluster inventory templates.....558
Sample multi-cluster upgrade inventory template.....565

Part I

Get Started

Topics:

- Overview
- Plan Your Deployment
- Choose Installation Workflow
- Installation Prerequisites

1 Overview

Topics:

- About this guide
- Audience
- Cisco Crosswork Network Controller
- Security

Summarizes the scope, intended audience, controller component descriptions, integration information, and security resources covered in the overview chapter of the Cisco Crosswork Network Controller installation guide.

This chapter introduces the scope and intended audience. It also summarizes Crosswork Network Controller components, integration information, and security guidance.

About this guide

Outlines the installation and upgrade scope of this guide for Cisco Crosswork Network Controller Releases 7.2, 7.2.1, and 7.2.1 SP, and identifies integrated-component installation as outside its coverage.

This guide provides the requirements and procedures to install or upgrade the Cisco Crosswork Network Controller solution.

Release applicability

Note

This guide applies to Cisco Crosswork Network Controller Releases 7.2, 7.2.1, and 7.2.1 Security Patch (SP). Use this guide to migrate from Release 7.2 or 7.2.1 to Release 7.2.1 SP. Release 7.2.1 SP supersedes the earlier releases.

This guide does not cover installing integrated components, such as Cisco NSO or Cisco SR-PCE. These components may already be installed or used independently. Refer to their installation documentation for details.

Audience

Provides the intended reader profile and lists assumed experience with Docker, Python scripts, VMware deployment tools, Amazon Web Services, CloudFormation templates, and KVM networking.

This guide is for experienced network users and operators who install Crosswork Network Controller in their network.

Assumed knowledge

This guide assumes that you are familiar with these areas:

- You can use a Docker container.
- You can run scripts in Python.
- You can deploy an OVA file using VMware vCenter.
- You can deploy an OVA file using VMware OVF tool.
- You are familiar with Amazon Web Services (AWS), Amazon EC2 concepts, and creating CloudFormation templates.
- You are familiar with Kernel-based Virtual Machine (KVM) functions and network configuration.

Cisco Crosswork Network Controller

Explains how Cisco Crosswork Network Controller automates intent-based service provisioning, health monitoring, and optimization across multivendor networks through a common GUI and API.

Cisco Crosswork Network Controller is a transport SDN controller. It simplifies and automates intent-based service provisioning, health monitoring, and optimization across multivendor networks. It supports these functions through a common GUI and API.

Operational functions

Crosswork Network Controller simplifies operational workflows by consolidating service lifecycle and device management functions in one integrated solution.

For details about solution components and packages, refer to the *Release Notes for Crosswork Network Controller, Release 7.2.1 SP*.

Cisco Crosswork Infrastructure

Explains the Crosswork Infrastructure platform, its VMware, AWS EC2, and KVM deployments, and the single VM or cluster options and resiliency differences described in this guide.

Cisco Crosswork Infrastructure is a resilient and scalable platform on which Crosswork components can be deployed.

Deployment options

The infrastructure uses a cluster architecture to support extensibility, scalability, and high availability. It supports deployment in VMware, AWS EC2, and KVM environments.

Crosswork Network Controller can run on one VM with all core functions, but at a limited scale. Larger deployments can use VM clusters, with or without geo redundancy, when they require greater resiliency. Refer to Determine how many VMs you need for deployment resource and capability details.

A single VM deployment operates supported functions on one machine, offering limited redundancy.

For more information about the deployment options, refer to Plan your deployment.

Crosswork Data Gateway

Explains how Crosswork Data Gateway collects information from managed devices, forwards it to Crosswork Network Controller or external destinations, and supports component data needs.

Crosswork Data Gateway collects information from managed devices and forwards it to Crosswork Network Controller. It can also forward data to external destinations.

Functions and deployment information

Crosswork applications analyze collected information for topology visualization, service health monitoring, element management, and network performance optimization. Third-party applications that use forwarded data are outside the scope of this guide.

Forwarding data to external destinations requires an additional license. For licensing and destination details, refer to the *Crosswork Data Gateway setup, management, and troubleshooting* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Data Gateway count depends on the number of devices, data volume, network topology, and redundancy requirements. Deploy each Data Gateway on an individual VM. For deployment guidance, consult the Cisco Customer Experience team.

Crosswork Data Gateway is an integral part of the Crosswork solution and does not require a separate license. This guide explains Data Gateway as a foundational component that must be installed with the Crosswork Infrastructure.

For single VM deployments, Data Gateway functions are available through Embedded Collectors. The Crosswork Network Controller single VM package includes these collectors. A separate Embedded Collectors package is not needed.

Cisco integrated components

Explains how Cisco NSO and Cisco SR-PCE support device configuration, maintenance, telemetry analysis, and optimal path computation for Crosswork Network Controller.

Cisco NSO and Cisco SR-PCE integrate with Crosswork Network Controller for device configuration, telemetry analysis, and path computation.

Component roles

Cisco Crosswork Network Service Orchestrator (NSO) provides device configuration for Crosswork Network Controller. It configures devices for their expected functions. NSO can optionally configure MDT sensor paths for data collection. It also provides device management, configuration, and maintenance services.

For single VM deployments with the Crosswork Network Controller Advantage tier, the Advantage package includes Embedded NSO. Embedded NSO provides NSO functions, so no separate installation is needed.

Cisco Segment Routing Path Computation Element (SR-PCE) runs on an IOS-XR physical or virtual device. SR-PCE supports Segment Routing Traffic Engineering (SR-TE) and Resource Reservation Protocol Traffic Engineering (RSVP-TE). Crosswork Network Controller uses Cisco SR-PCE telemetry and data to compute optimal paths for TE tunnels. It also uses this data to discover network devices.

Components that support integration with Crosswork Network Controller

Lists supported integration component types, including authentication servers, a DHCP server for Crosswork ZTP, external data destinations, and backup storage using SCP.

The listed component types support integration with Crosswork Network Controller.

Integration components

Supported integration component types include:

- TACACS+, LDAP, and RADIUS servers. For more information, refer to *External Authentication Integration* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
- DHCP server when using Crosswork ZTP.
- External Kafka for external data collection destinations.
- External gRPC for external data collection destinations.
- Storage server that supports SCP for backup storage.

Security

Explains Cisco product security commitments, the stated resources for product validation and advisories, and how to report security questions to Cisco Customer Experience.

Security is an end-to-end commitment.

Security information and support

Cisco works to ensure that its products conform to the latest industry recommendations and can help secure your entire environment. Work with your Cisco account team to review the security profile of your network.

For details about product validation, refer to *Cisco Secure Products and Solutions* and *Cisco Security Advisories*.

If you have security questions or concerns about a Cisco product, open a case with the Cisco Customer Experience team. Include details about the tool and any vulnerabilities it reports.

2 Plan Your Deployment

Topics:

- Prepare for installation
- Determine how many VMs you need
- Identify the resource footprint
- Special considerations

Summarizes the planning information for deploying Crosswork Network Controller 7.2.x.

Plan the deployment by identifying VM counts and resource requirements for the selected platform.

The chapter covers VM requirements, platform resource footprints, physical server requirements, and special considerations.

Prepare for installation

Explains the planning decisions required before installation on VMware, KVM, or AWS EC2.

This section explains the decisions you need to make before installing Crosswork Network Controller solution on your preferred platform.

Summary

You make the decisions needed before installing Crosswork Network Controller on your preferred platform.

Workflow

These are the stages for preparing to install Crosswork Network Controller.

1. You determine how many VMs you need. Determine how many VMs you need on page 18
2. You identify the resource footprint. Identify the resource footprint on page 20
3. You review special considerations. Special considerations on page 32

What's next

After completing the planning , follow the relevant installation workflow steps for your platform:

- **For VMware vCenter**:Install Cisco Crosswork Network Controller on VMware vCenter on page 34
- **For KVM**:Install Cisco Crosswork Network Controller on KVM on page 36
- **For AWS EC2**:Install Cisco Crosswork Network Controller on AWS EC2 on page 38

> **Note**
>
> **AWS EC2 support for Crosswork Network Controller**:
>
> - Crosswork Network Controller supports cluster deployments on AWS EC2 up to the Advantage license tier.
> - Single VM deployments of Crosswork Network Controller, regardless of license tier, are not supported on AWS EC2.
> - Geo High Availability deployments are not supported on AWS EC2.

> **Note**
>
> Starting with the Crosswork Network Controller version 4.1, Crosswork deployment is no longer supported for the Cisco CSP platform. For more information, refer to End-of-Life Announcement for the Cisco Cloud Services Platform Operating System.

Determine how many VMs you need

Lists the VM and Data Gateway counts required for each Crosswork Network Controller tier.

Note

This topic is applicable only when you need to deploy Crosswork Network Controller as a cluster. In case of single VM deployment, refer to Install Crosswork Network Controller on a Single VM

Note

Geo redundancy solution requires double the number of VMs. For more information, refer to Enable Geo Redundancy Solution section.

Crosswork Network Controller is available in these tiers:

Table 1: Crosswork Network Controller tiers

Cisco Crosswork Network Controller tier [1]	Contents	Crosswork Data Gateway deployment [2]	Minimum number of VMs needed [3]
Essentials	Element Management Functions	**On-Premise Standard**(default): Collectors only.	**3 Hybrid nodes**
	Path Analytics	**On-Premise Standard**	
Advantage	Crosswork Optimization Engine	**On-Premise Standard**(default): Collectors only.	**3 Hybrid nodes + 2 Worker nodes**
	Crosswork Active Topology	**On-Premise Standard**(default): Collectors only.	
	Crosswork Service Health	**On-Premise Extended**: Collectors and offload services.	

[3] The number of VMs mentioned is the minimum count. You can add more Worker nodes (maximum up to three worker nodes) as needed. If your requirements exceed the minimum count, please contact the Cisco Customer Experience team.

[2] The VM resource requirements for Crosswork Data Gateway are different for each type and cannot be modified. Therefore, if your requirements change, you must re-deploy the Crosswork Data Gateway to move from one type to another. For more information, refer to the *Redeploy a Crosswork Data Gateway VM* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

[1] There are licensing implications for different tiers, please consult your Cisco Account team to understand which tiers and licenses are required for your use cases.

Cisco Crosswork Network Controller tier [1]	Contents	Crosswork Data Gateway deployment [2]	Minimum number of VMs needed [3]
Premier	Crosswork Workflow Manager	**On-Premise Extended**: Collectors and offload services.	**3 Hybrid nodes + 2 Worker nodes**
	Crosswork Workflow Manager Solutions	**On-Premise Extended**: Collectors and offload services.	
Add-on	Crosswork Change Automation	**On-Premise Extended**: Collectors and offload services.	**3 Hybrid nodes + 2 Worker nodes**
	Crosswork Health Insights	**On-Premise Extended**: Collectors and offload services.	

Ensure that you have sufficient worker nodes in your cluster. You can always check the load in your cluster and choose to add new worker nodes post installation. For more information, refer to the *Deploy New Cluster Nodes* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Identify the resource footprint

Identifies per-VM resource requirements and links to platform-specific and physical server specifications.

Once you determine the Crosswork applications you want and the number of VMs you will need to deploy to host them, ensure that you have the resources needed for them. The resources required per VM such as CPU, Memory, and Storage vary based on the data center where your VMs will be hosted.

[3] The number of VMs mentioned is the minimum count. You can add more Worker nodes (maximum up to three worker nodes) as needed. If your requirements exceed the minimum count, please contact the Cisco Customer Experience team.

[2] The VM resource requirements for Crosswork Data Gateway are different for each type and cannot be modified. Therefore, if your requirements change, you must re-deploy the Crosswork Data Gateway to move from one type to another. For more information, refer to the *Redeploy a Crosswork Data Gateway VM* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

[1] There are licensing implications for different tiers, please consult your Cisco Account team to understand which tiers and licenses are required for your use cases.

Note

- The resources listed for NSO are higher than for other NSO use cases due to the additional requirements Crosswork Network Controller places on NSO.
- The NSO footprint depends on the type of deployment, standalone or LSA.
- The SR-PCE count will depend on the number of head-ends that need to be managed
- The values in **Storage** column is the space needed for storing Crosswork files and does not consider any additional overhead that may be required (for example, RAID configuration). Storage requirements vary based on factors such as the number of devices being supported and the type of deployment selected. However, 1 TB disk space should work for most deployments.
- The storage required for each backup will vary based on the your cluster size, applications in the cluster, and the scale requirements.
- Solid-state drives (SSDs) are required for all Crosswork Network Controller deployments due to their performance.
- Upgrade of the cluster (using parallel hardware) temporarily requires double the total disk space used by the cluster. refer to Upgrade using parallel hardware for more information.
- The number of Data Gateways needed depends on the number of devices you have in your network and the level of redundancy you want (1:n up to 1:1). To determine the number of Crosswork Data Gateways needed, contact the Cisco Customer Experience team.

refer to these topics to know more about resource requirements per VM for each deployment:

- VM requirements on page 21
- Resource footprint for VMware on page 24
- Resource footprint for AWS EC2 on page 25
- Resource footprint for KVM on page 26
- Resource footprint for single VM deployments on page 28
- Physical server requirements on page 29

VM requirements

Lists storage, disk, IOPS, and network bandwidth requirements for Hybrid and Worker VMs.

This topic explains the requirements for Hybrid and Worker VMs when deploying different tiers of the Crosswork Network Controller.

VM storage specifications

Table 2: Storage specifications per VM

Crosswork Network Controller tier	VM type	Device count	Recommended VM disk size (approximated)	File system to configure	Required file system size
Essentials	Hybrid	<= 15k	~ 1 TB	ManagerDataFs	500 GB
	Hybrid	> 15k	~ 2 TB	ManagerDataFs	1500 GB
	Worker	<= 15k	~ 1 TB	WorkerDataFs	500 GB
	Worker	> 15k	~ 2 TB	WorkerDataFs	1500 GB
Advantage / Premier	Hybrid	N/A	~ 1 TB	ManagerDataFs	500 GB
	Worker	<= 15k	~ 1 TB	WorkerDataFs	500 GB
	Worker	> 15k	~ 2 TB	WorkerDataFs	1500 GB

Here are the disk parameter values for each use case:

> **Note**
>
> It is recommended to use the default value unless otherwise specified.

Table 3: Disk parameter values

Disk parameter	1 TB VM disk storage	2 TB VM disk storage
Datafs disk size (GB)	485 GB (default) - 8000 GB	1500 GB (default) - 8000 GB
Logfs disk size (GB)	20 GB (default) - 1000 GB	
Corefs partition size (GB)	18 GB (default) - 1000 GB	
High speed disk size (GB)	15 GB (default)	
Minimum backup partition size (percentage of the data disk space)	1% - 80% Default value: 35%	

Table 4: Data file systems

Parameter	Description
`ManagerDataFsSize`	Size of data file system for Hybrid VMs.
`WorkerDataFsSize`	Size of data file system for Worker VMs.

Disk requirements

The VM data store(s) need to have disk access latency < 10 ms and > 4000 IOPs. For more information on the diagnostic checks, refer to Diagnostic assessment.

Table 5: IOPs requirements

IOPs value	Status	Response
IOPs value > 4000	PASS	Installation proceeds without the need for user acknowledgement.
1000 < IOPs value > 4000	NOT IDEAL	For production deployments, the installation requires ~ 4000 IOPs for optimal scaling performance. If the system operates in the suboptimal range of 1000 to 4000 IOPs, the installation process requires user acknowledgment before continuing. If the installation is blocked, you can choose to ignore the failure and continue. You must explicitly acknowledge the report before proceeding, and by doing so, you acknowledge this warning and accept the risks.
IOPs value < 1000	FAIL	Installation is blocked irrespective of user response.

Note

- When installing on shared network storage or datastores used by multiple Crosswork virtual machines (VMs), power on Crosswork nodes sequentially. Start with the seed node first, then power on each additional node one at a time. This sequential approach ensures accurate IOPs diagnostics across all Crosswork nodes.
- If you power on nodes simultaneously in the shared network storage or datastore environments, initial diagnostics might not accurately reflect IOPs availability, which can lead to incorrect system assessments.
- The interval time between powering on VMs during the initial installation depends on how long diagnostics take to complete. The **diagnostic all** command typically completes in 4–5 minutes in VMware environments and in 14–15 minutes in KVM environments.

Table 6: Network bandwidth requirements

Network bandwidth	Status	Response
Bandwidth > 8000 Mbps	PASS	Installation proceeds without the need for user acknowledgement.

Network bandwidth	Status	Response
1000 < Bandwidth < 8000 Mbps	NOT IDEAL	Values between 1000 and 8000 Mbps are considered suboptimal. Installation is blocked, but you can choose to ignore the failure and continue. You must explicitly acknowledge the report before proceeding, and by doing so, you acknowledge this warning and accept the risks.
Bandwidth < 1000 Mbps	FAIL	Installation is blocked irrespective of user response.

Resource footprint for VMware

Lists per-VM resource requirements for Crosswork Network Controller components deployed on VMware.

This topic explains the resource requirements per VM to deploy Crosswork Network Controller Hybrid or Worker node, Crosswork Data Gateways, NSO, and SR-PCE on VMware.

Note

- As Crosswork Network Controller cluster nodes place high demands on the VMs, ensure that you have not oversubscribed CPU or memory resources on the machines hosting the nodes.
- Ensure that you have a docker-capable host to load the installer tool.
- The VM data store(s) must have disk access latency < 10 ms and > 4000 IOPS.
- The network bandwidth from Node IP towards to Management VIP should be at least 8000 Mbps.

Table 7: Resource footprint for VMware

Component	vCPU	Clock Freq (GHz)	Memory (RAM)	Network Interface Controller (NIC)	Storage (Boot disk + Data disks)
Crosswork Network Controller Hybrid or Worker node	12 Minimum clock reservation: 18 GHz	>= 2.20	96 GB	10 Gbps	refer to VM requirements on page 21
Crosswork Data Gateway On-Premise Standard	12	>= 2.20	48 GB	10 Gbps	70 GB (50 GB + 20 GB)

Component	vCPU	Clock Freq (GHz)	Memory (RAM)	Network Interface Controller (NIC)	Storage (Boot disk + Data disks)
Crosswork Data Gateway On-Premise Extended	20	>= 2.20	96 GB	10 Gbps	570 GB (50 GB + 520 GB)
Cisco NSO	16	>= 2.20	256 GB	10 Gbps	1 TB
Cisco SR-PCE	8	>= 2.20	64 GB	10 Gbps	128 GB
Note The hardware specifications shown in this row are recommended by Cisco XRv9K platform for SR-PCE deployments on Cisco IOS XR release 25.3.x or later.					
Basic SCP Server (for storing backups)	-	-	-	-	At least 25 GB (recommended)

Resource footprint for AWS EC2

Lists per-VM resource requirements for Crosswork Network Controller components deployed on AWS EC2.

This topic explains the resource requirements per VM to deploy Crosswork Network Controller Hybrid or Worker node, Crosswork Data Gateways, NSO, and SR-PCE on AWS EC2.

When deploying in AWS EC2, choose an EC2 instance size that meets or surpasses these resource requirements listed:

> **Note**
>
> The VM data store(s) must have disk access latency < 10 ms and > 4000 IOPS.

Table 8: Resource footprint for AWS EC2

Component	vCPU	Memory (RAM)	Network Interface Controller (NIC)	Storage (Boot disk + Data Disks)
Crosswork Network Controller Hybrid or Worker node	12	96 GB	10 Gbps	refer to VM requirements on page 21
Crosswork Data Gateway On-Premise Standard	12	64 GB	10 Gbps	70 GB (50 GB + 20GB)
Crosswork Data Gateway On-Premise Extended	24	128 GB	10 Gbps	570 GB (50 GB + 520 GB)
Cisco NSO	16	256 GB	10 Gbps	1 TB
Cisco SR-PCE	8	64 GB	10 Gbps	128 GB
Note The hardware specifications shown in this row are recommended by Cisco XRv9K platform for SR-PCE deployments on Cisco IOS XR release 25.3.x or later.				
Basic SCP Server (for storing backups)	-	-	-	At least 25 GB (recommended)

Resource footprint for KVM

Lists per-VM resource requirements for Crosswork Network Controller components deployed on KVM.

This topic explains the resource requirements per VM to deploy Crosswork Network Controller Hybrid or Worker nodes, Crosswork Data Gateways, and NSO on KVM based hypervisor.

Note

- As Crosswork Network Controller cluster nodes place high demands on the VMs, ensure that you have not oversubscribed CPU or memory resources on the machines hosting the nodes.
- The VM data store(s) must have disk access latency < 10 ms and > 4000 IOPS.
- The network bandwidth from Node IP towards to Management VIP should be at least 8000 Mbps.

Table 9: Resource footprint for KVM

Component	vCPU	Clock Freq (GHz)	Memory (RAM)	Network Interface Controller (NIC)	Storage (Boot disk + Data disks)
Crosswork Network Controller Hybrid or Worker node	12 Minimum clock reservation: 18 GHz	>= 2.20	96 GB	10 Gbps	refer to VM requirements on page 21
Crosswork Data Gateway On-Premise Standard	12	>= 2.20	48 GB	10 Gbps	70 GB (50 GB + 20 GB)
Crosswork Data Gateway On-Premise Extended	20	>= 2.20	96 GB	10 Gbps	570 GB (50 GB + 520 GB)
Cisco NSO	16	>= 2.20	256 GB	10 Gbps	1 TB

Component	vCPU	Clock Freq (GHz)	Memory (RAM)	Network Interface Controller (NIC)	Storage (Boot disk + Data disks)
Cisco SR-PCE	8	>= 2.20	64 GB	10 Gbps	128 GB
Note The hardware specifications shown in this row are recommended by Cisco XRv9K platform for SR-PCE deployments on Cisco IOS XR release 25.3.x or later.					
Basic SCP Server (for storing backups)	-	-	-	-	At least 25 GB (recommended)

Resource footprint for single VM deployments

Lists recommended resources and datastore requirements for a single VM deployment.

This table lists the recommended resources to deploy Crosswork Network Controller on a single VM.

Note

- As Crosswork Network Controller node places high demands on the VM, ensure that you have not oversubscribed CPU or memory resources on the machine hosting the node.
- The VM data store need to have disk access latency < 10 ms and > 4000 IOPS.
- The network bandwidth from Node IP towards to Management VIP should be at least 8000 Mbps.

Table 10: Crosswork Network Controller on a single VM

VM profile	vCPU	Memory (RAM)	Storage	Use case
`XLarge`	24	128 GB	1 TB	The`XLarge` VM profile supports deployment of all tiers (Essential, Advantage, and Premier) of Crosswork Network Controller on a single VM.
`Large`	12	96 GB	1 TB	The`Large` VM profile only supports deployment of Crosswork Network Controller Essentials tier on a single VM.
`Small`	8	48 GB	650 GB	The`Small` VM profile is only used to deploy the arbiter VM on a single VM (to enable auto-arbitration in a geo HA setup).

Physical server requirements

Describes physical server requirements and resource calculations for Crosswork Network Controller tiers.

This topic details the recommended physical server (host machine) requirements for deploying the various tiers of the Crosswork Network Controller.

Formulas for resource calculation

When sizing your host machine, allocate a 20% buffer for CPU and memory resources, and add 100 GB of storage for each VM that will run on the host machine. This helps ensure smooth performance and reduces the risk of resource-related issues during operation.

Note

Please note that this is a recommendation for optimal performance, not a mandatory requirement.

- **RAM calculation**:

  Total RAM: (Number of VMs) * (RAM per VM)

  Buffer (for example, 20%): Total RAM * 0.20

  Estimated total RAM needed (GB) = Total RAM + Buffer

- **CPU (vCPU) calculation**:

  Total vCPUs: (Number of VMs) * (vCPUs per VM)

  Buffer (for example, 20%): Total vCPUs * 0.20

  Estimated total vCPUs needed = Total vCPUs + Buffer

- **Storage calculation**:

  Total Storage: (Number of VMs) * (Storage per VM)

  Buffer: additional 100 GB per VM

  Estimated total storage needed (GB) = Total Storage + Buffer

Cluster deployments

Cluster deployments support only the`Large` VM profile.

Note

The calculations in this table are performed for a minimum cluster size of 3 VMs and assume all 3 VMs reside on a single host with 1 TB of storage per VM. If you plan to distribute your VMs across multiple hosts (for example, 1 VM per host), you must adjust the calculated host resource requirements accordingly.

Table 11: Physical server requirements for cluster deployments

Component	Minimum requirement per physical server
RAM	Number of VMs: 3 RAM per VM: 96 GB (`Large` profile) Total RAM: 3 * 96 = 288 GB Buffer: 288 * 0.20 = 57.6 GB **Estimated total RAM needed: 288 + 57.6 = 345.6 GB**

Component	Minimum requirement per physical server
CPU	Number of VMs: 3 vCPUs per VM: 12 (`Large` profile) Total vCPUs needed for VMs: 3 * 12 = 36 Buffer: 36 * 0.20 = 7.2 **Estimated total vCPUs needed: 36 + 7.2 = 43.2**
Storage	Number of VMs: 3 Storage per VM: 1000 GB Total storage needed for VMs: 1000 * 3 = 3000 GB Buffer: 3 * 100 = 300 GB (additional 100 GB per VM) **Estimated total storage needed: 3000 + 300 = 3300 GB**

Single VM deployments

Single VM deployments support`Small` (for arbiter only), `Large`, and `XLarge` VM size profiles to suit different use cases. For more information, see Resource footprint for single VM deployments on page 28.

Table 12: Physical server requirements for single VM deployments

Component	Minimum requirement per physical server		
	`Small` VM profile	**`Large` VM profile**	**`XLarge` VM profile**
RAM	RAM per VM: 48 GB Buffer: 48 * 0.20 = 9.6 GB **Estimated total RAM needed: 48 + 9.6 = 57.6 GB**	RAM per VM: 96 GB Buffer: 96 * 0.20 = 19.2 GB **Estimated total RAM needed: 96 + 19.2 = 115.2 GB**	RAM per VM: 128 GB Buffer: 128 * 0.20 = 25.6 GB **Estimated total RAM needed: 128 + 25.6 = 153.6 GB**
CPU	vCPUs per VM: 8 Buffer: 8 * 0.20 = 1.6 **Estimated total vCPUs needed: 8 + 1.6 = 9.6**	vCPUs per VM: 12 Buffer: 12 * 0.20 = 2.4 **Estimated total vCPUs needed: 12 + 2.4 = 14.4**	vCPUs per VM: 24 Buffer: 24 * 0.20 = 4.8 **Estimated total vCPUs needed: 24 + 4.8 = 28.8**
Storage	Storage per VM: 650 GB Buffer: 100 GB **Estimated total storage needed: 650 + 100 = 750 GB**	Storage per VM: 1000 GB Buffer: 100 GB **Estimated total storage needed: 1000 + 100 = 1100 GB**	Storage per VM: 1000 GB Buffer: 100 GB **Estimated total storage needed: 1000 + 100 = 1100 GB**

Special considerations

Lists setup considerations for software versions, authentication, logging, and playbooks.

- **Check versions of NSO and SR-PCE**- Please refer to the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes* to know the NSO and SR-PCE versions compatible with Crosswork Network Controller.
- **Are you going to use self-signed certificates?**- if yes, you need to make the certificates available. For more information on the type of certificates supported and how to manage them, refer to the *Manage Certificates* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
- **Do you plan to integrate Crosswork with external authentication servers?**- Integration with TACACS+ or other external authentication servers will require you to have credentials created for the Crosswork user accounts and roles.
- **In Crosswork Data Gateway, do you want to configure a Syslog server to collect the syslog?**- if yes, then provide the host name or IPv4 or IPv6 address of an external syslog server. Or, you have the option to configure the syslog server using the interactive console after the installation is complete.

  > **Note**
  >
  > IPv6 deployments allow the use of both full-form and short-form IPv6 addresses during installation. In the backend, all IPv6 addresses are normalized to a standard format regardless of the input form.

- **In Crosswork Data Gateway, do you want to configure an Auditd server to collect the event logs?**- if yes, then provide the host name or IPv4 or IPv6 address of an external auditd server. As an alternative, you have the option to configure the auditd server using the interactive console after the installation is complete.
- **Do you plan to enable the automatic execution of linked playbooks?**- If yes, you must enable Playbook Job Scheduling and disable Credential Prompting for playbook execution in the Network Automation settings window. For more information, refer to the *Enable Automatic Playbook Execution* topic in the *Cisco Crosswork Network Controller 7.2.x Closed-Loop Network Automation* guide.

3 Choose Installation Workflow

Topics:

- Overview
- Install Cisco Crosswork Network Controller on VMware vCenter
- Install Cisco Crosswork Network Controller on KVM
- Install Cisco Crosswork Network Controller on AWS EC2
- Install Cisco Crosswork Network Controller on a single VM

Summarizes the installation paths for VMware vCenter, KVM, AWS EC2, and single VM deployments, and directs readers to the stage tables and linked installation guidance.

This chapter presents installation workflows for VMware vCenter, KVM, AWS EC2, and single VM deployments. Use the tables to review the major installation stages and linked task guidance.

Overview

Summarizes the supported VMware, KVM, AWS EC2, and single VM installation workflows, along with the chapter scope, typical installation timing, and release compatibility guidance.

This chapter outlines the installation workflows for each of the supported deployment options: VMware, KVM, AWS, and Single VM.

The workflows give a high-level description of the tasks necessary to install the Crosswork Network Controller and to integrate it with the required components of the solution. Integration with optional components is addressed in *Cisco Crosswork Network Controller 7.2.x Administration Guide*

These workflow steps should be used as your primary installation guidepost and roadmap for the end to end installation of Crosswork Network Controller. After completing each detailed step, it is recommended that you refer to the workflow chart for the next step to perform.

> **Note**
>
> The entire installation can typically take up to a few hours based on the size of your deployment profile and the performance characteristics of your hardware.

Please refer to the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes* to know the NSO and SR-PCE versions compatible with Crosswork Network Controller.

Install Cisco Crosswork Network Controller on VMware vCenter

Outlines the VMware vCenter installation stages and points to related guidance for preparing the environment, installing the controller and Data Gateway, and completing integrations.

The table describes the stages to install Crosswork Network Controller on VMware vCenter.

Before you begin:

- Ensure you have identified the Crosswork components you need and arranged for the resources required to complete the installation. If not, please refer to the guidelines in Plan Your Deployment on page 17.

Table 13: vCenter installation workflow

Step	Action
Prepare for installation	
1. Ensure that your VMware environment meets all the requirements.	Refer to the guidelines in Installation Prerequisites on page 43 and VMware settings on page 82.
Install the Crosswork Network Controller cluster	

Step	Action
2. Install the Cisco Crosswork cluster on VMware vCenter.	Install using your preferred method: • *Using Docker installer tool:* Install Cisco Crosswork on VMware vCenter using the Docker installer tool on page 94 • *Manual Installation:* Manual installation of Cisco Crosswork using vCenter vSphere UI on page 102
3. Verify if the installation was successful, and log into the Cisco Crosswork UI.	Refer to the guidelines in: • Monitor cluster activation on page 113 • Log into the Cisco Crosswork UI on page 120
Install Crosswork Data Gateway	
4. Install one or more Data Gateways on VMware vCenter.	Choose the profile for the Data Gateway VM (Standard or Extended) and install as per your preferred method: • *Using vSphere:* Install Crosswork Data Gateway using vCenter vSphere Client on page 144 • *Using OVF tool:* Install Crosswork Data Gateway via the OVF Tool on page 154 **Note** If you plan to install multiple Data Gateway VMs due to load or scale requirements, or you wish to leverage Data Gateway High Availability, you are recommended to install all the Data Gateway VMs first and then proceed with adding them to a Data Gateway pool.
5. Verify that the Data Gateway VM or multiple VMs have enrolled successfully with *Cisco Crosswork Network Controller 7.2.x Administration Guide*.	Follow the steps in Crosswork Data Gateway authentication and enrollment on page 158. After verifying that the Data Gateway VM has enrolled successfully with Cisco Crosswork, set up the Data Gateway for collection by creating a Data Gateway pool. Refer to the *Create a Crosswork Data Gateway Pool* section in *Cisco Crosswork Network Controller 7.2.x Administration Guide*
6. Complete the Data Gateway post-installation tasks.	Follow the steps in Crosswork Data Gateway post-installation tasks on page 159.
Install the Crosswork Network Controller Applications	
7. Install the Crosswork Applications	Follow the instructions in Install Crosswork Applications on page 265.
Integrate NSO with Crosswork Network Controller	

Step	Action
8. Do you have Cisco NSO already installed?	If yes, proceed to step 9. If no, please follow the **system install** instructions in the NSO documentation. **Note** The NSO requirements for Crosswork integration exceed those specified in the NSO installation document.
9. Add NSO Provider and verify that it is reachable	Follow the instructions in Integrate Cisco NSO on page 273.
10. Install the latest NSO Function Packs	Follow the instructions in Deploy NSO function pack bundles on page 289.
Integrate SR-PCE with Crosswork Network Controller	
11. Is your SR-PCE installed?	If yes, please proceed to step 12. If no, please choose the type of SR-PCE you wish to use (physical or virtual device) and follow the appropriate instructions to get the device (or virtual device) deployed. For more information, refer to the Cisco IOS XRv 9000 Router Installation Guide for the XR device you are going to use as your PCE. **Note** For the rest of the document, we will refer to the physical or virtual device(s) as the SR-PCE(s).
12. Configure SR-PCE	Follow the instructions in Configure SR-PCE.
13. Add SR-PCE Provider and verify that it is reachable.	Follow the instructions in Add Cisco SR-PCE Providers.
14. (Recommended) Create a backup of your Crosswork Network Controller.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide*

After completing the installation, follow the instructions in the *Setup workflow* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

Install Cisco Crosswork Network Controller on KVM

Outlines the KVM installation stages and links to related guidance for preparing the environment, installing the controller and Data Gateway, and completing integrations.

This table describes the stages to install Crosswork Network Controller on a KVM-based hypervisor.

Kernel-based Virtual Machine (KVM) is a virtualization solution for Linux on x86 hardware. It uses the hardware virtualization extensions of Intel (VT-x) and AMD (AMD-V) processors to allow multiple operating systems to run simultaneously on a host machine.

> **Note**
>
> Ensure you have identified the Crosswork components you need and arranged for the resources required to complete the installation. If not, please refer to the guidelines in Plan Your Deployment on page 17.

Table 14: KVM installation workflow

Step	Action
Prepare for installation	
1. Ensure that your KVM environment meets all the requirements.	Refer to the guidelines in Installation Prerequisites on page 43 and KVM host bare metal requirements on page 170.
Install the Crosswork Network Controller cluster	
2. Install the Crosswork Network Controller cluster on KVM.	Refer to the guidelines in Install Crosswork Cluster on KVM on page 169.
Install Crosswork Data Gateway	
3. Install one or more Data Gateway instances on VMware vCenter.	Refer to the guidelines in Install Crosswork Data Gateway on KVM on page 199 **Note** If you plan to install multiple Data Gateway VMs due to load or scale requirements, or you wish to leverage Data Gateway High Availability, you are recommended to install all the Data Gateway VMs first and then proceed with adding them to a Data Gateway pool.
Install the Crosswork Network Controller Applications	
4. Install the Crosswork Network Controller Applications	Follow the instructions in Install Crosswork Applications on page 265.
Integrate NSO with Crosswork Network Controller	

Step	Action
5. Do you have Cisco NSO already installed?	If yes, proceed to step 6. If no, please follow the **system install** instructions in the NSO documentation. **Note** The NSO requirements for Crosswork integration exceed those specified in the NSO installation document.
6. Add NSO Provider and verify that it is reachable	Follow the instructions in Integrate Cisco NSO on page 273.
7. Install the latest NSO Function Packs	Follow the instructions in Deploy NSO function pack bundles on page 289.
Integrate SR-PCE with Crosswork Network Controller	
8. Is your SR-PCE installed?	If yes, please proceed to step 9. If no, follow the instructions in Installation workflow for deploying SR-PCE (Cisco IOS XRv 9000) on RHEL 8.10 KVM on page 193 **Note** For the rest of the document, we will refer to the physical or virtual device(s) as the SR-PCE(s).
9. Configure SR-PCE	Follow the instructions in Configure SR-PCE.
10. Add SR-PCE Provider and verify that it is reachable.	Follow the instructions in Add Cisco SR-PCE Providers.
11. (Recommended) Create a backup of your Crosswork Network Controller.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide*

After completing the installation, follow the instructions in the *Setup workflow* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

Install Cisco Crosswork Network Controller on AWS EC2

Outlines the AWS EC2 installation stages that use CloudFormation templates and links to related guidance for preparing the environment, installing the controller components, and completing integrations.

The table describes the stages to install Crosswork Network Controller on AWS EC2 using CloudFormation (CF) templates.

Before you begin:

- Ensure you have identified the Crosswork components you need and arranged for the resources required to complete the installation. If not, please refer to the guidelines in Plan Your Deployment on page 17.

Crosswork Network Controller enables you to selectively install components of the Cisco Crosswork solution, including hybrid and worker nodes for the Crosswork cluster, one or more Data Gateways, and NSO, based on your deployment needs.

Table 15: Crosswork installation workflow

Step	Action
Prepare for installation	
1. Ensure that your AWS EC2 environment meets all the requirements.	Refer to the guidelines in Installation Prerequisites on page 43 and Amazon EC2 requirements for Crosswork on page 226.
2. Extract the CF template package	Follow the instructions in Extract CF template image on page 229.
Install Crosswork components	
3. **Module deployment:**Install the Cisco Crosswork components using module deployment.	Install the Crosswork components individually: • Install the Crosswork cluster: Cisco Crosswork cluster deployment on Amazon on page 244 • Install one or more Data Gateways: Crosswork Data Gateway deployment on Amazon on page 245 Crosswork Data Gateway is deployed with the default parameter values if you missed configuring the values during deployment. For more information, refer to Auto-configuration for deploying Crosswork Data Gateway on page 246.
4. Verify the installation and access Crosswork UI	Refer to the guidelines in Access the Crosswork UI on page 253
Install the Crosswork Applications	
5. Install the Crosswork Applications	Follow the instructions in Install Crosswork Applications on page 265.
Integrate NSO with Crosswork	

Step	Action
6. Do you have Cisco NSO already installed?	If yes, proceed to step 7. If no, please follow the **system install** instructions in the NSO documentation. **Note** The NSO requirements for Crosswork integration exceed those specified in the NSO installation document.
7. Install the NSO Function Packs	Follow the instructions in Deploy NSO function pack bundles on page 289.
8. Add NSO Provider and verify that is reachable.	Follow the instructions in Integrate Cisco NSO on page 273.
Integrate SR-PCE with Crosswork	
9. Is your SR-PCE installed?	If yes, please proceed to step 10. If no, please select the SR-PCE type (for AWS) and follow the relevant install instructions in the Cisco IOS XRv 9000 Router Installation Guide.
10. Configure SR-PCE	Follow the instructions in Configure SR-PCE.
11. Add SR-PCE Provider and verify that SR-PCE is reachable.	Follow the instructions in Add Cisco SR-PCE Providers.
12. (Recommended) Create a backup of your Crosswork Network Controller.	Follow the instructions in *Manage Backups* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

After completing the installation, follow the instructions in the *Setup workflow* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

Install Cisco Crosswork Network Controller on a single VM

Lists the stages for installing Cisco Crosswork Network Controller on a single VM and links to the requirements and task guidance for each stage.

The table lists the stages to install Crosswork Network Controller on a single VM or node.

Table 16: Installation workflow

Step	Action
Prepare for installation	
1. Ensure that your environment meets all the requirements for single VM deployment.	Refer to the guidelines in Installation requirements on page 332.

Step	Action
Install Crosswork Network Controller	
2. Install the Crosswork Network Controller on a single VM.	Install using your preferred method: • **VMware vCenter**: • *Manual Installation:* Install Crosswork Network Controller using the vCenter vSphere UI on page 342 • *Using the OVF tool:* Install Crosswork Network Controller via the OVF Tool on page 348 • *Using Docker installer tool:* Install Crosswork Network Controller using the Docker installer tool on page 351 • **KVM**: Install Crosswork Network Controller VM using CLI on page 357
Integrate SR-PCE with Crosswork	
3. Is your SR-PCE installed?	If yes, please proceed to step 4. If no, please choose the type of SR-PCE you wish to use (physical or virtual device) and follow the appropriate instructions to get the device (or virtual device) deployed. For more information, refer to the Cisco IOS XRv 9000 Router Installation Guide for the XR device you are going to use as your PCE. **Note** For the rest of the document, we will refer to the physical or virtual device(s) as the SR-PCE(s).
4. Configure SR-PCE	Follow the instructions in Configure SR-PCE
5. Add SR-PCE Provider and verify that it is reachable.	Follow the instructions in Add Cisco SR-PCE Providers.
6. (Recommended) Create a backup of your Crosswork Network Controller.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide*

Step	Action
(Optional) **Install custom NSO Function Packs**	

> **Note**
>
> This step applies only when deploying Crosswork Network Controller Advantage tier on a single VM. In this case, the standard NSO Function Packs are installed automatically.

(Optional) 7. Install custom NSO Function Packs.	Contact the Cisco Customer Experience team for assistance on bundling the custom Function Pack. For installation instructions, refer to Deploy NSO function pack bundles on page 289.

After completing the installation, follow the instructions in the *Setup workflow* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

4 Installation Prerequisites

Topics:

- Installation requirements overview
- VM requirements
- Validate storage performance and review diagnostics
- Crosswork network ports and communication details
- IP address restrictions
- Supported network topology models

Summarizes the prerequisites covered in this chapter, including VM resources, storage performance, network communication, IP address restrictions, and supported topologies.

This chapter lists the requirements that apply before you install Cisco Crosswork Network Controller and its components. It covers deployment resources, storage performance, network ports, IP addresses, and supported topology models.

Installation requirements overview

Outlines the general and platform-specific prerequisites for installing Crosswork components, including VM, port, application, and deployment requirements for supported environments.

This chapter explains the general (such as VM requirements, port requirements, application requirements, etc.) and platform-specific prerequisites

The data center resources needed to operate other integrated components or applications (such as WAE, DHCP, and TFTP servers) are not addressed in this document. Refer to the respective installation documentation of those components for more details.

VM requirements

Summarizes the per-VM resources required for the Crosswork cluster and Crosswork Data Gateway, and links to detailed host and gateway specifications.

This section explains the resource requirements per VM to deploy the Crosswork Cluster and Crosswork Data Gateway.

- Crosswork cluster VM requirements on page 44
- Crosswork Data Gateway VM requirements on page 46

Crosswork cluster VM requirements

Lists VM host resource, network, SSH password, and IP address requirements for Crosswork cluster deployments, including network interface and addressing details.

The Crosswork cluster consists of three VMs or nodes operating in a hybrid configuration. This is the minimum configuration necessary to support the applications in a standard network. Additional VMs or nodes in a worker configuration can be added during the installation or later to scale your deployment. For more information on VM count for each Crosswork Network Controller package, refer to Table 1: Crosswork Network Controller tiers on page 19. Please consult with the Cisco Customer Experience team for guidance on your deployment to best meet your needs.

The table below explains the network requirements per VM host:

Table 17: Network requirements (per VM)

Requirement	Description
Network Connections	For production deployments, we recommend that you use dual interfaces, one for the Management network and one for the Data network. For optimal performance, the Management and Data networks should use links configured at a minimum of 10 Gbps with a latency of less than 10 milliseconds.

Requirement	Description
IP Addresses	The number of IP addresses required may vary based on the selection of the IP stack (IPv4, IPv6, or dual stack). **When using dual NICs** (one for the Management network and one for the Data network): A management and data IP address (IPv4 or IPv6) for each node being deployed (Hybrid or Worker) and two additional IP addresses to be used as the Virtual IP (VIP) address (one for the Management network and one for the Data network). For example, in the case of a single stack deployment of a cluster with 3 Hybrid VMs and 1 Worker VM using dual NICs, you need 10 IP addresses (5 for the Management network and 5 for the Data network). This number would double in the case of a dual stack deployment. **Note** • The IP addresses must be able to reach the gateway address for the network where Crosswork Data Gateway will be installed, or the installation will fail. • When deploying a IPv6 cluster, the installer needs to run on an IPv6 enabled container/VM. • At this time, your IP allocation is permanent and cannot be changed without re-deployment. For more information, contact the Cisco Customer Experience team.
NTP Server	The IPv4 or IPv6 addresses (based on the selection of the IP stack) or host names of the NTP server you plan to use. If you want to enter multiple NTP servers, separate them with spaces. These should be the same NTP servers you use to synchronize the Crosswork application VM clock, devices, clients, and servers across your network. Ensure that the NTP servers are reachable on the network before attempting installation. The installation will fail if the servers cannot be reached.
DNS Servers	The IPv4 or IPv6 addresses ((based on the selection of the IP stack) of the DNS servers you plan to use. These should be the same DNS servers you use to resolve host names across your network. Ensure that the DNS servers are reachable on the network before attempting installation. The installation will fail if the servers cannot be reached.
DNS Search Domain	The search domain you want to use with the DNS servers, for example, cisco.com. You can have only one search domain.

Requirement	Description
Backup Server	Cisco Crosswork will back up the configuration of the system to an external server using SCP. The SCP server storage requirements will vary slightly but you must have at least 25 GB of storage.
FQDN (Optional)	The installation process supports using either a VIP (Virtual IP address) or a FQDN (Fully Qualified Domain Name) to access the cluster. If you choose to use the FQDN, you will need one for the Management and one for the Data network. **Note** Secure ZTP and Secure Syslog require the Crosswork cluster to be deployed with FQDN.

Cisco Crosswork Infrastructure and applications are built to run as a distributed collection of containers managed by Kubernetes.

SSH password guidelines

- When updating the SSH password for the cluster, change the password on all cluster nodes simultaneously, ensuring each node uses the same password.
- In a geo setup, after updating the SSH password on the cluster nodes, promptly update the password in the geo inventory.

Crosswork Data Gateway VM requirements

Lists Crosswork Data Gateway deployment types and the VM requirements for supported on-premises Crosswork products, including interface and resource specifications.

This section provides information about the general guidelines and minimum requirements for installing Crosswork Data Gateway.

- Crosswork Data Gateway deployment type on page 46
- VM requirements on page 47

Crosswork Data Gateway deployment type

this table lists the deployment profile that must be used for installing Crosswork Data Gateway in each Crosswork product:

Table 18: Crosswork Data Gateway deployment types

Cisco Crosswork Network Controller package [4]	Contents	Crosswork Data Gateway Deployment [5]
Essentials	Element Management Functions	**On-Premise Standard** (default): Collectors only.
Advantage	Cisco Crosswork Optimization Engine	**On-Premise Standard** (default): Collectors only.
	Cisco Crosswork Active Topology	**On-Premise Standard** (default): Collectors only.
	Cisco Crosswork Service Health	**On-Premise Extended**: Collectors and offload services.
Add-on	Cisco Crosswork Change Automation	**On-Premise Extended**: Collectors and offload services.
	Cisco Crosswork Health Insights	**On-Premise Extended**: Collectors and offload services.

VM requirements

The VM requirements for Crosswork Data Gateway are listed in this table.

[5] The VM resource requirements for Crosswork Data Gateway are different for each type and cannot be modified. Therefore, if your requirements change, you must redeploy the Crosswork Data Gateway to move from one type to another. For more information, refer to the *Redeploy a Crosswork Data Gateway VM* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

[4] There are licensing implications for different packages, please consult your Cisco Account team to understand which packages and licenses are required for your use cases.

Table 19: Crosswork Data Gateway Requirements for on-premise applications

<table>
<tr><th>Requirement</th><th>Description</th></tr>
<tr><td>Interfaces</td><td>Minimum: 2
Maximum: 3
Crosswork Data Gateway can be deployed with 2 or 3 interfaces as per the combinations after this:
Note
If you use two interfaces on your Crosswork Cluster, then you can use two or three interfaces on the Data Gateway as per your network requirements.
<table><tr><th>No. of NICs</th><th>vNIC0</th><th>vNIC1</th><th>vNIC2</th></tr><tr><td>2</td><td>Management Traffic</td><td>• Control/Data Traffic
• Device Access Traffic</td><td>–</td></tr><tr><td>3</td><td>Management Traffic</td><td>Control/Data Traffic</td><td>Device Access Traffic</td></tr></table>• Management traffic: for accessing the Interactive Console and passing the Control/Data information between servers (for example, a Crosswork application to Crosswork Data Gateway).
• Control/Data traffic: for data and configuration transfer between Crosswork Data Gateway and Crosswork applications and other external data destinations.
• Device access traffic: for device access and data collection.
Note
Due to security policies, traffic from subnets of a vNIC received on other vNICs is dropped. For example, in a 3 vNIC model setup, all device traffic (incoming and outgoing) must be routed through default vNIC2. Crosswork Data Gateway drops device traffic received over vNIC0 and vNIC1.</td></tr>
<tr><td>IP Addresses</td><td></td></tr>
</table>

Requirement	Description
	1 or 2 IPv4 or IPv6 addresses based on the number of interfaces that you choose to use. An additional IP address to be used as the Virtual IP (VIP) address. For each active Data Gateway, a unique VIP is required. For more information, refer to the *Interfaces* section in the Crosswork Data Gateway Deployment Parameters and Scenarios. **Note** In a 3-NIC deployment, you must provide an IP address for Management interface (vNIC0) and Control/Data interface (vNIC1) during installation. A virtual IP address for Device Access Traffic (vNIC2) is assigned when you create a Data Gateway to a pool as explained in the *Create a Crosswork Data Gateway Pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
NTP Servers	The IPv4 or IPv6 addresses or hostnames of the NTP servers you plan to use. If you want to enter multiple NTP servers, separate them with spaces. These should be the same NTP servers you use to synchronize devices, clients, and servers across your network. Verify that the NTP IP address or host name is reachable on the network else the installation fails. Also, the ESXi hosts that run the Crosswork application and Crosswork Data Gateway VM must have NTP configured, or the initial handshake may fail with "certificate not valid" errors.
DNS Servers	The IPv4 or IPv6 addresses of the DNS servers you plan to use. These should be the same DNS servers you use to resolve hostnames across your network. Confirm that the DNS servers are reachable on the network before attempting installation. The installation fails if the servers cannot be reached.
DNS Search Domain	The search domain you want to use with the DNS servers, for example cisco.com. You can have only one search domain.
FQDN	The FQDN addresses are configured for Amazon EC2 deployments.
Internet Control Message Protocol (ICMP)	The Crosswork uses ICMP in the communications with Crosswork Data Gateway. Ensure that the firewall between Crosswork and Crosswork Data Gateway passes this traffic.

Validate storage performance and review diagnostics

Use this topic for VMware and KVM deployments to validate storage performance before installation, review diagnostic results during installation, and understand the risk of bypassing failed diagnostic checks.

Refer to this topic before you install Crosswork Network Controller on VMware or KVM, whether you are deploying a cluster or a single virtual machine, so that you can verify storage readiness, interpret diagnostic results, and decide whether installation can proceed.

When to perform validation

- Validate storage performance after you complete the hypervisor and storage setup and before you start the Crosswork Network Controller installation.

- Review the diagnostic results again if the installation reports a diagnostic failure.
- For VMware deployments, perform this validation after the datastore and virtual machine environment are prepared and before you continue with installation.
- For KVM deployments, perform this validation after KVM host setup is complete and before you continue with installation.
- When you use shared network storage or datastores that are used by multiple Crosswork virtual machines, power on Crosswork nodes sequentially so that the IOPs diagnostics reflect the actual available storage performance.

Storage requirements and diagnostic outcomes

The VM datastore must provide disk access latency of less than 10 milliseconds and more than 4000 IOPs for recommended operation.

Table 20: Storage diagnostic thresholds

Storage result	Status	Installation response
More than 4000 IOPs	PASS	Installation proceeds without user acknowledgment.
1000 to 4000 IOPs	NOT IDEAL	Installation requires user acknowledgment before continuing. This range is suboptimal for production deployments and can affect scaling and runtime performance.
Less than 1000 IOPs	FAIL	Installation is blocked irrespective of user response.

Warning

For production deployments, do not treat a NOT IDEAL result as equivalent to a passed diagnostic check. Suboptimal storage performance can lead to later deployment failures, cluster instability, slow performance, or service issues.

Run the storage validation command

Use this backend check to verify directory IOPs for the Crosswork storage locations in on-premises deployments:

```
fio --randrepeat=1 --fdatasync=1 --ioengine=sync --name=test --rw=rw \
--filename=<DIR>/mytest --bs=8k --size=600M --runtime=10 --time_based=1
```

Replace `<DIR>` with the applicable directory for the storage that you want to validate.

Table 21: Directories used for the storage validation command

Directory placeholder	Use
`/mnt/cw_ssd/`	Use when validating the high-speed storage location.
`/mnt/datafs`	Use when validating the data filesystem location.

Directory placeholder	Use
`/mnt/cw_glusterfs`	Use when validating the applicable clustered storage location.

Run the validation for the storage path that maps to your deployment. If the measured result is below the recommended threshold, correct the underlying storage issue before you continue with installation. Keep any platform-specific user interface actions, such as VMware vCenter parameter updates, in the corresponding deployment chapter.

Review diagnostic results during installation

If the installer reports a diagnostic failure, use these commands and files to review the detailed results:

Table 22: Diagnostic commands and file locations

Command or file	Purpose
diagnostic all	Runs the diagnostic check manually.
diagnostic history	Displays the previous diagnostic report.
`/var/log/diagnostic_stdout.log`	Stores the standard output from the diagnostic run.
`/var/log/diagnostic_stderr.log`	Stores the standard error output from the diagnostic run.
`/home/cw-admin/diagnosis_report.txt`	Stores the result from each diagnostic execution.

Understand the risk of IgnoreDiagnosticsCheckFailure

The `IgnoreDiagnosticsCheckFailure` parameter controls whether installation stops when the diagnostic check reports an error.

Table 23: IgnoreDiagnosticsCheckFailure behavior

Parameter value	Behavior
`"False"`	The installation terminates if the diagnostic check reports an error. This is the default value and the recommended setting for production deployments.
`"True"`	The installation ignores the diagnostic failure and continues.

Warning

Setting `IgnoreDiagnosticsCheckFailure` to `"True"` can hide insufficient IOPs, high disk latency, or other host and storage issues. Use this setting only when you understand the reported failure and have explicitly accepted the operational risk.

If installation fails while `IgnoreDiagnosticsCheckFailure` is set to `"False"`, review the diagnostic report and correct the infrastructure issue before retrying the installation. For production deployments, do not use `"True"` as a substitute for meeting the documented storage requirements.

Crosswork network ports and communication details

Lists ports and communication details for internal cluster traffic, northbound and southbound services, Crosswork Data Gateway, single VM deployments, and geo redundancy.

This section provides details of the network ports used by the Crosswork Network Controller. To ensure the successful installation and operation of Crosswork Network Controller, allow these ports through your firewalls and security groups.

Warning

- The list of ports in this section applies only to on-premises deployments where all the intra-site components are able to communicate without any firewall restrictions. Firewalls between the individual VMs hosting Crosswork components is not supported. For any inter-site communication between the Crosswork geo redundancy components (if configured) or with any other external tools in the integration, firewall may be present.
- All IP addresses (including Virtual IP addresses) between Crosswork cluster, Crosswork applications, and Crosswork Data Gateway need to be reachable (to be pinged to/from) between each other.
- As a general policy, ports that are not needed should be disabled. To view a list of all the open listening ports once all the applications are installed and active, log in as a Linux CLI admin user on any Crosswork cluster VM, and run the **netstat -aln** command.

- Internal cluster infrastructure on page 52
- Northbound and application services (NBI) on page 56
- Southbound interface (SBI) and external management on page 57
- Crosswork Data Gateway ports on page 62
- Single VM communication on page 64
- Geo redundancy communication on page 68

Internal cluster infrastructure

These ports must be open between all nodes within the Crosswork Network Controller cluster to allow for cluster formation, data synchronization, and health monitoring.

Note

Crosswork cluster ports allow bidirectional flow of information.

Table 24: Internal cluster infrastructure ports

Port	Protocol	Direction	Communication flow	Purpose	Network interface
7	TCP / UDP	Inbound	Cluster node ↔ Cluster node	Used for discovering endpoints and verifying reachability using ICMP.	Management (vNIC0)
22	TCP	Inbound	Admin workstation → Cluster nodes	Used for initial node joining via kubeadm, Terraform installer validation, and administrative SSH access.	Management (vNIC0)
179	TCP	Inbound	Cluster node ↔ Cluster node	Used by the Calico BGP to manage networking between the Kubernetes cluster nodes.	Management (vNIC0)
279	TCP	Inbound	Cluster node ↔ Cluster node	Used by Calico BGP instead of port 179 only if the "UseNonDefaultCalicoBgpPort" setting is enabled.	Management (vNIC0)
500	UDP	Inbound	Cluster node ↔ Cluster node	Used for the IPSec control protocol to manage secure, encrypted communication between cluster nodes.	Management (vNIC0)
4500	UDP	Inbound	Cluster node ↔ Cluster node	Used for the IPSec control protocol (NAT traversal) to manage secure communication between nodes.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
2379	TCP	Inbound	Cluster node ↔ Cluster node	Used for Kubernetes etcd database synchronization to maintain cluster configuration and state.	Management (vNIC0)
2380	TCP	Inbound	Cluster node ↔ Cluster node	Used for Kubernetes etcd database peer communication during cluster state synchronization.	Management (vNIC0)
6443	TCP	Inbound	Cluster node ↔ Cluster node	Used for the Kubernetes API server (kube-apiserver) to manage and orchestrate cluster operations.	Management (vNIC0)
9100	TCP	Inbound	Cluster node ↔ Cluster node	Used for Kubernetes metamonitoring to ensure node health and performance tracking.	Management (vNIC0)
10250	TCP	Inbound	Cluster node ↔ Cluster node	Used by the kubelet service to manage and monitor the state of pods on each individual node.	Management (vNIC0)
24007	TCP	Inbound	Cluster node ↔ Cluster node	Used for GlusterFS control communication to manage distributed storage across the cluster.	Management (vNIC0)
30606	TCP	Inbound	Cluster node ↔ Cluster node	Used for the internal Docker Registry to manage and distribute container images within the cluster.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
31560	TCP	Inbound	Cluster node ↔ Cluster node	Used for iPerf validation of network bandwidth during node joining and system diagnostics.	Management (vNIC0)
49152: 49570	TCP	Inbound	Cluster node ↔ Cluster node	Dedicated to GlusterFS data traffic for storage synchronization across nodes.	Management (vNIC0)
N/A	ESP (50)	Inbound	Cluster node ↔ Cluster node	The ESP protocol provides the encryption for pod-to-pod traffic within the cluster.	Management (vNIC0)
N/A	IP-in-IP (Protocol 4)	Inbound	Cluster node ↔ Cluster node	This protocol handles encapsulated pod-to-pod traffic in IPv4 deployments during initial cluster formation before IPSec is activated.	Management (vNIC0)
N/A	IPv6 ULA	Inbound	Cluster node ↔ Cluster node	This traffic handles pod-to-pod traffic using the IPv6 ULA prefix (fdfb:85ef:26ff::/48) in IPv6 deployments before IPSec is enabled.	Management (vNIC0)
N/A	IGMP (2)	Inbound	Cluster node ↔ Cluster node	IGMP Membership Reports required for joining the VRRP multicast group 224.0.0.18.	Management (vNIC0) / Data (vNIC1)
N/A	AH (51) / VRRP (112)	Inbound	Cluster node ↔ Cluster node	Used for VIP control protocol.	Management (vNIC0) / Data (vNIC1)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
N/A	ICMP/v6	Inbound	Any → Cluster nodes	This protocol is essential for network health checks, path MTU discovery, and IPv6 neighbor discovery.	Management (vNIC0) / Data (vNIC1)

Northbound and application services (NBI)

These ports allow users to access the management interfaces and allow external systems to communicate with Crosswork Network Controller.

Table 25: Northbound and application services (NBI)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30160:31560	TCP / UDP	Inbound	External → Cluster VIP **Note** "External" in this context refers to the customer environment. This includes any external system, browser, or management tool within the customer's network that needs to access the Crossork platform services.	Generic range of TCP/UDP ports enabled in Kubernetes; specific ports within this range are used by various Crosswork Network Controller applications.	Management (vNIC0) / Data (vNIC1)
30602	TCP	Inbound	User browser → Cluster VIP	This port provides access to the Grafana interface for monitoring the status of the CNC installation.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30603	TCP	Inbound	User browser → Cluster VIP	This port provides access to the primary Crosswork Network Controller GUI.	Management (vNIC0) / Data (vNIC1)

Southbound interface (SBI) and external management

These ports are used by the Crosswork Network Controller cluster to communicate with Cisco NSO, SR-PCE, and standard enterprise services.

Table 26: Southbound interface (SBI) and external management

Port	Protocol	Direction	Communication flow	Purpose	Network interface
22	TCP	Outbound	Cluster → Managed device	Used to initiate SSH connections with managed devices for configuration, management, and zero touch provisioning.	Management (vNIC0) / Data (vNIC1)
22	TCP	Outbound	Cluster → Backup server	Used by the backup and restore service to reach the user-defined external storage destination.	Management (vNIC0) / Data (vNIC1)
22	TCP	Outbound	Cluster → User provided IP	Used by the Application Manager to download Crosswork Network Controller application bundles and security patches via SCP.	Management (vNIC0) / Data (vNIC1)
49	TCP	Outbound	Cluster → TACACS+	Used for TACACS+ authentication and authorization for user login and role-based access control.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
53	TCP/ UDP	Outbound	Cluster → DNS	Used for connection to DNS servers for system-wide name resolution.	Management (vNIC0)
80	TCP	Outbound	Cluster → User provided IP	Used by the Application Manager to download Crosswork Network Controller application bundles and security patches via HTTP.	Management (vNIC0) / Data (vNIC1)
123	UDP	Outbound	Cluster → NTP	Used for Network Time Protocol (NTP) synchronization across the cluster.	Management (vNIC0)
389	TCP	Outbound	Cluster → LDAP	Used for non-secure LDAP authentication for user login and role-based access control.	Management (vNIC0)
443	TCP	Outbound	Cluster → User provided IP	Used by the Application Manager to download Crosswork Network Controller application bundles and security patches via HTTP.	Management (vNIC0) / Data (vNIC1)
443	TCP	Outbound	Cluster → vCenter/EC2	Used to connect to the vCenter REST API for VM orchestration. Used for outgoing communication to access the EC2 API.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
443	HTTPS	Outbound	Nesla → Cloud	Used to connect to the Cisco Cloud transport through the nesla agent.	Management (vNIC0)
636	TCP	Outbound	Cluster → LDAP	Used for secure LDAP (SSL/TLS) authentication for user login and role-based access control.	Management (vNIC0)
830	TCP	Outbound	Cluster → NSO	Used for NETCONF over SSH communication between the Crosswork platform and Cisco NSO.	Management (vNIC0)
1645	UDP	Outbound	Cluster → RADIUS	Used for legacy RADIUS authentication and authorization for user login and role-based access control.	Management (vNIC0)
1812	UDP	Outbound	Cluster → RADIUS	Used for RFC standard RADIUS authentication and authorization for user login and RBAC.	Management (vNIC0)
2022	TCP	Outbound	Cluster → NSO	Used for NETCONF over SSH communication specifically for NSO integration.	Management (vNIC0)
8080	TCP	Outbound	Cluster → SR-PCE	Used for REST API communication between the Crosswork cluster and SR-PCE.	Management (vNIC0) / Data (vNIC1)
8443	HTTPS	Outbound	Nesla → On-prem	Used to connect to the on-premise transport through the nesla agent.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
8888	TCP	Outbound	Cluster → NSO	Used for HTTPS-based API communication between Crosswork and Cisco NSO.	Management (vNIC0)
20243	TCP	Outbound	Cluster → NSO	Used by the DLM Function Pack for communication between DLM and Cisco NSO.	Management (vNIC0)
20244	TCP	Outbound	Cluster → NSO	Used to internally manage the DLM Function Pack listener during NSO package reload scenarios.	Management (vNIC0)
57400	TCP	Outbound	Cluster → SR-PCE	Used for gRPC-based communication between the Crosswork cluster and SR-PCE.	Management (vNIC0) / Data (vNIC1)
30604	TCP	Inbound	Device → Cluster VIP	Used for Classic Zero Touch Provisioning (Classic ZTP) on the NGINX server.	Management (vNIC0)
30609	TCP	Outbound	Data Gateway VM → Cluster VIP	Used by the Expression Orchestrator (Service Health).	Management (vNIC0)
30610	TCP	Outbound	Data Gateway VM → Cluster VIP	Used by the Metric Scheduler (Service Health).	Management (vNIC0)
30611	TCP	Outbound	Data Gateway VM → Cluster VIP	Used by the Expression Tracker component (Service Health).	Management (vNIC0)
30612	TCP	Outbound	Data Gateway VM → Cluster VIP	Used by the robot-pulse service for health monitoring.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30617	TCP	Inbound	Device → Cluster VIP	Used for Secure Zero Touch Provisioning (Secure ZTP) on the ZTP server.	Management (vNIC0)
30620	TCP	Inbound	Device → Cluster VIP	Used to receive plug-and-play HTTP traffic on the ZTP server.	Management (vNIC0)
30621 31121 31122 31123	TCP	Inbound	Managed device → Cluster VIP	Used for FTP services (available only when the Element Management Functions is installed and FTP settings are enabled). 31121:31123 are the additional ports used for FTP file transfer.	Data (vNIC1)
30622	TCP	Inbound	Managed device → Cluster VIP	Used for SFTP services (available only when the Element Management Function is installed and SFTP settings are enabled).	Data (vNIC1)
30649	TCP	Inbound	Data Gateway VM → Cluster VIP	Used to set up and monitor Data Gateway collection status (cluster perspective).	Data (vNIC1)
30650	TCP	Outbound	Data Gateway VM → Cluster VIP	Used for the astack gRPC channel with astack-client running on the Data Gateway VMs.	Data (vNIC1)

Crosswork Data Gateway ports

These ports are required for the Crosswork Data Gateway to communicate with the Crosswork Network Controller Cluster and collect data from devices.

Table 27: Management traffic ports

Port	Protocol	Direction	Communication flow	Purpose	Network interface
22	TCP	Inbound	Admin → Data Gateway VM	Used for SSH access for Data Gateway administration and maintenance.	Management
22	TCP	Outbound	Data Gateway VM → Remote file server or storage location	Used by the SCP client for secure file transfers from the Data Gateway.	Management
53	UDP	Outbound	Data Gateway VM → DNS	Used by the DNS client for name resolution on the Data Gateway.	Management
123	UDP	Outbound	Data Gateway VM → NTP	Used by the NTP client for time synchronization on the Data Gateway.	Management
30607	TCP	Outbound	Data Gateway VM → Cluster VIP	Used for the collection of Data Gateway vitals to monitor its operational health.	Management
30608	TCP	Outbound	Data Gateway VM → Cluster VIP	Used for dg-manager gRPC server communication between the Data Gateway and Crosswork Network Controller cluster.	Management

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30649	TCP	Outbound	Data Gateway VM → Cluster VIP	Used for magellan gRPC server communication for data collection control. Used for the magellan gRPC Server to facilitate control communication with the Crosswork Controller.	Management

Note

SCP port can be tuned.

Table 28: Data traffic ports

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30649	TCP	Outbound	Data Gateway VM → Cluster VIP	Used for control communication with the Crosswork Network Controller.	Management
30993 30994 30995	TCP	Outbound	Data Gateway VM → Cluster VIP	Used by the Crosswork Data Gateway to send collected data to the Crosswork Kafka destinations.	Data

Table 29: Device traffic ports

Port	Protocol	Direction	Communication flow	Purpose	Network interface
22	TCP	Outbound	Data Gateway VM → Device	Used by the CLI Collector to initiate SSH connections for command-line data collection.	Data

Port	Protocol	Direction	Communication flow	Purpose	Network interface
161	UDP	Outbound	Data Gateway VM → Device	Used for SNMP polling of managed devices to collect performance data.	Data
1062	UDP	Inbound	Device → Data Gateway VM	Used to receive SNMP Trap traffic from managed devices.	Data
6514	TLS	Inbound	Device → Data Gateway VM	Used for secure, TLS-protected Syslog collection from managed devices.	Data
9010	TCP	Inbound	Device → Data Gateway VM	Used by the Crosswork Data Gateway to receive MDT traffic from devices.	Data
9514	UDP	Inbound	Device → Data Gateway VM	Used for UDP-based Syslog collection from managed devices.	Data
9898	TCP	Inbound	Device → Data Gateway VM	Used for TCP-based Syslog collection from managed devices.	Data
N/A	Site-specific	Outbound	Data Gateway VM → Device	gNMI Collector: Used for outbound gNMI data collection from devices.	Data

Single VM communication

These ports are specifically required when deploying Crosswork Network Controller on a single VM where the controller receives device traffic directly.

Table 30: Single VM communication ports - management network

Port	Protocol	Direction	Communication flow	Purpose	Network interface
7	TCP / UDP	Inbound	Device → Single VM VIP	Used for discovering endpoints and verifying reachability using ICMP.	Management (vNIC0)
22	TCP	Inbound	Device → Single VM VIP	Remote SSH connection	Management (vNIC0)
53	TCP/ UDP	Outbound	Single VM → DNS	Used for connection to DNS servers.	Management (vNIC0)
123	UDP	Outbound	Single VM → NTP	Used for Network Time Protocol (NTP).	Management (vNIC0)
830	TCP	Outbound	Single VM → NSO	Used for initiating NETCONF.	Management (vNIC0)
30602	TCP	Inbound	User browser → Single VM	Monitoring the installation (Crosswork Network Controller).	Management (vNIC0)
30603	TCP	Inbound	User browser → single VM VIP	Crosswork Network Controller Web user interface (NGINX server listens for secure connections on port 443).	Management (vNIC0)
30604	TCP	Inbound	Device → Single VM VIP	Used for Classic Zero Touch Provisioning (Classic ZTP) on the NGINX server.	Management (vNIC0)
30617	TCP	Inbound	Device → Single VM VIP	Used for Secure Zero Touch Provisioning (Secure ZTP) on the ZTP server.	Management (vNIC0/vNIC1)
30620	TCP	Inbound	Device → Single VM VIP	Used to receive plug-and-play HTTP traffic on the ZTP (PnP) server.	Management (vNIC0/vNIC1)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30653	TCP	Inbound	Site A ↔ Site B	Raft peer cluster communication port in geo HA deployments.	Management (vNIC0)
30695 30696	TCP	Inbound	Device → Single VM VIP	Embedded NSO	Management (vNIC0)

When configuring the ports for Embedded Collectors, ensure that the ports mentioned in these tables are configured on the device. For example, in case the port used for sending traps was previously set to 1062, change it to a port that is within the acceptable range for deploying a single VM. The acceptable range is provided with the port number in this table.

Table 31: Single VM communication ports - device network

Port	Protocol	Direction	Communication flow	Purpose	Network interface
22	TCP	Outbound	Single VM VIP → Device	• Used to initiate SSH connections with managed devices for configuration, monitoring, software upgrade & zero touch provisioning. • Used by embedded collector. • Used by embedded NSO (all CLI communications).	Data (vNIC1)
161	UDP	Outbound	Single VM VIP → Device	Used by embedded collector (all SNMP communications).	Data (vNIC1)
30514 Accepted range of ports is 30160-31560	TCP/ UDP	Inbound	Device → Single VM VIP	Used to receive standard Syslog traffic directly within Crosswork Network Controller in single VM deployments.	Data (vNIC1)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30614 Accepted range of ports is 30160-31560	TCP	Inbound	Device → Single VM VIP	Used to receive secure, TLS-protected Syslog traffic directly within Crosswork Network Controller in single VM deployments.	Data (vNIC 1)
30621 31121 31122 31123	TCP	Inbound	Device → Single VM VIP	Used for FTP services. This port is available only when the supported application is installed on Cisco Crosswork and the FTP settings are enabled. 31121:31123 are the additional ports used for FTP file transfer.	Data (vNIC 1)
30622	TCP	Inbound	Device → Single VM VIP	Used for SFTP services (available on data interface only). This port is available only when the supported application is installed on Cisco Crosswork and the SFTP settings are enabled.	Data (vNIC 1)
30898 Accepted range of ports is 30160-31560	TCP	Inbound	Device → Single VM VIP	Used to receive standard Syslog traffic directly within Crosswork Network Controller in single VM deployments.	Data (vNIC 1)
30910	TCP	Inbound	Device → Single VM VIP	Used to receive MDT traffic directly within Crosswork Network Controller in single VM deployments.	Data (vNIC 1)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
31062 Accepted range of ports is 30160-31560	TCP/ UDP	Inbound	Device → Single VM VIP	Used to receive SNMP and Trap traffic directly within Crosswork Network Controller in single VM deployments.	Data (vNIC 1)
50051	TCP	Outbound	Single VM VIP → Device	Used by embedded collector (all gRPC/gNMI communication).	Data (vNIC 1)
Site-specific [6] Ensure that port number on the device is the same as that configured on **Device Management > Network Devices > Edit Device**.	TCP	Outbound		gNMI collector	Data (vNIC 1)
Site-specific [7] To modify the user-defined destination ports, edit the port number from **Administration > Data destinations > Edit destination**.	Site-specific	Outbound		Kafka and gRPC destination	Data (vNIC 1)

Geo redundancy communication

These ports must be enabled between the active and standby Crosswork Network Controller cluster sites for high availability.

[6] For default port information of a device, refer to the platform-specific documentation.
[7] You cannot modify the port numbers of system-created destinations as they are created with predefined ports.

Table 32: Geo redundancy communication

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30180	TCP	Inbound	Site A ↔ Site B	Used for Postgres database replication via the PQ binary protocol, secured by SSL.	Management (vNIC0)
30190	TCP	Inbound	Site A ↔ Site B	Used for Timeseries database replication via the PQ binary protocol, secured by SSL.	Management (vNIC0)
30651 30653	TCP	Inbound	Site A ↔ Site B	Used for gRPC ping and Raft consensus, protected by TLS in geo-HA deployments.	Management (vNIC0)
30652	TCP	Inbound	Load balancer → Cluster VIP	Used by the Network Load Balancer for health checks in geo HA deployments.	Management (vNIC0)
30690	TCP	Inbound	Site A ↔ Site B	Used for NSO replication traffic between single VM deployments in geo HA, secured by mTLS.	Management (vNIC0)

Port	Protocol	Direction	Communication flow	Purpose	Network interface
30996 30997 30998	TCP	Inbound	Data Gateway → Cluster VIP	These ports enable communication between Crosswork clusters over the Management link in a geo-redundant deployment, allowing the Crosswork Data Gateway to send collected data to the Crosswork Kafka destination. This traffic occurs between Crosswork cluster sites.	Management (vNIC0)

For scenario-based port information for two-NIC and three-NIC Geo Redundancy deployments, refer to Geo redundancy port requirements on page 427.

IP address restrictions

Lists protected IPv4 and IPv6 subnets and describes network ranges that cannot be used for other devices or purposes in a Crosswork cluster.

Crosswork cluster uses specific IP ranges for internal communications by default, which cannot be used for other devices or purposes within your network. It is recommended to isolate your Crosswork cluster to ensure all communications remain internal and that there are no address space overlaps with external integration points (e.g., connections to devices, external servers, or the NSO server).

If these default IP ranges conflict with your network, collaborate with the Cisco Customer Experience team for assistance in modifying the settings before deployment.

Note

This is applicable for cluster installation and for adding a static route.

Note

The default values for the `K8sServiceNetwork` (10.96.0.0) and `K8sPodNetwork` (10.244.0.0) parameters can be changed.

Table 33: Protected IP subnets

IP Type	Subnet	Remarks
IPv4	172.17.0.0/16	Docker Subnet (Infrastructure)
	169.254.0.0/16	Link local address block
	127.0.0.0/8	Loopback address
	192.88.99.0/24	Reserved, used for relay servers to do IPv6 over IPv4
	240.0.0.0/4	Reserved for future use (previously class E block)
	224.0.0.0/4	MCAST-TEST-NET
	0.0.0.0/8	Current network, valid as source address only

IP Type	Subnet	Remarks
IPv6	2001:db8:1::/64	Docker Subnet (Infrastructure)
	fdfb:85ef:26ff::/48	Pod Subnet (Infrastructure)
	fd08:2eef:c2ee::/110	Service Subnet (Infrastructure)
	::1/128	Loopback address
	fe80::/10	Link local
	ff00::/8	IPv6 Multicast
	2002::/16	Reserved, used for relay servers to do IPv6 over IPv4
	2001:0000::/32	Terredo tunnel and relay
	2001:20::/28	Used by ORCHID and not IPv6 routable
	100::/64	Discard prefix, used in specific use-cases not applicable to Crosswork Zero Touch Provisioning
	::/128	Unspecified address, cannot be assigned to hosts
	::ffff:0:0/96	IPv4 mapped addresses
	::ffff:0:0:0/96	IPv4 translated addresses

What to do next:

- For installation instructions on VMware, refer to Install Cluster on VMware vCenter on page 81.
- For installation instructions on AWS EC2, refer to Install Cisco Crosswork Network Controller on AWS EC2 on page 225.
- For installation instructions for Single VM deployment, refer to Install Crosswork Network Controller on a Single VM on page 331.

Supported network topology models

Describes supported Crosswork topology models and lists NIC deployment modes for Crosswork VMs and integrated components in supported data center environments, including two- and three-NIC models.

This section introduces the different topology models that are supported when deploying Cisco Crosswork and the other solution components on a data center.

Routed and device networks

this table describes the types of traffic that comes from the Crosswork Network Controller. This traffic can use dual NICs.

Table 34: Types of Crosswork network traffic

Traffic	Description
Management	For accessing the UI and Crosswork Network Controller command line, and passing information between servers (for example, Cisco Crosswork to Crosswork Data Gateway or NSO).
Data	Data and configuration transfer between Cisco Crosswork and Crosswork Data Gateway and other data destinations (external Kafka/gRPC).
Device Access	The device access that the servers (Crosswork, NSO, Crosswork Data Gateway, or others) use to communicate with the managed devices in the network.

Connectivity between the various components should be accomplished via an external routing entity.

The Network Topology figures in this section show various line styles suggesting possible routing domains within the routed network.

- Solid - Management routing domain.
- Dotted - Data/Control routing domain (information transferred between Cisco Crosswork and Crosswork Data Gateway, and other data destinations (external Kafka or gRPC)).
- Dashes - Device access routing domain (from Crosswork Data Gateway and NSO).
- Blue dotted/dashed line - Alternate SR-PCE configuration path

The IP/subnet addressing scheme on each of these domains depends on the type of deployment.

Routing between domains is needed for Crosswork and NSO to reach the devices. However, proper firewall rules need to be in place to allow only select sources (for example, Crosswork and NSO) to reach the devices.

Note

- It is vital to have secure firewalls between Crosswork Network Controller and the network devices. However, the firewalls are not provided by Crosswork Network Controller and must be set up separately by the user. This topic highlights what application flows need to be allowed through the user-provided firewall system.
- On the device network, devices can be reached in-band or using out-of-band management interfaces, depending on the local security policies of each deployment.

Crosswork cluster configurations

The supported configurations are:

- **2 NIC Network Topology**: The Crosswork cluster, Crosswork Data Gateway, NSO, and SR-PCE use one network interface to communicate between their management interfaces, a second interface to pass the data between Crosswork Network Controller and Crosswork Data Gateway, and a routed interface to communicate with the network devices.
- **3 NIC Network Topology**: The Crosswork Data Gateway, NSO, and SR-PCE use one network interface to communicate between their management interfaces, a second interface to pass the data between Crosswork Network

Controller and Crosswork Data Gateway, and a third interface for Crosswork Data Gateway to communicate with the network devices.

> **Note**
>
> For simplicity, all these diagrams depict NSO using a management interface and a device-facing interface. Other configurations are also possible.

Figure 1: Cisco Crosswork two-NIC network topology

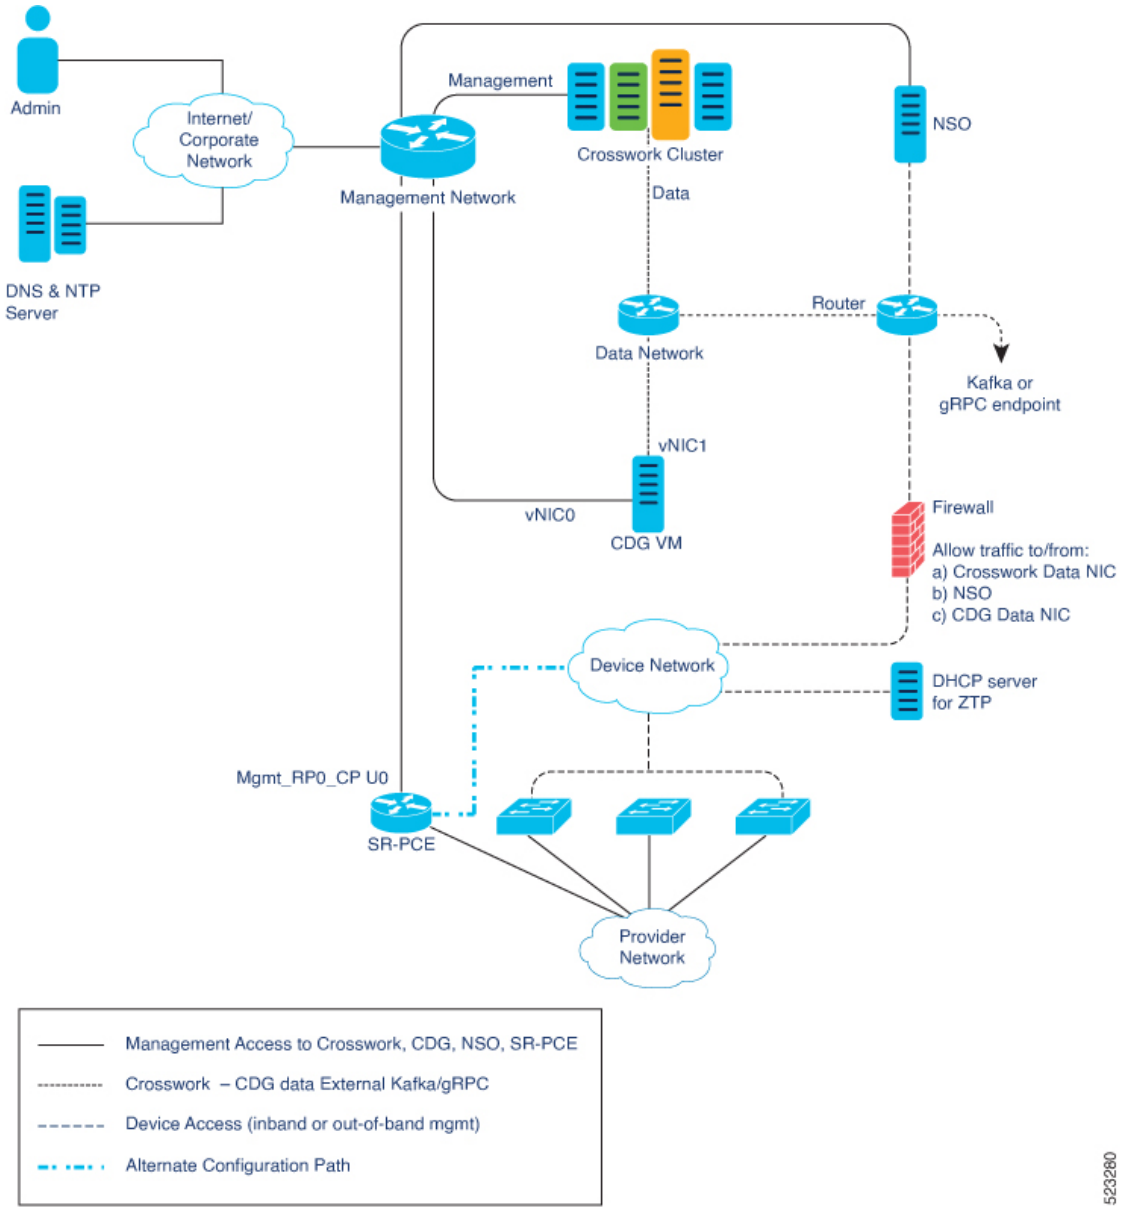
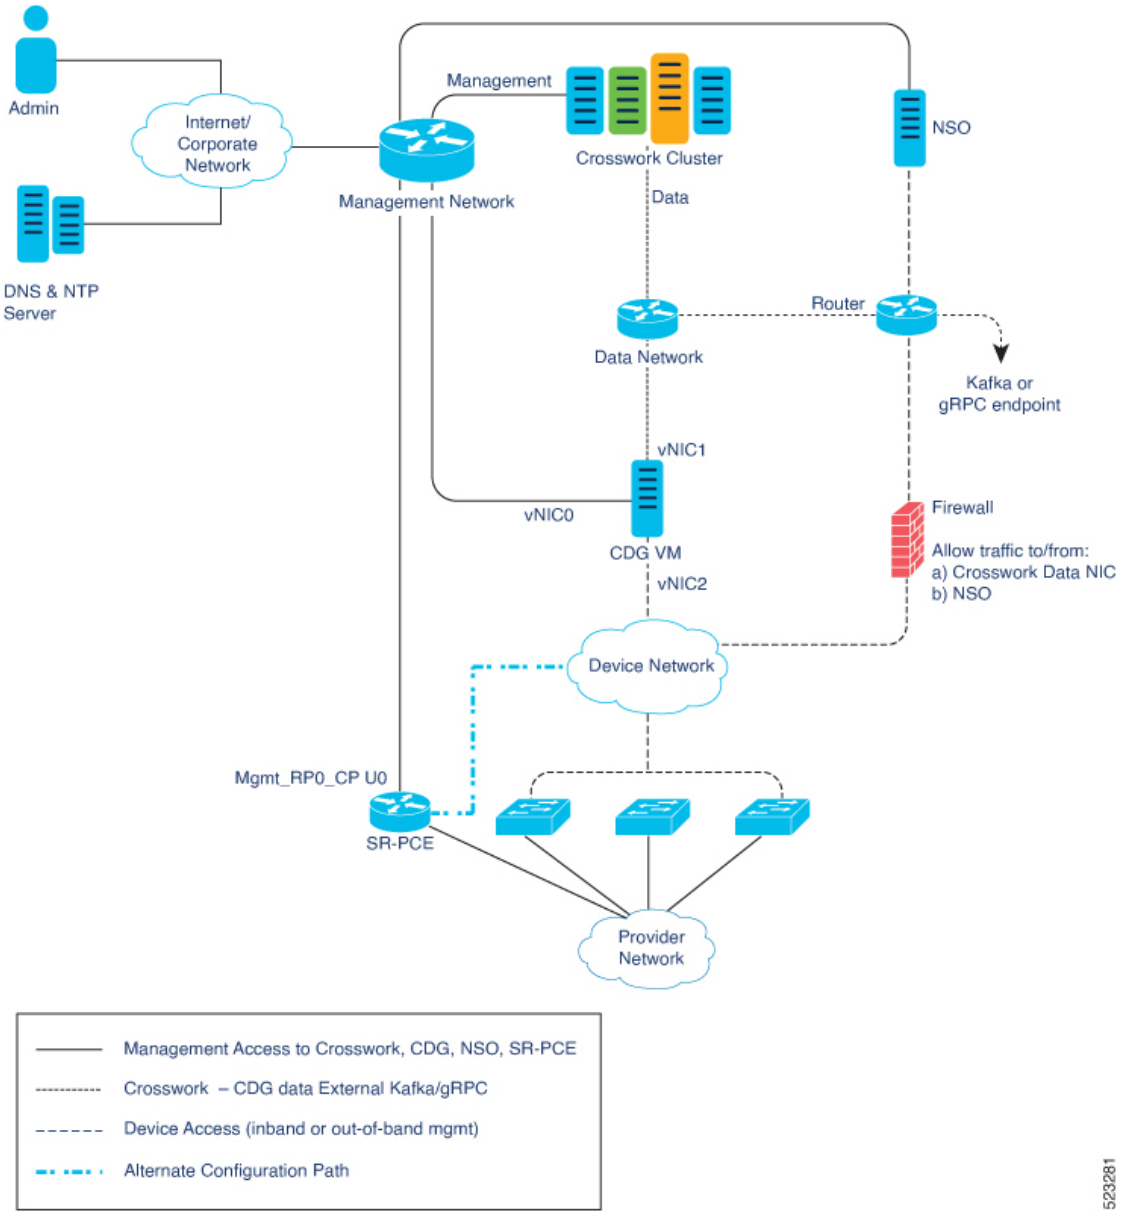


Figure 2: Cisco Crosswork three-NIC network topology



Cisco Crosswork virtual machine (VM)

The Cisco Crosswork VM has this vNIC deployment options:

Table 35: Cisco Crosswork vNIC deployment modes

No. of vNICs	vNIC	Description
2	Management	Management
	Data	Data and Device access

Crosswork Data Gateway VM

Overview of vNIC configuration

When configuring your Crosswork Data Gateway, the number of vNICs you use depends on your specific network requirements. If your Crosswork Cluster is set up with two interfaces, you can configure the Data Gateway to use either two or three vNICs, based on the needs of your deployment.

The number of vNICs may vary from one deployment to another. Factors such as security requirements and traffic isolation typically influence this decision.

vNIC mapping

The vNICs in Crosswork Data Gateway are identified by both their system interface names and their corresponding labels:

- vNIC0 corresponds to the eth0 interface
- vNIC1 corresponds to the eth1 interface
- vNIC2 corresponds to the eth2 interface

These system interface names (eth0, eth1, eth2) are used to refer to the respective vNICs in the operating system.

Available vNIC deployment options

The Crosswork Data Gateway VM has this vNIC deployment options:

Table 36: Crosswork Data Gateway default vNIC deployment modes

No. of vNICs	vNIC	Roles
2	vNIC0	Default Gateway, Administration, External Logging, and Management traffic.
	vNIC1	Control, Northbound External Data, and Southbound Data traffic.
3	vNIC0	Default Gateway, Administration, External Logging, and Management traffic.
	vNIC1	Control and Northbound External Data traffic.
	vNIC2	Southbound Data traffic

SR-PCE configuration

The Segment Routing Path Computation Element (SR-PCE) is both a device and a Software-Defined Networking (SDN) controller. Some deployments may want to treat an SR-PCE instance as a device, in which case they would need access via the device network. Some deployments may want to treat an SR-PCE instance as an SDN controller and access it on the Management routing domain. Crosswork supports both models. By default, Crosswork will use `eth0` (Management) to access SR-PCE as an SDN controller on the Management domain (shown in the figures). For more information on enabling Crosswork access to an SR-PCE instance as a device on the device network (shown as alternate path in the figures), please refer to Add Cisco SR-PCE Providers.

ZTP requirements

If you plan to use Zero Touch Provisioning, the device network needs to be equipped with a DHCP server (not provided as part of Cisco Crosswork). Some forms of ZTP also require a TFTP server (not provided as part of Cisco Crosswork). Additionally, all devices that use ZTP must have network connectivity to the Crosswork cluster as they will pull files (software and/or configuration) directly from the Crosswork cluster. For more information on Zero Touch Provisioning

concepts and features, refer to the *Zero Touch Provisioning* chapter in the *Cisco Crosswork Network Controller 7.2.x Device Lifecycle Management* .

Cisco Network Services Orchestrator (NSO) VM

The NSO VM has these vNICs:

- Management: Used for Crosswork applications to reach NSO.
- Device Access: Used for NSO to reach devices or NSO Resource Facing Services (RFS).

Part II

Install Cisco Crosswork Network Controller on VMware vCenter

Topics:

- Install Cluster on VMware vCenter
- Install Cisco Crosswork Data Gateway on VMware vCenter

5 Install Cluster on VMware vCenter

Topics:

- Installation overview
- VMware settings
- Installation parameters
- Install Cisco Crosswork on VMware vCenter using the Docker installer tool
- Manual installation of Cisco Crosswork using vCenter vSphere UI
- Monitor cluster activation
- Log into the Cisco Crosswork UI
- Import cluster inventory
- Troubleshoot the cluster

Summarizes VMware cluster installation options and prerequisites, procedures for using the Docker installer or vSphere UI, and guidance for monitoring and troubleshooting the cluster.

This chapter describes VMware cluster installation through the Docker installer tool or the vSphere UI. It also covers installation settings, activation monitoring, inventory import, diagnostics, and troubleshooting.

Installation overview

Explains the available VMware cluster installation methods, including the Docker installer tool and manual vSphere UI deployment, and identifies the preferred method.

The Crosswork Network Controller cluster can be installed using these methods:

- **Docker installer tool**: This day-0 tool used to deploy the Crosswork cluster with user specified parameters supplied via a template file. The tool is run from a Docker container which can be hosted on any Docker capable platform including a regular PC/laptop. The Docker container includes template files for each IP stack (IPv4, IPv6, and dual stack), which can be edited to provide deployment-specific data.
- **Manual installation (via the VMware UI)**: This option is available for deployments that cannot use the installer tool.

The installer tool method is the preferred option as it is faster and easier to use.

VMware settings

Lists supported VMware vCenter and ESXi versions, datastore types, and data center requirements that determine whether you can use the Docker installer or deploy the VMs individually.

these requirements are mandatory if you are planning to install Crosswork Network Controller using the Docker installer tool. If your vCenter data center does not meet these requirements, the VMs have to be deployed individually. For more information on manual installation, refer to Manual installation of Cisco Crosswork using vCenter vSphere UI on page 102.

- Hypervisor and vCenter supported:
  - VMware vCenter Server 8.0 (U2c or later) and ESXi 8.0 (U2b or later)
  - VMware vCenter Server 7.0 (U3p or later) and ESXi 7.0 (U3p or later)
- Crosswork Network Controller supports only VMFS and NFS datastore types.
- If you plan to use the docker installer tool, the machine where you run the installer must have network connectivity to the vCenter data center where you plan to install the cluster. If this mandatory requirement cannot be met, you must manually install the cluster. For more information on manual installation, refer to Manual installation of Cisco Crosswork using vCenter vSphere UI on page 102.
- Crosswork Network Controller cluster VMs (Hybrid nodes and Worker nodes) must be hosted on hardware with Hyper Threading disabled.
- As Crosswork Network Controller cluster nodes place high demands on the VMs, ensure that you have not oversubscribed CPU or memory resources on the machines hosting the nodes.
- All the physical host machines must be organized within the same VMware Data Center, and while it is possible to deploy all the cluster nodes on a single physical host (provided it meets the requirements), it is recommended that the nodes be distributed across multiple physical hosts. This prevents the host from being a single point of failure and improves solution resilience.
- Ensure that profile-driven storage is enabled by the vCenter admin user. Query permissions for the vCenter user at the root level (for all resources) of the vCenter.
- The networks required for the Crosswork Management and Data networks need to be built and configured in the data centers, and must allow low latency L2 communication (latency with RTT <= 10 ms).

> **Note**
>
> The same network names must be used and configured on all the ESXi host machines hosting the Crosswork VMs.

- To allow use of VRRP (Virtual Router Redundancy Protocol), the DVS Port group needs to be set in this list:

Property	Value
Promiscuous mode	Reject
MAC address changes	Reject

- The VRRP protocol requires unique router_id advertisements to be present on the network segment. The IDs can vary based on the deployment. For example, Crosswork usually uses the ID 169 on the management and ID 170 on the data network segments when multicast is used in discovery. If there is a symptom of conflict such as the VIP address not being reachable, check if any of the router IDs is duplicated and remove the conflicting VRRP router machines or use a different network.
- Ensure the user account you use for accessing vCenter has these privileges:
  - VM (Provisioning): Clone VM on the VM you are cloning.
  - VM (Provisioning): Customize on the VM or VM folder if you are customizing the guest operating system.
  - VM (Inventory): Create from the existing VM on the data center or VM folder.
  - VM (Configuration): Add a new disk on the data center or VM folder.
  - Resource: Assign a VM to a resource pool on the destination host, cluster, or resource pool.
  - Datastore: Allocate space on the destination datastore or datastore folder.
  - Network: Assign the network to which the VM will be assigned.
  - Profile-driven storage (Query): This permission setting needs to be allowed at the root of the data center tree level.
- We also recommend you to enable vCenter storage control.

Installation parameters

Outlines the parameters used to define a VMware cluster deployment, including network and VM values, installer settings, and Auto Action configuration.

This section explains the essential parameters that must be specified during the installation of the Crosswork cluster. Please ensure you have the relevant information for each parameter listed in the table and verify that your environment meets all prerequisite requirements.

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.

Note

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

Note

In case of dual stack deployment, you need to configure the IPv4 and IPv6 values for the Management, Data, and DNS parameters.

- `ManagementIPv4Address, ManagementIPv6Address`
- `ManagementIPv4Netmask, ManagementIPv6Netmask`
- `ManagementIPv4Gateway, ManagementIPv6Gateway`
- `ManagementVIPv4, ManagementVIPv6`
- `DataIPv4Address, DataIPv6Address`
- `DataIPv4Netmask, DataIPv6Netmask`
- `DataIPv4Gateway, DataIPv6Gateway`
- `DataVIPv4, DataVIPv6`
- `DNSv4, DNSv6`

Table 37: General parameters

Parameter name	Description
`ClusterName`	Name of the cluster file.
`ClusterIPStack`	The IP stack protocol: IPv4, IPv6, or DUALSTACK.
`ManagementIPAddress`	The Management IP address of the VM (IPv4 and/or IPv6).
`ManagementIPNetmask`	The Management IP subnet in dotted decimal format (IPv4 and/or IPv6).
`ManagementIPGateway`	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 (IPv4 and/or IPv6) for the cluster.
`DataIPAddress`	The Data IP address of the VM (IPv4 and/or IPv6).

Parameter name	Description
`DataIPNetmask`	The Data IP subnet in dotted decimal format (IPv4 and/or IPv6).
`DataIPGateway`	The Gateway IP on the Data Network (IPv4 and/or IPv6). The address must be reachable, otherwise the installation will fail.
`DataVIP`	The Data Virtual IP address (IPv4 and/or IPv6) for the cluster.
`DNS`	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`	Password to log into Cisco Crosswork. When setting up a VM, ensure the password is strong and meets these criteria: • It must be at least 8 characters long and include uppercase and lowercase letters, numbers, and at least one special character. • these special characters are not allowed: backslash (\), single quote ('), or double quote ("). • 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. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`VMSize`	Sets the VM size for the cluster. For cluster deployments, the only supported option is `"Large"`. **Note** • If you leave this field blank, the default value (`"Large"`) is automatically selected. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`VMName`	Name of the VM. A unique VM name is required for each node on the cluster (Hybrid or Worker).

Parameter name	Description
`NodeType`	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. **Note** The Crosswork cluster requires at least three VMs operating in a hybrid configuration.
`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. **Note** The final backup partition size will be calculated dynamically. This parameter defines the minimum.
`ThinProvisioned`	Set to **false** for production deployments.
`SchemaVersion`	The configuration Manifest schema version. This indicates the version of the installer to use with this template. Schema version should map to the version packaged with the sample template in the Docker installer tool on cisco.com. 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.
`LogFsSize`	Log partition size (in Giga Bytes). Minimum value is 20 GB and Maximum value is 1000 GB. You are recommended to use the default value.

Parameter name	Description
`EnableSkipAutoInstallFeature`	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"`. The recommended value for cluster deployment is `"False"`. **Note** • If left blank, the default value is automatically selected. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`EnforcePodReservations`	Enforces minimum resource reservations for the pod. If left blank, the default value (`"True"`) is selected. This parameter accepts a string value, so be sure to enclose the value in double quotes.
`K8sServiceNetwork`	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.
`K8sPodNetwork`	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.

Parameter name	Description
`IgnoreDiagnosticsCheckFailure`	Used to set the system response in case of a diagnostic check failure. 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. You are recommended to select the default value. This parameter accepts a string value, so be sure to enclose the value in double quotes. **Note** • The log files (`diagnostic_stdout.log` and `diagnostic_stderr.log`) can be found at **/var/log**. The result from each diagnostic execution is kept in a file at **/home/cw-admin/diagnosis_report.txt**. • Use **diagnostic all** command to invoke the diagnostic manually on day N. • Use **diagnostic history** command to view previous test report.
`ManagementVIPName`	Name of the Management Virtual IP for the Crosswork VM. This is an optional parameter used to reach Crosswork Management VIP via DNS name. If this parameter is used, the corresponding DNS record must exist in the DNS server.
`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.
`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. This parameter is optional for installing using the Docker installer tool.
`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. For more information on VM count, refer to Table 1: Crosswork Network Controller tiers on page 19. This parameter is optional for installing using the Docker installer tool.

Parameter name	Description
`InitMasterCount`	Total number of Hybrid nodes in the cluster. The default value is 3. This parameter is optional for installing using the Docker installer tool.
`EnableHardReservations`	Determines the enforcement of VM CPU and memory profile reservations. This is an optional parameter and the default value is `"True"`, if not explicitly specified. **Note** This parameter accepts a string value, so be sure to enclose the value in double quotes. 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. 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.

Parameter name	Description
`RamDiskSize`	Size of the Ram disk. **Warning** Use the `RamDiskSize` option only for temporary lab testing or product demonstrations. This option is intended for environments where disk performance does not meet the application requirements and where rebuilding the VMs after a restart is an acceptable trade-off. Do not use this option for deployments that require resiliency or reliability. The minimum value for `RamDiskSize` is `2`. When you set `RamDiskSize` to a non-zero value, leave `HSDatastore` unset as `""`. If the VM is powered down or restarted after installation with `RamDiskSize` configured, the system may be unable to launch applications, and Kubernetes may report errors.
`ManagerDataFsSize`	This parameter is applicable only when installing with the Docker installer tool. Refers to the data disk size for Hybrid 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.
`WorkerDataFsSize`	This parameter is applicable only when installing with the Docker installer tool. Refers to the data disk size for Worker nodes (in Gigabytes). 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.

Parameter name	Description
`Timezone`	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. This is an optional parameter. **Note** 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 `Timezone` parameter in Crosswork.

Table 38: VMware template parameters

Parameter name	Description
`VCenterAddress`	The vCenter IP or host name.
`VCenterUser`	The username needed to log into vCenter.
`VCenterPassword`	The password needed to log into vCenter.
`DCname`	The name of the Data Center resource to use. Example: `DCname = "WW-DCN-Solutions"`
`MgmtNetworkName`	The name of the vCenter network to attach to the VM's Management interface. This network must already exist in VMware or the installation will fail.
`DataNetworkName`	The name of the vCenter network to attach to the VM's Data interface. This network must already exist in VMware or the installation will fail.
`Host`	The ESXi host, or ONLY the vcenter cluster/resource group name where the VM is to be deployed. The primary option is to use the host IP or name (all the hosts should be under the data center). If the hosts are under a cluster in the data center, only provide the cluster name (all hosts within the cluster will be picked up). The subsequent option is to use a resource group. In this case, a full path should be provided. Example: `Host = "Main infrastructure/Resources/00_trial"`

Parameter name	Description
`Datastore`	The datastore name available to be used by this host or resource group. The primary option is to use host IP or name. The subsequent option is to use a resource group. Example: `Datastore = "SDRS-DCNSOL-prodexsi/bru-netapp-01_FC_Prodesx_ds_15"`
`HSDatastore`	The high speed datastore available for this host or resource group. When not using a highspeed data store, set to same value as Datastore.
`DCfolder`	The resource folder name on vCenter. To be used if you do not have root access as a VMware user, or when you need to create VMs in separate folders for maintenance purposes. You must provide the complete path as value for the DCfolder. Example: `DCfolder = "/WW-DCN-Solutions/vm/00_trial"` Please contact your VMware administrator for any queries regarding the complete folder path. Leave as empty if not used.
`Cw_VM_Image`	The name of Crosswork cluster VM image in vCenter. This value is set as an option when running the Docker installer tool and does not need to be set in the template file.
`HostedCwVMs`	The IDs of the VMs to be hosted by the ESXi host or resource.

After you have decided the installation parameters values for Crosswork Network Controller, choose the method you prefer and begin your deployment:

- Install Cisco Crosswork on VMware vCenter using the Docker installer tool on page 94
- Manual installation of Cisco Crosswork using vCenter vSphere UI on page 102

Automate application installation

Outlines Auto Action functionality for automating application installation and activation in day-0 and day-N deployments, including definition files and sample templates.

The auto-action functionality is an optional feature that allows you to automate the installation and activation of applications as needed. Auto-action is supported for both cluster and single VM deployments.

Day-0 deployment

The auto-action feature simplifies installation process and can be configured with the Docker installer, direct OVA installation, and the OVF tool.

To enable auto-action, you must configure a definition file (JSON format) that lists the tar bundles to be imported and activated. The JSON file is submitted alongside the day-0 installer and overrides the default file bundled in the OVA during installation.

> **Note**
>
> Please ensure that the file path specified in the auto-action file (the location on your local machine where the files are downloaded) is accessible from the Crosswork cluster.

The auto-action definition file customizes two actions:

- `add_to_repository_requests`: The auto-action functionality supports HTTP, HTTPS, and SCP protocols with basic authentication (HTTP/HTTPS) to add the application file to the repository.
- `install_activate_requests`: The application files to be installed and activated are identified using the **version** and **id** parameters.

Auto-action feature can be customized with various deployment methods:

- **Using Docker Installer**: While installing via the docker installer, if the JSON file is successfully validated, the installation will proceed. If there are syntax errors in the file, you will be prompted with an error message, and the installation will be halted. Once the errors are corrected, you can retry the installation.

  Syntax to execute the auto-action file:

  ```
  ./cw-installer.sh install -p -m /data/<template file name> -a <path to json
   def file> -o /data/<.ova file>
  ```

  Example:

  ```
  ./cw-installer.sh install -m /data/deployment.tfvars -a
  https://example.com/path/to/crosswork_auto_action.json -o
  /data/cnc-platform-cluster-deployment-7.2.1.ova
  ```

- **Using vCenter UI or OVF tool**: While installing via the vCenter UI, the CDATA JSON content is validated in the backend. If there are syntax errors in the data, the auto-action instructions are skipped and Crosswork Cluster is installed as per the regular installation workflow. vCenter and the OVF tool do not support the direct upload of JSON format files due to issues with handling special characters. To resolve this, you must compress the JSON file and enclose it in CDATA format.

  ```
  <![CDATA[{auto-action json compressed content}]]>
  ```

  > **Note**
  >
  > The CDATA example has line breaks for the purpose of readability. During production deployment, kindly execute the CDATA without any line breaks.

  Example:

  ```
  <![CDATA[{"auto_action":{"add_to_repository_requests":[{"file_location":
  {"uri_location":{"uri":"<file path>/<filename.tar.gz>"}}}],
  "install_activate_requests":[{"package_identifier":{"id":"capp-coe","version":"7.2.1"}}]}}]]>
  ```

Product-specific definition

You can define any product-specific parameters in a JSON file and execute it with the auto-action file.

```
./cw-installer.sh install -m /data/<template file name> -o /data/<.ova file>
-c /data/<.json file> -y
```

Day-N deployment

Auto-action can be used on day-N to automate installation of Crosswork Network Controller components or application patches. For more information, refer to Add the auto-action file on page 269.

Sample templates

Refer to the example in Sample auto action templates on page 555

Install Cisco Crosswork on VMware vCenter using the Docker installer tool

Deploys a Cisco Crosswork cluster on VMware vCenter using the Docker installer tool, with guidance on preparing files, configuring the template, and responding to installer prompts.

Use the Docker installer tool to install a Cisco Crosswork cluster on VMware vCenter.

This section explains the procedure to install Cisco Crosswork on VMware vCenter using the Docker installer tool.

> **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

Pointers to know when using the Docket installer tool:

- Make sure that your environment meets all the vCenter requirements specified in Installation Prerequisites on page 43.
- If you intend to use a dual stack configuration for your deployment, make sure that the host machine running the Docker installer meets these requirements:
  - 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 `--network host` when running the Docker installer.
  - Confirm that the provided IPv6 network CIDR and gateway are valid and reachable.
- The edited template in the `/data` directory will contain sensitive information (VM passwords and the vCenter password). The operator needs to manage access to this content. Store the templates used for your install in a secure environment or edit them to remove the passwords after the install is complete.
- The `install.log`, `install_tf.log`, and `.tfstate` files will be created during the install and stored in the `/data` directory. If you encounter any trouble with the installation, provide these files to the Cisco Customer Experience team when opening a case.
- The install script is safe to run multiple times. Upon error, input parameters can be corrected and rerun. Running the Docker installer tool multiple times may result in the deletion and re-creation of VMs.

- In case you are using the same Docker installer tool for multiple Crosswork cluster installations, it is important to run the tool from different local directories, allowing for the deployment state files to be independent. The simplest way for doing so is to create a local directory for each deployment on the host machine and map each one to the container accordingly.
- Docker version 19 or higher is required while using the Docker installer tool option. For more information on Docker, refer to https://docs.docker.com/get-docker/
- In order to change install parameters or to correct parameters following installation errors, it is important to distinguish whether the installation has managed to deploy the VMs or not. Deployed VMs are evidenced by the output of the installer similar to: *vsphere_virtual_machine.crosswork-IPv4-vm["1"]: Creation complete after 2m50s [id=4214a520-c53f-f29c-80b3-25916e6c297f]*
- In case of deployed VMs, changes to the Crosswork VM settings or the Data Center host for a deployed VM are NOT supported. To change a setting using the installer when the deployed VMs are present, the clean operation needs to be run and the cluster redeployed. For more information, refer to Delete the VM using the Docker installer tool on page 414.
- A VM redeployment will delete the VM's data, hence caution is advised. We recommend you perform VM parameter changes from the Crosswork UI, or alternatively one VM at a time.
- Installation parameter changes that occur prior to any VM deployment, e.g. an incorrect vCenter parameter, can be performed by applying the change and simply re-running the install operation.
- If you want to use the auto action functionality, the definition file (JSON format) must be specified while executing the OVA file. For more information, refer to Automate application installation on page 92.

Known limitations:

These following scenarios are the caveats for installing the Crosswork cluster using the Docker installer tool.

- The vCenter host VMs defined must use the same network names (vSwitch) across all hosts in the data center.
- The vCenter storage folders or datastores organized under a virtual folder structure, are not supported currently. Please ensure that the datastores referenced are not grouped under a folder.
- Any VMs that are not created by the Day 0 installer (for example, manually brought up VMs), cannot be changed either by the Day 0 installer or via the Crosswork UI later. Similarly, VMs created via the Crosswork UI cannot be modified using the Day 0 installer. When making modifications after the initial deployment of the cluster, ensure that you capture the inventory information.
- vCenter UI provides a service where a user accessing via IPv4 can upload images to the IPv6 ESXi host. Docker installer tool cannot use this service. Follow either of these workarounds for IPv6 ESXi hosts:
  1. Install the OVA template image manually (for more information, refer to Manual installation of Cisco Crosswork using vCenter vSphere UI on page 102).
  2. Run the Docker installer tool from an IPv6 enabled machine. To do this, configure the Docker daemon to map an IPv6 address into the docked container.

Note

The Docker installer tool will deploy the software and power on the virtual machines. If you wish to power on the virtual machines yourself, use the manual installation.

Procedure

1. In your vCenter data center, go to **Host** > **Configure** > **Networking** > **Virtual Switches** and select the virtual switch. In the virtual switch, select **Edit** > **Security**, and configure these DVS port group properties:
   - Set **Promiscuous mode** as *Reject*
   - Set **MAC address changes** as *Reject*

   Confirm the settings and repeat the process for each virtual switch used in the cluster.
2. Prepare the installation files.
   a) In your Docker capable machine, create a directory where you will store everything you will use during this installation.

   Note

   If you are using a Mac, please ensure that the directory name is in lower case.

   b) Download the installer bundle (.tar.gz file) and the OVA file from cisco.com to the directory you created previously.

   Table 39: Crosswork network controller installation files

Release	Files
7.2.0	**Installer**: `cnc-platform-cluster-docker-deployment-7.2.0-45.tar.gz` **OVA**: `cnc-platform-cluster-deployment-7.2.0-45.ova`
7.2.1 SP	**Installer**: `cnc-platform-cluster-docker-deployment-7.2.1.tar.gz` **OVA**: `cnc-platform-cluster-deployment-7.2.1.ova`

   For the purpose of these instructions, we will use the file names as **cnc-platform-cluster-docker-deployment-7.2.1.tar.gz** and **cnc-platform-cluster-deployment-7.2.1.ova** respectively.

> **Note**
>
> The file names mentioned in this topic are sample names and may differ from the actual file names in cisco.com.

3. Unpack the installer bundle and review its contents.
   a) Use these command to unzip the installer bundle:

   ```
   tar -xvf cnc-platform-cluster-docker-deployment-7.2.1.tar.gz
   ```

   The contents of the installer bundle are unzipped to a new directory (e.g. `cnc-platform-cluster-docker-deployment-7.2.1`). This new directory will contain the installer image (**cnc-platform-cluster-docker-deployment-7.2.1-123.tar.gz**) and files necessary to validate the image.

   b) Change the directory to the directory created by opening the file and then review the contents of the README file in order to understand everything that is in the package and how it will be validated in these steps.

4. Use these command to verify the signature of the installer image:
   a) Ensure you have python installed. If not, go to python.org and download the version of python that is appropriate for your work station.
   b) Use `python --version` to find out the version of python on your machine.
   c) Depending on the python version use one of these commands to validate the file.

   If you are using python 2.x, use these command to validate the file:

   ```
   python cisco_x509_verify_release.py -e <.cer file> -i <.tar.gz file> -s
   <.tar.gz.signature file>
   -v dgst -sha512
   ```

   If you are using python 3.x, use these command to validate the file:

   ```
   python cisco_x509_verify_release.py3 -e <.cer file> -i <.tar.gz file> -s
   <.tar.gz.signature file>
   -v dgst -sha512
   ```

5. Load the installer image and identify its Docker image ID.
   a) Use these command to load the installer image file into your Docker environment.

   ```
   docker load -i <.tar.gz file>
   ```

   For example:

   ```
   docker load -i cnc-platform-cluster-docker-deployment-7.2.1-123.tar.gz
   ```

   b) Run Docker image list or Docker images command to get the "image ID" (which is needed in the next step).

   For example:

   ```
   docker images
   ```

The result will be similar to these: (section we will need is underlined for clarity)

```
My Machine% docker images
REPOSITORY                          TAG
             IMAGE ID             CREATED        SIZE
dockerhub.cisco.com/cw-installer
cnc-platform-cluster-docker-deployment-7.2.1-123   a4570324fad30   7 days
ago     276MB
```

Note

Pay attention to the "CREATED" time stamp in the table presented when you run `docker images`, as you might have other images present from the installation of prior releases. If you wish to remove these, the `docker image rm {image id}` command can be used.

6. Start the installer and configure the deployment parameters.
   a) Launch the Docker container using these command:

```
docker run --rm -it -v `pwd`:/data {image id of the installer container}
```

To run the image loaded in our example, the command would be:

```
docker run --rm -it -v `pwd`:/data a4570324fad30
```

Note

- You do not have to enter that full value. In this case, "docker run --rm -it -v `pwd`:/data a45" was adequate. Docker requires enough of the image ID to uniquely identify the image you want to use for the installation.
- In the above command, we are using the backtick (`). Do not use the single quote or apostrophe (') as the meaning to the shell is very different. By using the backtick (recommended), the template file and OVA will be stored in the directory where you are on your local disk when you run the commands, instead of inside the container.
- When deploying a IPv6 or dual stack cluster, the installer needs to run on an IPv6 enabled container/VM. This requires additionally configuring the Docker daemon before running the installer, using these method:
  - **Linux hosts (ONLY)**: Run the Docker container in host networking mode by adding the "-network host" flag to the Docker run command line.

    ```
    docker run --network host <remainder of docker run
     options>
    ```

- Centos/RHEL hosts, by default, enforce a strict SELinux policy which does not allow the installer container to read from or write to the mounted data volume. On such hosts, run the Docker volume command with the Z option as shown below:

  ```
  docker run --rm -it -v `pwd`:/data:Z <remainder of docker
   options>
  ```

b) Navigate to the directory with the VMware template.

`cd /opt/installer`

c) Copy the template file found under `/opt/installer/deployment_template_tfvars` to the `/data` folder using a different name.

For example: `cp deployment_template_tfvars /data/deployment.tfvars`

For the rest of this procedure, we will use `deployment.tfvars` in all the examples.

d) Edit the template file located in the `/data` directory in a text editor, to match your planned deployment. Refer to the Installation parameters on page 83 table for details on the required and optional fields and their proper settings. The Crosswork Network Controller cluster deployment templates for VMware vCenter on page 526 includes an example that you can reference for proper formatting. The example is more compact due to the removal of descriptive comments:

7. Run the installation and respond to the installer prompts.

a) From the `/opt/installer` directory, run the installer.

```
./cw-installer.sh install -p -m /data/<template file name> -o /data/<.ova
 file> -y
```

For example:

```
./cw-installer.sh install -p -m /data/deployment.tfvars -o
/data/cnc-platform-cluster-deployment-7.2.1.ova -y
```

> **Note**
>
> If you want to use the auto action functionality, the definition file (JSON format) must be specified while executing the OVA file in these format:
>
> ```
> ./cw-installer.sh install -p -m /data/<template file name> -a <path to
> json def file> -o /data/<.ova file>
> ```

b) Read, and then enter "yes" if you accept the End User License Agreement (EULA). Otherwise, exit the installer and contact your Cisco representative.
c) Enter "yes" when prompted to confirm the operation.

Note

It is not uncommon to refer to some warnings like these during the install:

```
Warning: Line 119: No space left for device '8' on parent controller
'3'.
Warning: Line 114: Unable to parse 'enableMPTSupport' for attribute 'key'
 on element 'Config'.
```

If the install process proceeds to a successful conclusion (refer to sample output below), these warnings can be ignored.

Sample output:

```
cw_cluster_vms = <sensitive>
INFO: Copying Day 0 state inventory to CW
INFO: Waiting for deployment status server to startup on 10.90.147.66.
Elapsed time 0s, retrying in 30s
Crosswork deployment status available at
http://{VIP}:30602/d/NK1bwVxGk/crosswork-deployment-readiness?orgId=1&refresh=10s&theme=dark

Once deployment is complete login to Crosswork via:
https://{VIP}:30603/#/logincontroller
INFO: Cw Installer operation complete.
```

Note

If the installation fails due to a timeout, you should try rerunning the installation (step 13) without the `-p` option. This will deploy the VMs serially rather than in parallel.

If the installer fails for any other reason (for example, mistyped IP address), correct the error and rerun the install script.

If the installation fails (with or without the -p), open a case with Cisco and provide the .log files that were created in the /data directory (and the local directory where you launched the installer docker container), to Cisco for review. The two most common reasons for the install to fail are: (a) password that is not adequately complex, and (b) errors in the template file.

What to do next

- refer to Monitor cluster activation on page 113 to know how you can check the status of the installation.
- refer to Troubleshoot the cluster on page 122 for common troubleshooting scenarios.

Manual installation of Cisco Crosswork using vCenter vSphere UI

Summarizes manual Cisco Crosswork installation through the VMware vSphere UI and provides related procedures for importing the certificate chain, building the OVF template, and deploying it.

This section explains how to build the cluster using the vCenter user interface. This same procedure can be used to add or replace nodes if necessary.

The manual installation workflow is broken into two parts. In the first part, you create a template. In the second part, you deploy the template as many times as needed to build the cluster of 3 Hybrid nodes (typically) along with any Worker nodes that your environment requires.

1. Build the OVF template on page 103
2. Deploy the template on page 109

> **Note**
>
> If the cluster has already been installed (no matter the method used) the template file will already exist unless it was deleted. In this case, you can directly go to deploying the template (the second part of this procedure).

The manual installation is preferred if you face any of these reasons:

- Owing to your data center configuration, you cannot deploy the cluster using the installer tool.
- You need to add nodes to the existing cluster.
- You need to replace a failed node.
- You want to migrate a node to a new host machine.

> **Note**
>
> Anytime the configuration of the cluster is changed manually–whether to install the Crosswork cluster, add nodes, or move nodes to new hosts using the procedures detailed in this section–you must import the cluster inventory file (.tfvars file) to the Crosswork UI. You must set the parameter `OP_Status = 2` to enable manual import of the inventory. For more information, refer to Import cluster inventory on page 122.

Import the CA certificate chain on VMware

Imports the certificate chain VMware uses to validate a signed OVA before you deploy it through the vSphere UI. The procedure identifies the certificate files to import.

Use this procedure to import the CA certificate chain that VMware uses to validate signed OVA images during deployment from the VMware vSphere UI.

Starting with VMware ESXi 7.0 U2, VMware validates the certificate chain of signed OVA images during deployment. If the required certificate chain is not available in the VMware trust store, VMware can display an invalid certificate warning.

This requirement applies when you deploy signed OVA images from the VMware vSphere UI, including Crosswork Network Controller cluster and Crosswork Data Gateway deployments.

> **Note**
>
> The code-signing certificate validity is 15 months. If the certificate expires and the OVA image is re-signed, use the PEM file associated with the reposted image. For any assistance, contact the Cisco Customer Experience team.

Follow these steps to import the CA certificate chain:

Before you begin

Obtain the intermediate CA certificate and the root CA certificate for the signed OVA image that you plan to deploy.

- Confirm that the PEM files apply to the currently posted OVA image.
- Perform this procedure in VMware vCenter before you deploy the OVA image from the VMware vSphere UI.
- For Crosswork Network Controller Release 7.2.1 SP, use these PEM files:
  - *DigiCert Trusted G4 Code Signing RSA4096 SHA384 2021 CA1*: https://cacerts.digicert.com/DigiCertTrustedG4CodeSigningRSA4096SHA3842021CA1.crt.pem
  - *DigiCert Trusted Root G4*: https://cacerts.digicert.com/DigiCertTrustedRootG4.crt.pem

Procedure

1. Log in to the vCenter UI as an administrator.
2. Choose **Menu** > **Administration**.
3. Under **Certificates**, open **Certificate Management**.
4. Locate **Trusted Root Certificates** and select **Add**.
5. Import the intermediate certificate.

   Use the PEM file for *DigiCert Trusted G4 Code Signing RSA4096 SHA384 2021 CA1*.
6. Import the root certificate.

   Use the PEM file for *DigiCert Trusted Root G4*.
7. Verify that the imported certificate chain appears as a trusted certificate.

The VMware trust store contains the intermediate and root CA certificates that sign the OVA image.

What to do next

Proceed with the OVA deployment from the VMware vSphere UI.

Build the OVF template

Builds a Cisco Crosswork OVF template through vSphere, including the configuration, storage, and network selections needed before you deploy the cluster VMs.

Build an OVF template for Cisco Crosswork in VMware vSphere.

Before you begin

- Make sure that your environment meets all the vCenter requirements specified under Crosswork cluster VM requirements on page 44 and Installation Prerequisites on page 43.

- Import the required CA certificate chain into the VMware trust store before you deploy the signed OVA image from the VMware vSphere UI. For more information, refer to Import the CA certificate chain on VMware on page 102.

Procedure

1. Download the platform image and prepare the vSphere host and network.
   a) Download the latest available Cisco Crosswork platform image file (*.ova) to your system.
   b) With VMware ESXi running, log into the VMware vSphere Web Client. On the left navigation pane, choose the ESXi host or cluster where you want to deploy the VM.
   c) In the vSphere UI, go to **Host** > **Configure** > **Networking** > **Virtual Switches** and select the virtual switch. In the virtual switch, select **Edit** > **Security**, and configure these DVS port group properties:
      - Set **Promiscuous mode** as *Reject*
      - Set **MAC address changes** as *Reject*

      Confirm the settings and repeat the process for each virtual switch used in the cluster.
2. Review and confirm that your network settings meet the requirements.

   Ensure that the networks that you plan to use for Management network and Data network are connected to each host where VMs will be deployed.
3. Start the OVF deployment and provide the initial deployment details.
   a) Choose **Actions** > **Deploy OVF Template**.

   > ⚠ **Caution**
   >
   > The default VMware vCenter deployment timeout is 15 minutes. If vCenter times out during deployment, the resulting VM will not be bootable. To prevent this, we recommend that you document the choices (such as IP address, gateway, DNS server, etc.) so that you can enter the information quickly and avoid any issues with the VMware configuration.

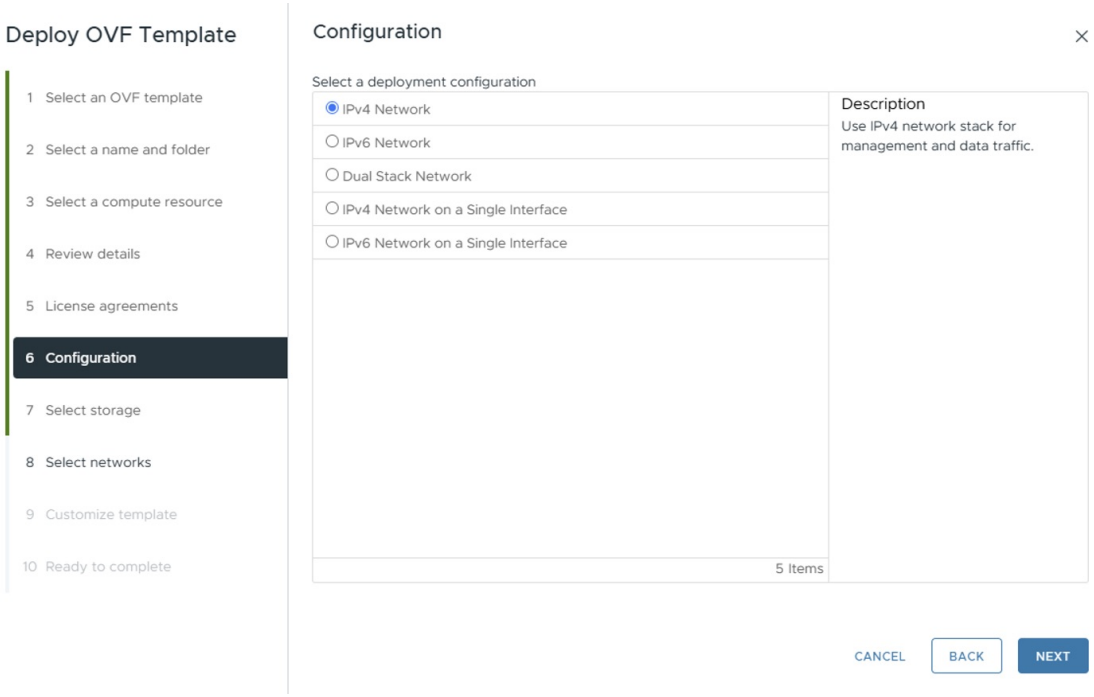
   b) The VMware **Deploy OVF Template** window appears, with the first step, **1 – Select an OVF template**, highlighted. select **Choose Files** to navigate to the location where you downloaded the OVA image file and select it. Once selected, the file name is displayed in the window.
   c) select **Next**. The **Deploy OVF Template** window is refreshed, with **2 – Select a name and folder** now highlighted. Enter a name and select the respective data center for the Cisco Crosswork VM you are creating.

      We recommend that you include the Cisco Crosswork version and build number in the name, for example: Cisco Crosswork 7.2.x Build 48.
   d) select **Next**. The **Deploy OVF Template** window is refreshed, with **3 – Select a compute resource** highlighted. Select the host or cluster for your Cisco Crosswork VM.
   e) select **Next**. The VMware vCenter Server validates the OVA. Network speed will determine how long validation takes. After the validation is complete, the **Deploy OVF Template** window is refreshed, with **4 – Review details** highlighted.
4. Review the template details and license.
   a) Review the OVF template that you are deploying. Note that this information is gathered from the OVF, and cannot be modified.

b) select **Next**. The **Deploy OVF Template** window is refreshed, with **5 – License agreements** highlighted. Review the End User License Agreement and if you agree, select the **I accept all license agreements** checkbox. Otherwise, contact your Cisco Experience team for assistance.

c) select **Next** The **Deploy OVF Template** window is refreshed, with **6 – Configuration** highlighted. Crosswork supports only these deployment configurations: **IPv4 Network**, **IPv6 Network**, and **Dual Stack Network**. Please select your preferred deployment configuration.

Figure 3: Select a deployment configuration



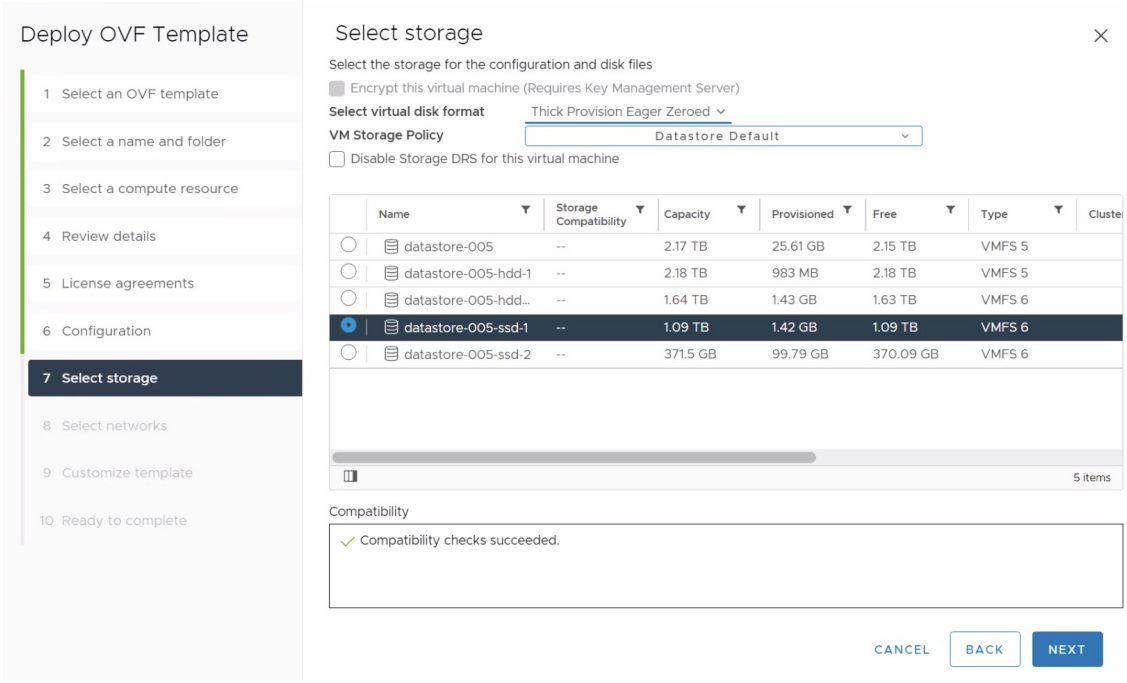
5. Select storage and map the deployment networks.

a) select **Next**. The **Deploy OVF Template** window is refreshed, with **7 – Select Storage** highlighted. Choose the relevant option from the **Select virtual disk format** drop-down list.

> **Note**
>
> For production deployment, choose the **Thick Provision Eager Zeroed** option because this will preallocate disk space and provide the best performance. For lab purposes, we recommend the **Thin Provision** option because it saves disk space.

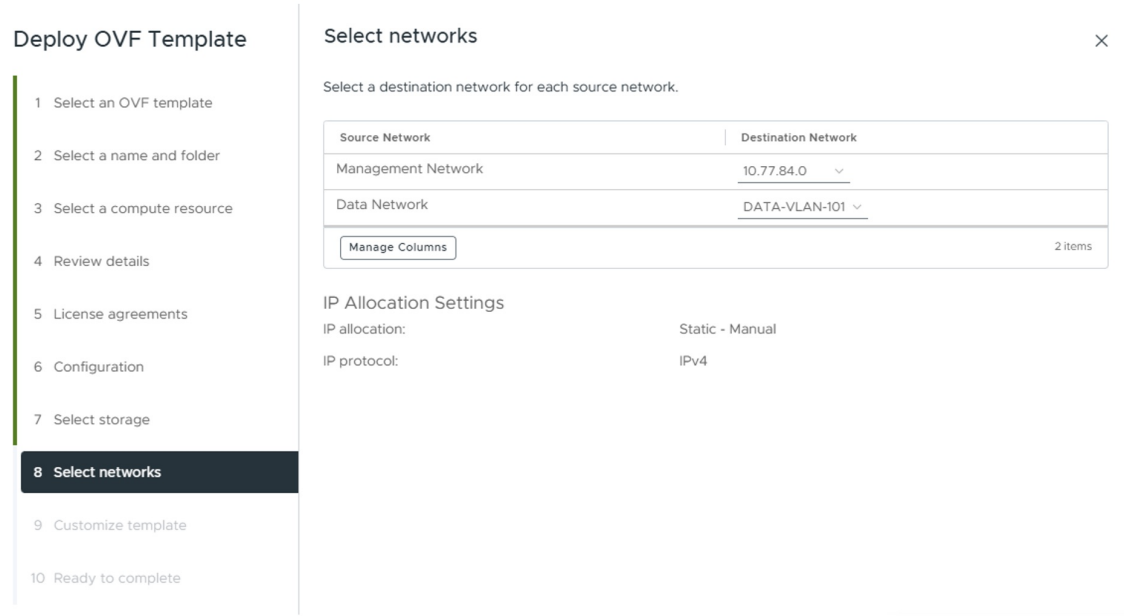
From the table, choose the datastore you want to use, and review its properties to ensure there is enough available storage.

Figure 4: Select storage



b) select **Next**. The **Deploy OVF Template** window is refreshed, with **8 – Select networks** highlighted. From the **Destination Network** drop-down list, select the proper networks for the Management Network and the Data Network.

Figure 5: Select networks



6. select **Next**. The **Deploy OVF Template** window is refreshed, with **9 – Customize template** highlighted.

> **Note**
>
> As you are creating a template now, enter the IP information for the first node.

a) Expand the **Management Network** settings. Provide information for the IPv4, IPv6 or dual stack deployment (as per your selection) such as IP address, IP netmask, IP gateway, and virtual IP address.
b) Expand the **Data Network** settings. Provide information for the IPv4, IPv6 or dual stack deployment (as per your selection) such as IP address, IP netmask, IP gateway, and virtual IP address.
c) Expand the **DNS and NTP Servers** settings. According to your deployment configuration (IPv4, IPv6 or dual stack), the fields that are displayed are different. Provide information in these three fields:

- **DNS IP Address**: The IP addresses of the DNS servers you want the Cisco Crosswork server to use. Separate multiple IP addresses with spaces.
- **DNS Search Domain**: The name of the DNS search domain.
- **NTP Servers**: The IP addresses or host names of the NTP servers you want to use. Separate multiple IPs or host names with spaces.

> **Note**
>
> The DNS and NTP servers must be reachable using the network interfaces you have mapped on the host. Otherwise, the configuration of the VM will fail.

d) Expand **Crosswork Cluster Configuration**. Provide relevant values for these fields:

- **VM Type**:
  - Choose **Hybrid** if this is one of the 3 Hybrid nodes.
  - Choose **Worker** if this is a Worker node.
- **Cluster Seed node**:
  - 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.
- **Initial hybrid node count**: Set to the default value, which is 3.
- **Initial total node count**: Set to the default value, which is 3.

e) Expand the **Deployment Credentials** settings. Enter relevant values for the VM Username and Password.

> **Note**
>
> Use a strong VM Password (8 characters long, including upper & 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.

f) The default **Disk Configuration** settings should work for most environments. Change the settings only if you are instructed to by the Cisco Customer Experience team.

g) Expand **Advanced Configuration**. Provide relevant values for these fields:

- **Timezone**: Enter the timezone details. Default value is UTC.
- **Disclaimer**: Enter your legal disclaimer text (users will refer to this text if they log into the CLI).
- **Crosswork Management Cluster Virtual IP Name**: Enter the Management Virtual IP DNS name.
- **Crosswork Data Cluster Virtual IP Name**: Enter the Data Virtual IP DNS name.
- **Enable Skip Auto Install Feature**: Any pods marked as skip auto install will not be brought up until a dependent application/pod explicitly asks for it. If left blank, the default value ("False") is selected.
- **Ignore Diagnostic Failure**: Use the default value ("False").
- **Location of VM**: Enter the location of VM.
- **K8 Orchestrator**: Enforces minimum resource reservations for the pod. If left blank, the default value ("True") is selected.
- (Optional) **Kubernetes Service Network**: The network address for the kubernetes service network. By default, the CIDR range is fixed to '10.96.0.0/16'.
- (Optional) **Kubernetes Pod Network**: The network address for the kubernetes pod network. By default, the CIDR range is fixed to '10.224.0.0/16'.
- **Installation type**:
  - *For new cluster installation*: Do not select the check box.
  - *Replacing a failed VM*: Select the check box.
  - *Adding a new worker node to an existing cluster*: Do not select the check box.
- **Auto Action Manifest Definition**: Auto action functionality enables you to customize the installation of applications along with the cluster installation. For more information, refer to Automate application installation on page 92.
  - If you plan to use the auto action functionality, enter manifest definition details. You must compress or minimize the auto action JSON file and enclose it in CDATA format. The format is **<![CDATA[{auto action json compressed content}]]>**.

    Sample auto action CDATA:

    ```
    <![CDATA[{"auto_action":{"add_to_repository_requests":[{"file_location":
    {"uri_location":{"uri":"file:///example.com/path/to/cw-na-cncadvantage-7-2-sp.tar.gz.signature"}}}],
    "install_activate_requests":[{"package_identifier":{"id":"capp-coe","version":"7-2-1-sp"}}]}}]]>
    ```

  - If you do not plan to use this functionality, leave the field blank.

- **CA Private Key**: Use the default value (Empty).
- **CA Public Key**: Use the default value (Empty).
- **Use NonDefault Calico Bgp Port**: Leave the checkbox unchecked.

7. select **Next**. The **Deploy OVF Template** window is refreshed, with **10 - Ready to Complete** highlighted.
8. Review your settings and then select **Finish** if you are ready to begin deployment. Wait for the deployment to finish before continuing. To check the deployment status:
   a) Open a VMware vCenter client.
   b) In the **Recent Tasks** tab of the host VM, view the status of the **Deploy OVF template** and **Import OVF package** jobs.
9. To finalize the template creation, select the host and right-select on the newly installed VM and select **Template** > **Convert to Template**. A prompt confirming the action is displayed. select **Yes** to confirm. The template is created under the **VMs and Templates** tab in the vSphere Client UI.

   This is the end of the first part of the manual installation workflow. In the second part, use the newly created template to build the cluster VMs.

Deploy the template

Deploys Cisco Crosswork VMs from the OVF template and describes the resource checks and follow-up actions for completing cluster deployment.

Deploy Cisco Crosswork VMs from the prepared OVF template.

Procedure

1. To build a VM, right-select on the template and select **New VM from This Template**.

   Note

   If the template is no longer present, go back and create the template. For more information, refer to Build the OVF template on page 103.

2. Set the VM name, folder, data center, and compute resource.
   a) The VMware **Deploy From Template** window appears, with the first step, highlighting the **1 - Select a name and folder** section. Enter a name and select the respective data center for the VM.

      Note

      If this is a new VM, the name must be unique and cannot be the same name as the template. If this VM is replacing an existing VM (for example, `CW-VM-0`) give the VM a unique temporary name (for example, `CW-VM-0-New`).

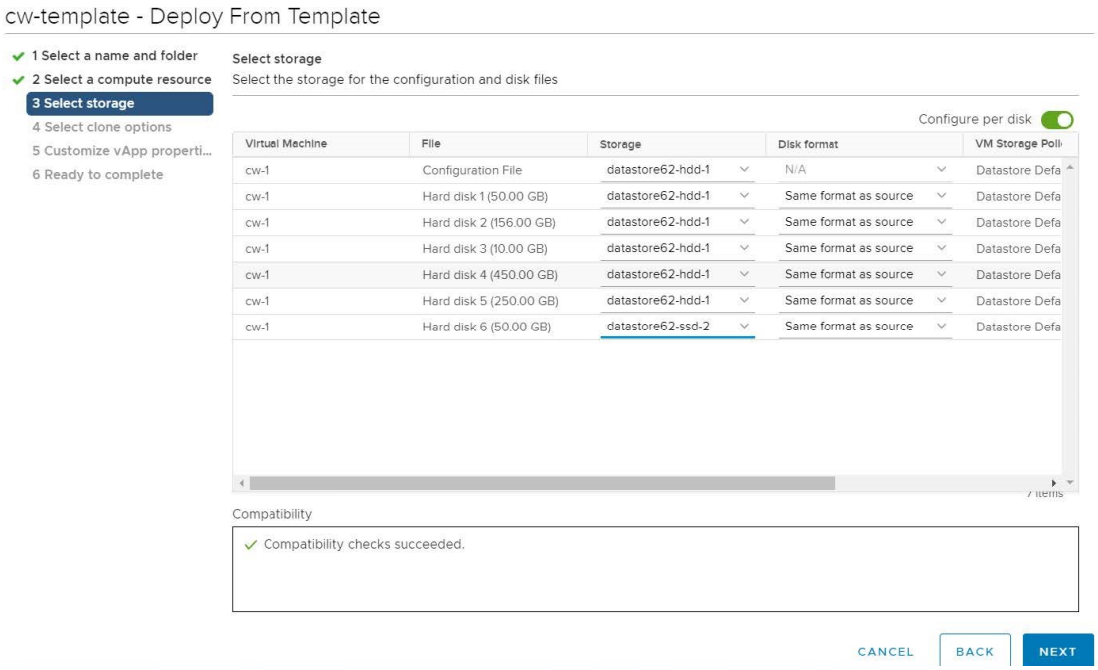
   b) select **Next**. The **Deploy From Template** window will refresh, highlighting the **2 - Select a compute resource** section. Select the host for your Cisco Crosswork VM.

3. Select storage and clone options.
   a) select **Next**. The **Deploy From Template** window will refresh, highlighting the **3 – Select Storage** section. Choose **Same format as source** option as the virtual disk format (recommended).

      The recommended configuration for the nodes uses a combination of high-speed (typically SSD based) and normal (typical disks) storage. If you are following the recommended configuration follow the steps for two data stores. Otherwise, follow the steps for using a single data store.

      - *If you are using two data stores (regular and high speed)*:
        - Enable **Configure per disk** option.
        - Select same data store (regular) as the **Storage** setting for disks 1 through 5. This data store must have 916 GB of space.
        - Select the host's high speed (ssd) data store as the **Storage** setting for disk 6. The high speed data store must have at least 50 GB of space.
        - select **Next**.

        Figure 6: Select storage - configure per disk



      - *If you are using a single data store*: Select the data store you wish to use, and select **Next**.
      - *If your data center uses shared storage*: Configure all the drives to utilize the shared storage, and select **Next**.

   b) The **Deploy From Template** window will refresh, highlighting the **4 – Select clone options** section with these checkboxes visible on the screen. Unless you have been given specific instructions to make modifications, select **Next**.

      - **Customize the operating system**: Check this box if you want to customize the operating system to avoid conflicts when deploying the VM. This step is optional.
        - If you check this box, the **Deploy From Template** window will refresh, highlighting the **Customize guest OS** section. Make the necessary changes, and select **Next**.
      - **Customize this virtual machine's hardware**: Check this box if you want to modify the IP settings or resource settings of this VM. This step is optional.

- If you check this box, the **Deploy From Template** window will refresh, highlighting the **Customize hardware** section. Make the necessary changes, and select **Next**.

- **Power on virtual machine after creation**: Leave this checkbox unselected.

4. select **Next**. The **Deploy From Template** window will refresh, highlighting the **5 – Customize vApp properties** section. The vApp properties are prepopulated with the values entered during the template creation. Some of the values will need to be updated with the appropriate values for each node being deployed.

   Note

   - It is recommended to change only the fields that are unique to each node. Leave all other fields at the default values.
   - If this VM is being deployed to replace a failed VM, or to migrate the VM to a new host, the IP and other settings must match the machine being replaced.

   - Set the node type (Hybrid/Worker).
   - **Management Network settings**: Enter correct IP values for each VM in the cluster.
   - **Data Network settings**: Enter correct IP values for each VM in the cluster.
   - **Deployment Credentials**: Enter same deployment credentials for each VM in the cluster.
   - **DNS and NTP Servers**: Enter correct values for the DNS and NTP servers.
   - **Disk Configuration**: Leave at the default settings unless directed otherwise by the Cisco Customer Experience team.
   - **Crosswork Configuration**: Enter the disclaimer message.
   - **Crosswork Cluster Configuration**:
     - **VM Type**: Select Hybrid or Worker
     - **Cluster Seed node**:
       - 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 or moving the VM to a new host.
     - **Crosswork Management Cluster Virtual IP**: The Virtual IP will remain same for each cluster node.
     - **Crosswork Data Cluster Virtual IP**: The Virtual IP will remain same for each cluster node.

5. select **Next**. The **Deploy From Template** window will refresh, highlighting the **6 – Ready to complete** section. Review your settings and then select **Finish** if you are ready to begin deployment.
6. Check VM resources and deploy the remaining VMs.
   a) For the newly created VM, confirm that the resource settings allocated for the VM match those specified in Identify the resource footprint on page 20.
   b) Repeat from **Step 1** to **Step 8** to deploy the remaining VMs in the cluster.

> **Note**
>
> When deploying the cluster for the first time, make sure the IP addresses and Seed node settings are correct. When replacing or migrating a node make sure the settings match the original VM.

7. Choose the relevant action:
   - If you are deploying a new VM, power on the VM selected as the cluster seed node. After a delay of few minutes, power on the remaining nodes. To power on, expand the host's entry, select the Cisco Crosswork VM, and then choose **Actions** > **Power** > **Power On**.
   - If this VM is replacing an existing VM, perform these:
     - Power down the existing VM.
     - Change the name of the original VM (for example, change to `CW-VM-0-Old`)
     - Change the name of the replacement VM (for example change to `CW-VM-0-New`) to match the name of the original VM (for example, `CW-VM-0`).
     - Power on the new VM. Monitor the health of the cluster using the Crosswork UI.
     - Once the cluster is healthy and stable, delete the original VM (now named as `CW-VM-0-Old`)

8. The time taken to create the cluster can vary based on the size of your deployment profile and the performance characteristics of your hardware. refer to Monitor cluster activation on page 113 to know how you can check the status of the installation.

> **Note**
>
> If you are running this procedure to replace a failed VM, then you can check the status from the Cisco Crosswork GUI (go to **Administration** > **Crosswork Manager** and select on the cluster tile to check the *Crosswork Cluster* status.

> **Note**
>
> If you are using the process to build a new Worker node, the node will automatically register itself with the existing Kubernetes cluster. For more information on how the resources are allocated to the Worker node, refer to the *Rebalance Cluster Resources* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

What to do next

Return to the installation workflow: Overview on page 34

Monitor cluster activation

Outlines how to review Cisco Crosswork cluster activation and deployment readiness, including diagnostic reports, failure-handling settings, and IOPS calculation details.

Steps to monitor deployment progress

Here are the critical steps to monitor to ensure progress is on track:

1. The installer uploads the crosswork image file to the vCenter data center.

   Note

   On running, the installer will upload the image file into the vCenter if it is not already present, and convert it into a VM template. After the installation is completed successfully, you can delete the template file from the vCenter UI (located under *VMs and Templates*) if the image is no longer needed.

2. The installer creates the VMs, and displays a success message (e.g. "Creation Complete") after each VM is created.

   Note

   For VMware deployments, this activity can also be monitored from the vSphere UI.

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 registed 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.

After the Cisco Crosswork UI becomes accessible, you can also monitor the status from the UI. For more information, refer to Log into the Cisco Crosswork UI on page 120.

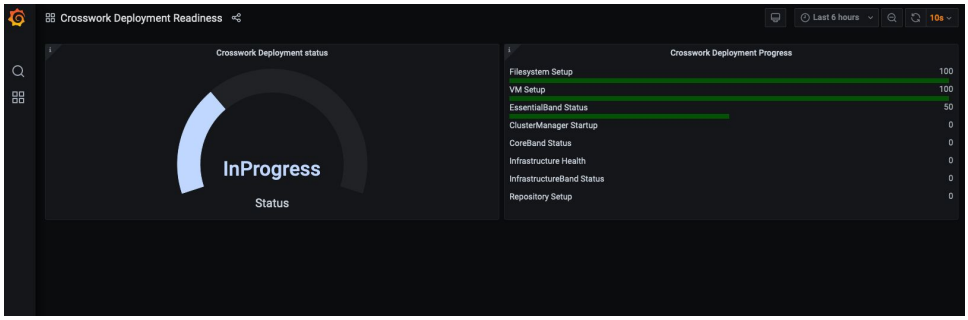
Monitor startup process

You can monitor startup progress using these 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/NK1bwVxGk/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 7: 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 refer to if the nodes are ready) and `kubectl get pods` (to refer to the list of active running pods) commands.
5. Repeat the `kubectl get pods` command until you refer to `robot-ui` in the list of active pods. For example, you may refer to `robot-ui-0` and `robot-ui-1` as active pods.
6. At this point, you can try to access the Cisco Crosswork UI.

Diagnostic assessment

During deployment, the system verifies VM datastore resource values, including disk latency, IOPs, and network bandwidth. If any value falls below the recommended threshold, the diagnostic assessment reports a failure, requiring user action to proceed with the installation. refer to Diagnostic assessment on page 115 for more information.

Deployment failure

In the event of one of the failure scenarios 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.

What to do next:

Return to the installation workflow: Overview on page 34

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.

Diagnostic assessment

Outlines deployment-time checks for datastore latency, IOPS, and bandwidth; explains how IgnoreDiagnosticsCheckFailure affects failures and shows several diagnostic report examples.

Provides the source information for this topic.

This topic explains the diagnostic checks performed during the Crosswork Network Controller deployment.

During deployment, the system checks the VM data store(s) for disk latency, IOPs, and network bandwidth (refer to VM requirements on page 21 for the recommended values for each parameter). If a fetched value is lower than the recommended value for any parameter, the diagnostic assessment reports a failure. At this point, you can choose to either ignore the report and proceed with the installation—accepting the risk of failure—or update your VM resources to meet the required criteria and retry the installation.

The outcome of the diagnostic assessment depends on the value set for the **IgnoreDiagnosticsCheckFailure** parameter:

- If set to `"false"` (default value), the installation will be blocked if the diagnostic check reports an error.
- If set to `"true"`, the diagnostic check result is ignored, and the installation will proceed.

Diagnostic failure scenario

This is a breakdown of the failure scenario for the diagnostic assessment when set to its default value (`IgnoreDiagnosticsCheckFailure` = `false`).

In this scenario, the fetched value is lower than the recommended value (such as IOPs < 4000 and/or network bandwidth < 8000 Mbps), resulting in a diagnostic assessment failure.

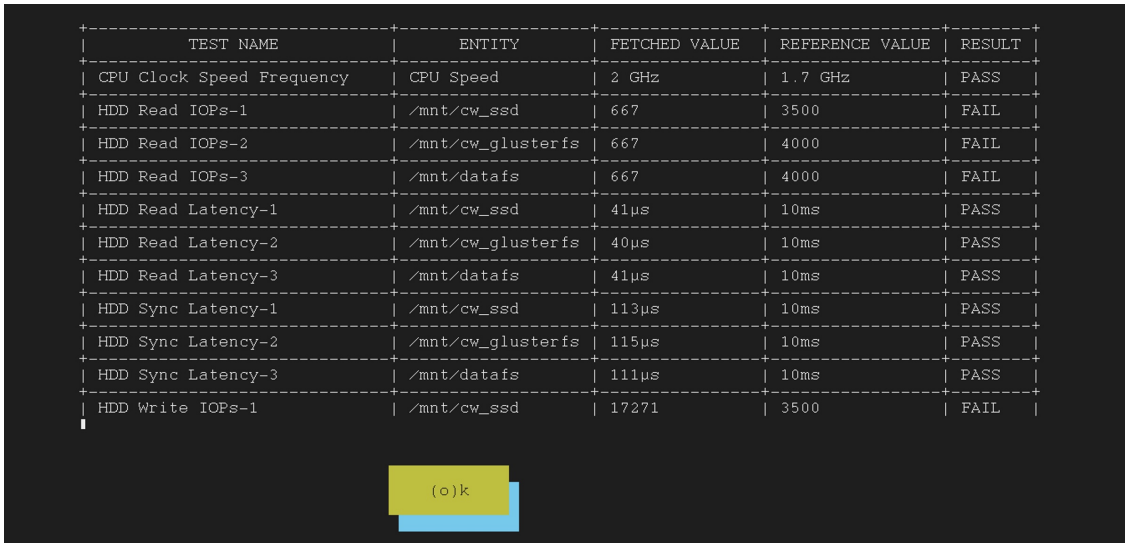
1. A banner message is displayed to notify you about the failure.

   Figure 8: Sample diagnostic failure report

2. You can use `diagnostics all` command to invoke diagnostic check, and the `diagnostics history` command to view the detailed diagnostic report.

   Figure 9: Diagnostic history output



3. To ignore the failure report and proceed with the installation, you must change the value of the **IgnoreDiagnosticsCheckFailure** parameter.

   a. Log in to the vCenter UI.

   b. Power off the VM reporting the failure. Right-select on the VM, and select **Power** > **Power Off**. select **Yes** in the confirmation popup window.

   c. select on the **Configure** tab and select **vApp Options** in the **Settings** dropdown menu.

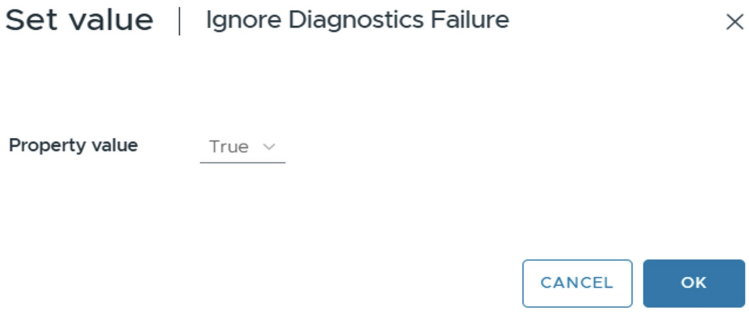
   d. Under **Properties**, select the **IgnoreDiagnosticsCheckFailure** parameter and select **Set Value**.

   Figure 10: Select parameter



   e. Set the **Property value** as **True**. select **OK** to confirm.

   Figure 11: Set value



   f. Power on the VM. Right-select on the VM, and select **Power** > **Power On**. select **Yes** in the confirmation popup window.

4. (Applicable only to Docker-based installations) If the diagnostic check reports failures for multiple nodes and the installation is blocked, you can also proceed with a full installation using the skip option (`-s`) command in the Docker installer.

   Example:

   ```
   ./cw-installer.sh install -p -m /data/<template file name> -o /data/<.ova file> -y -s
   ```

5. After `IgnoreDiagnosticsCheckFailure` is set to `True`, a banner message informs you of the decision to skip the diagnostic check failure.

   Figure 12: Skip install check banner message



> **Note**
>
> If the parameter value is lower than suboptimal (IOPs < 1000 and/or network bandwidth < 1000 Mbps), the installation will fail irrespective of your choice to ignore the diagnostic check.

Diagnostic success scenario

The diagnostic check is successful, and the installation will proceed without requiring any user action.

Backend checks

These backend checks are used to verify the resource values of the VM:

- **Directory IOPs**:

  ```
  fio --randrepeat=1  --fdatasync=1 --ioengine=sync --name=test --rw=rw --filename=<DIR>/mytest --bs=8k --size=600M --runtime=10 --time_based=1
  ```

  Replace `<DIR>` with `/mnt/cw_ssd/`, `/mnt/datafs`, or `/mnt/cw_glusterfs`, as appropriate.

- **iPerf**:

  ```
  iperf3 <IP Version> --client <VIP address> -t 5 -P 5 -p 31560 -json
  ```

  Replace `<IP Version>` with `4` (for IPv4) or `6` (for IPv6), and `<VIP address>` with the VM's VIP address.

  Ensure that the iPerf server is running on the seed node (on the VIP). To start the iPerf server on the seed node, execute `systemctl start iperf3.service`.

> **Note**
>
> - During installation, the non-seed node functions as the iPerf client and the seed node functions as the iPerf server. The iPerf test runs on the non-seed node to measure network bandwidth between the non-seed node and the seed node.
> - The iPerf check is skipped for single VM deployments.

IOPs calculation

Lists the weighted-average IOPS calculation method, the FIO commands, and sample command output used to assess storage performance in deployment diagnostics.

Provides the source information for this topic.

IOPs diagnostics use a weighted average method with **4,000 IOPs** as the reference benchmark. This calculation incorporates results from both **sync** and **libaio** engines to provide a comprehensive performance assessment.

Calculation methodology

1. **Fetch values**: Run FIO commands for both sync and libaio engines
2. **Average results**: Execute each test **2 times** and calculate the average
3. **Apply weights**:
   - 75% of sync engine value
   - 25% of libaio engine value
4. **Calculate final score**: Sum of the weighted values

Formula:

```
Final IOPs = (Sync_IOPs × 0.75) + (Libaio_IOPs × 0.25)
```

Example:

```
Sync engine result: 1,500 IOPs
Libaio engine result: 3,000 IOPs

Final Result: (1500 x 0.75) + (3000 x 0.25) = 1,875 IOPs (compared against
4,000 IOPs reference)
```

FIO command examples

1. **Sync engine test command**:

```
fio --randrepeat=1 --iodepth=64 --fdatasync=1 --ioengine=sync --name=test
--rw=rw --filename=/mnt/cw_ssd/mytest --bs=8k --size=600M --runtime=10
--time_based=1
```

Table 40: Sync engine test command

Option	Value	Description
`--randrepeat`	1	Use repeatable random I/O patterns for consistency
`--iodepth`	64	Number of I/O units to keep in flight (queue depth)
`--fdatasync`	1	Use `fdatasync()` to flush data after each write
`--ioengine`	sync	Use synchronous I/O engine (blocking operations)
`--name`	test	Job name identifier
`--rw`	rw	Mixed read/write workload pattern
`--filename`	/mnt/cw_ssd/mytest	Target file path for testing
`--bs`	8k	Block size of 8 KB per I/O operation
`--size`	600M	Total size of file to use (600 MB)
`--runtime`	10	Run test for 10 seconds
`--time_based`	1	Continue running for full runtime duration regardless of file size

2. **Libaio engine test command**:

```
fio --name=libaio --directory=/mnt/cw_ssd/mytest --ioengine=libaio --direct=1
 --sync=1 --rw=rw --bs=128k --size=512m --numjobs=4 --iodepth=4,8,16,32
--runtime=60 --time_based=1 --group_reporting=1 --norandommap=1
--output-format=json
```

Table 41: Libaio engine test command

Option	Value	Description
`--name`	libaio	Job name identifier
`--directory`	/mnt/cw_ssd/mytest	Directory path for test files
`--ioengine`	libaio	Use Linux asynchronous I/O engine (non-blocking)
`--direct`	1	Enable direct I/O (bypass page cache)
`--sync`	1	Use synchronous I/O (O_SYNC flag)
`--rw`	rw	Mixed read/write workload pattern
`--bs`	128k	Block size of 128 KB per I/O operation
`--size`	512m	Total size of file to use (512 MB)

Option	Value	Description
`--numjobs`	4	Number of parallel jobs/threads to run
`--iodepth`	4,8,16,32	Test multiple queue depths sequentially
`--runtime`	60	Run test for 60 seconds
`--time_based`	1	Continue running for full runtime duration
`--group_reporting`	1	Aggregate statistics across all jobs
`--norandommap`	1	Do not track previously accessed blocks
`--output-format`	json	Output results in JSON format for parsing

Log into the Cisco Crosswork UI

Verifies platform and cluster health by logging in to the Cisco Crosswork UI after activation, with additional guidance for manual installations.

Log in to the Cisco Crosswork UI and check cluster health after activation.

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, refer to the *Compatibility Information* section in the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes*.

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

Note

If the Cisco Crosswork UI is not accessible, during installation, please access the host's console from the data center UI to confirm if there was any problem in setting up the VM. When logging in, if you are directed to review the `firstboot.log` file, please check the file to determine the problem. If you are able to identify the error, rectify it and restart the node(s). If you require assistance, please contact the Cisco Customer Experience team.

Procedure

1. Launch one of the supported browsers.

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, refer to the *Manage Certificates* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

3. Log into the Cisco Crosswork in this list:
   a) Enter the Cisco Crosswork administrator username **admin** and the default password **admin**.
   b) select **Log In**.
   c) When prompted to change the administrator's default password, enter the new password in the fields provided and then select **OK**.

   > **Note**
   >
   > Use a strong password that is at least eight characters long and includes uppercase and lowercase letters, numbers, and at least one special character. Although allowed, avoid using passwords similar to dictionary words (for example, "Pa55w0rd!") or other easily guessable terms.

   The **Crosswork Manager** window is displayed.

4. select on the **Crosswork Health** tab, and select the **Platform Infrastructure** tile to view the health status of the microservices running on Cisco Crosswork.

5. **In case of manual installation:** 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. Then, import the file back into the Crosswork UI. For more information, refer to Import cluster inventory on page 122.

What to do next

Return to the installation workflow: Overview on page 34

Import cluster inventory

Imports a .tfvars inventory file for a manually installed VMware cluster so Cisco Crosswork can track the cluster VM and data center details.

Import the inventory file for a manually installed cluster into Cisco Crosswork.

Note

- If the manual installation was performed to replace a failed VM, you must delete the original VM after importing the cluster inventory file.
- Please uncomment or set the "*OP_Status = 2*" parameter while importing the cluster inventory file manually. If you fail to do this, the status of the VM will incorrectly appear as "Initializing" even after the VM becomes functional.
- Crosswork cannot deploy or remove VM nodes in your cluster until you complete this operation.
- You must remove the leading zeroes from the IPv6 addresses in the cluster inventory file.

Procedure

1. From the main menu, choose **Administration** > **Crosswork Manager**.
2. On the **Crosswork Summary** tab, select the **System summary** tile to display the **Cluster Management** window. Ensure the cluster is healthy.
3. Choose **Actions** > **Import inventory** to display the **Import inventory** dialog box.
4. (Optional) select **Download sample template file** to download the template. Update the file with information about the VMs in your cluster, along with the data center parameters. For more details on the installation parameters, refer to Installation parameters on page 83.
5. select **Browse** and select the cluster inventory file.
6. select **Import** to complete the operation.

Troubleshoot the cluster

Lists common cluster installation and operation scenarios with their possible resolutions, including installer, network, and platform issues. It also includes VMware-specific troubleshooting scenarios.

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.

What to do next:

Return to the installation workflow: Overview on page 34

Table 42: General scenarios

Scenario	Possible Resolution
Certificate Error	The ESXi 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.
Image upload takes a long time or upload is interrupted.	The image upload duration depends on the link and datastore performance and can be expected to take around 10 minutes or more. If an upload is interrupted, the user needs to manually remove the partially uploaded image file from vCenter via the vSphere UI.
vCenter authorization	The vCenter user needs to have authorization to perform the actions as described in Installation Prerequisites on page 43.
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.
Crosswork VM is not allowing the admin user to log in OR these error is displayed: *Error: Invalid value for variable on cluster_vars.tf line 113:* *├────────────────* *This was checked by the validation rule at cluster_vars.tf:115,3-13. Error: expected length of name to be in the range (1 - 80), got with data.vsphere_virtual_machine.template_from_ovf, on main.tf line 32, in data "vsphere_virtual_machine" "template_from_ovf": 32: name = var.Cw_VM_Image Mon Aug 21 18:52:47 UTC 2023: ERROR: Installation failed. Check installer and the VMs' log by accessing via console and viewing /var/log/firstBoot.log*	This happens when the password is not complex enough. Create a strong password, update the configuration manifest and redeploy. Use a strong VM Password (8 characters long, including upper & 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.
Deployment fails with: *Failed to validate Crosswork cluster initialization.*	The clusters' seed VM is either unreachable or one or more of the cluster VMs have failed to get properly configured. 1. Check whether the VM is reachable, and collect logs from `/var/log/firstBoot.log` and `/var/log/vm_setup.log` 2. Check the status of the other cluster nodes.

Scenario	Possible Resolution

The VMs are deployed but the Crosswork cluster is not being formed.	A successful deployment allows the operator logging in to the VIP or any cluster hybrid node IP address to run these command to get the status of the cluster:
	`sudo kubectl get nodes`
	Note
	This command will not work on worker nodes.
	A sample output for a 3-node cluster is:
	`NAME STATUS ROLES AGE VERSION`
	`172-25-87-2-hybrid.cisco.com Ready master 17d v1.33.0`
	`172-25-87-3-hybrid.cisco.com Ready master 17d v1.33.0`
	`172-25-87-4-hybrid.cisco.com Ready master 17d v1.33.0`
	In case of a different output, collect these logs: `/var/log/firstBoot.log` and `/var/log/vm_setup.log`
	In addition, for any cluster nodes not displaying the Ready state, collect:
	`sudo kubectl describe node <name of node>`

these error is displayed while uploading the image: *govc: The provided network mapping between OVF networks and the system network is not supported by any host.*	The vSphere Distributed Switch (DSwitch) on the vCenter is misconfigured. Please check whether it is operational and mapped to the ESXi hosts.
VMs deploy but install fails with *Error: timeout waiting for an available IP address*	Most likely cause would be an issue in the VM parameters provided or network reachability. Enter the VM host through the vCenter console. and review and collect these logs: `/var/log/firstBoot.log` and `/var/log/vm_setup.log`
When deploying on a vCenter, these error is displayed towards the end of the VM bringup: Error processing disk changes post-clone: *disk.0: ServerFaultCode: NoPermission: RESOURCE (vm-14501:2000), ACTION (queryAssociatedProfile): RESOURCE (vm-14501), ACTION (PolicyIDByVirtualDisk)*	Enable Profile-driven storage. Query permissions for the vCenter user at the root level (i.e. for all resources) of the vCenter.

Scenario	Possible Resolution
On running or cleaning, installer reports *Error: cannot locate virtual machine with UUID "xxxxxxx": virtual machine with UUID "xxxxxxxx" not found*	The installer uses the `tfstate` file stored as `/data/crosswork-cluster.tfstate` to maintain the state of the VMs it has operated upon. If a VM is removed outside of the installer, that is through the vCenter UI, this state is out of synchronization. To resolve, remove the `/data/crosswork-cluster.tfstate` file.
In a cluster with five or more nodes, the databases move to hybrid nodes during a node reload scenario. Users will refer to these alarm: *"The robot-postgres/cw-timeseries-db pods are currently running on hybrid nodes. Please relocate them to worker nodes if they are available and health."*	To resolve the alarm, invoke the move API to move the databases to worker nodes. Use these script to place services. It returns a job ID that can be queried to ensure the job is completed. For more information, refer to these API commands and responses: • Placement Services • Get Jobs • Get Events

Table 43: Installer tool scenarios

Scenario	Possible Resolution
Missing or invalid parameters	The installer provides a clue as regards to the issue; however, in case of errors in the manfiest file HCL syntax, these can be misguiding. If you refer to "Type errors", check the formatting of the configuration manifest. The manifest file can also be passed as a simple JSON file. Use any appropriate HCL to JSON conversion tool to validate/convert.
Error conditions such as: *Error: Error locking state: Error acquiring the state lock: resource temporarily unavailable* *Error: error fetching virtual machine: vm not found* *Error: Invalid index*	These errors are common when re-running the installer after an initial run is interrupted (Control C, or TCP timeout, etc). Remediation steps are: 1. Run the clean operation (`./cw-installer.sh clean -m <your manifest here>`) OR remove the VM files manually from the vCenter. 2. Remove the state file (`rm /data/crosswork-cluster.tfstate`). 3. Retry the installation (`./cw-installer.sh clean -m <your manifest here>`).
The VMs take a long time to deploy	The time needed to clone the VMs during the installation will be determined by the workload on the disk drives used by the host machines. Running the install serially (without the [-p] flag) will lessen this load while increasing the time needed to deploy the VMs.

Scenario	Possible Resolution
Installer reports plan to add more resources than the current number of VMs	Other than the Crosswork cluster VMs, the installer tracks other meta-resources. Thus, when doing an installation of, say a 3-VM cluster, the installer may report a "plan" to add more resources than the number of VMs.
On running or cleaning, installer reports *Error: cannot locate virtual machine with UUID "xxxxxxx": virtual machine with UUID "xxxxxxxx" not found*	To resolve, remove the `/data/crosswork-cluster.tfstate` file. The installer uses the `tfstate` file stored as `/data/crosswork-cluster.tfstate` to maintain the state of the VMs it has operated upon. If a VM is removed outside of the installer, that is through the vCenter UI, this state is out of synchronization.
Encountered one of these errors during execution: Error 1: *% docker run --rm -it -v `pwd`:/data a45 docker: invalid reference format: repository name must be lowercase. refer to 'docker run --help'* Error 2: *docker: Error response from daemon: Mounts denied: approving /Users/Desktop: file does not exist ERRO[0000] error waiting for container: context canceled*	Move the files to a directory where the path is in lowercase (all lowercase, no spaces or other special characters). Then navigate to that directory and rerun the installer.

Table 44: Dual stack scenarios

Scenario	Possible Resolution
During deployment, these error message is displayed: *ERROR: No valid IPv6 address detected for IPv6 deployment.*	If you intend to use a dual stack configuration for your deployment, make sure that the host machine running the Docker installer meets these requirements: • 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 `--network host` when running the Docker installer. • Confirm that the provided IPv6 network CIDR and gateway are valid and reachable.
During deployment, these error message is displayed: *ERROR: seed v4 host empty*	Ensure you use the approved version of Docker installer (19 or higher) to run the deployment.

Scenario	Possible Resolution
During deployment, these error message is displayed: *ERROR: Installation failed. Check installer and the VMs' log by accessing via console and viewing /var/log/firstBoot.log*	Common reasons for failed installation are: • Incorrect IPv4 or IPv6 Gateway IP for either Management or Data interfaces. • Unreachable IPv4 or IPv6 Gateway IP for either Management or Data interfaces. • Errors in mapping the vCenter networks in MgmtNetworkName and DataNetworkName parameters in .tfvars file Check the `firstBoot.log` file for more information, and contact Cisco Customer Experience team for any assistance.

6 Install Cisco Crosswork Data Gateway on VMware vCenter

Topics:

- Crosswork Data Gateway workflows
- Crosswork Data Gateway parameters and deployment scenarios
- Install Crosswork Data Gateway using vCenter vSphere Client
- Install Crosswork Data Gateway via the OVF Tool
- Log in and Log out of the Data Gateway VM
- Crosswork Data Gateway authentication and enrollment
- Crosswork Data Gateway post-installation tasks
- Troubleshoot installation and enrollment

Summarizes Crosswork Data Gateway deployment choices, installation procedures, enrollment, post-installation configuration, and troubleshooting.

This chapter covers VMware-based Crosswork Data Gateway installation, from choosing a deployment workflow through enrollment and troubleshooting.

Crosswork Data Gateway workflows

Provides installation and upgrade workflows for Crosswork Data Gateway, including deployment choices, verification checkpoints, and links to related procedures. It also outlines upgrade steps and post-installation checkpoints.

Use these workflows to select an installation or upgrade path for Crosswork Data Gateway.

Crosswork Data Gateway is installed as a base VM that contains only enough software to register itself with Crosswork Network Controller.

Installation workflow

Follow these steps to install a Data Gateway VM for use with Crosswork Network Controller:

1. Choose the deployment profile for the Crosswork Data Gateway VM. refer to Crosswork Data Gateway VM requirements on page 46.
2. Review the installation parameters at Crosswork Data Gateway Parameters and Deployment Scenarios and make sure that you have all the required information to install Data Gateway using the preferred deployment scenario.
3. Install Data Gateway using a preferred method:

   Table 45: Crosswork Data Gateway installation options

VMware	Install Crosswork Data Gateway using vCenter vSphere Client on page 144 Install Crosswork Data Gateway via the OVF Tool on page 154

   Note

   If you plan to install multiple Crosswork Data Gateway VMs due to load or scale requirements or you wish to leverage Crosswork Data Gateway High Availability, we recommend that you install all the Data Gateway VMs first and then proceed with adding them to a Data Gateway pool.

4. Complete the post-installation tasks mentioned in the section Crosswork Data Gateway post-installation tasks on page 159.
5. Verify that the Data Gateway VM has enrolled successfully with Crosswork Network Controller. For information on how to verify the enrollment process, refer to Crosswork Data Gateway authentication and enrollment on page 158.

   After the Data Gateway VM has enrolled successfully with Crosswork Network Controller, set up the data collection by creating a Data Gateway pool. For more information, refer to the *Create a Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Upgrade workflow

The upgrade procedure is required only for the Data Gateway Base VM upgrade. Upgrade of other components, such as collectors, is performed by Crosswork Network Controller.

Follow these steps to upgrade a Data Gateway VM for use with Crosswork Network Controller:

1. Shut down and delete the existing Data Gateway VMs. For detailed instructions, refer to Delete Crosswork Data Gateway VM from Cisco Crosswork.

2. Install new Data Gateway VMs with the same number of VMs as the old configuration. Also, configure the new VMs with identical settings, including the Management interface. For detailed installation instructions, refer to Installation workflow on page 130.
3. Verify that the newly installed Data Gateway VMs are enrolled with Crosswork Network Controller.

   On the **Data Gateway Management** page, confirm that Data Gateway status has the status as:

   - **Admin State** is in the **Up** state.
   - **Operational State** is in the **Not Ready** state.

   For information on how to verify the enrollment process, refer to Crosswork Data Gateway authentication and enrollment on page 158.
4. Complete all required post-installation tasks, such as configuring network settings. For more information, refer to Crosswork Data Gateway post-installation tasks on page 159.
5. Verify that devices are attached to the data gateways in the Crosswork Network Controller UI.

   For detailed information on upgrading the Data Gateway VM, refer to Upgrade Cisco Crosswork Network Controller on page 369.

Crosswork Data Gateway parameters and deployment scenarios

Lists Crosswork Data Gateway account details, deployment parameters, supported scenarios, and configuration values for VMware deployments. It explains required and optional settings, defaults, user roles, and scenario-specific information.

Review these deployment parameters and scenarios when configuring Crosswork Data Gateway on VMware.

Before you begin installing Crosswork Data Gateway, read through this section to understand the deployment parameters and possible deployment scenarios.

During installation, Crosswork Data Gateway creates these user accounts:

- Crosswork Data Gateway administrator, with the username, dg-admin, and the password set during installation. The administrator uses this user ID to log in and troubleshoot the Data Gateway. If the password is forgotten, the Data Gateway must be reinstalled to restore access, as it cannot be recovered or reset.
- Crosswork Data Gateway operator, with the username, dg-oper and the password set during installation. The dg-oper user has permissions to perform all 'read' operations and limited 'action' commands.

  To know what operations an admin and operator can perform, refer to the *Supported user roles* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
- Crosswork Data Gateway technical assistance center, with the username, dg-tac. The password for this user is set when one of the other users of the Data Gateways enables this account.

The **dg-admin**, **dg-oper**, and **dg-tac** user accounts are reserved usernames and cannot be changed. You can change the passwords for the **dg-admin** and **dg-oper** accounts through the Crosswork Data Gateway user interface. The **dg-tac** account is inactive by default. It can be temporarily activated by an authorized user who enables the account and sets a password.

Note

Activating the **dg-tac** account requires a token that only Cisco can provide. If CLI access to the Data Gateway system is required for troubleshooting, Cisco Customer Experience (CX) will assist you in enabling the **dg-tac** account.

For more information, refer to the *Change passphrase* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. In case of lost or forgotten passwords, destroy the current VM, you have to create a new VM, and re-enroll the new VM with Cisco Crosswork.

these table provides the label and key values of deployment parameters. Labels represent the parameters that can be configured in the VMware UI and Keys corresponds to field values in the OVF script that match your configuration.

In these table:

***** Denotes the mandatory parameters. Parameters without this mark are optional. You can choose them based on your deployment scenario. Deployment scenarios are explained (wherever applicable) in the **Additional Information** column.

****** Denotes parameters that you can enter during install or address later using additional procedures.

Note

When entering the parameters for deployment, ensure that you add the correct parameters. If the parameter values are incorrect, you have to destroy the current Data Gateway VM, create a new VM, and re-enroll the new VM with Cisco Crosswork.

Table 46: Crosswork Data Gateway Deployment Parameters and Scenarios

Label	Key	Description	Additional information
DeploymentOption	`DeploymentOption`	Parameter conveys the VM resource profile. For an on-premise installation, choose either: • onpremise-standard • onpremise-extended	This parameter is applicable only for the OVF tool.
Host Information			
Hostname*	`Hostname`	Name of the Data Gateway VM specified as a fully qualified domain name (FQDN). In larger systems, you are likely to have more than one Data Gateway VM. The hostname must, therefore, be unique and created in a way that makes identifying a specific VM easy.	

Label	Key	Description	Additional information
Description*	`Description`	A detailed description of the Data Gateway.	
Crosswork Data Gateway Label	`Label`	Label used by Cisco Crosswork to categorize and group multiple Data Gateway VMs.	
Allow Usable RFC 8190 Addresses*	`AllowRFC8190`	Choose how to validate interface addresses that fall in a usable RFC 8190 range. Options are: `Yes`, `No`, or `Ask`, where the initial configuration script prompts for confirmation. The default value is `Yes`.	The default value is `Yes` to automatically allow interface addresses in an RFC 8190 range.
Crosswork Data Gateway Private Key URI	`DGCertKey`	SCP URI to private key file for session key signing. You can retrieve this using SCP (`user@host:path/to/file`).	Cisco Crosswork uses self-signed certificates for handshake with Crosswork Data Gateway. These certificates are generated at installation. However, if you want to use third party or your own certificate files, then enter these parameters. Certificate chains override any preset or generated certificates in the Data Gateway VM and are given as an SCP URI (user:host:/path/to/file). The host with the URI files must be reachable on the network (from the vNIC0 interface via SCP) and the files must be present at the time of install.
Crosswork Data Gateway Certificate File and Key Passphrase	`DGCertChainPwd`	Passphrase of the SCP user to retrieve the Crosswork Data Gateway PEM formatted certificate file and private key.	
Data Disk Size	`DGAppdataDisk`	Indicates the size in GB of a second data disk. The default value of this parameter in each profile is: • 20 GB for Standard. • 520 GB for Extended. Do not change the default value without consulting a Cisco representative.	

Label	Key	Description	Additional information
High Availability Network Mode*	`HANetworkMode`	Indicates the mode for the high-availability network. Options are: • L2 • L3 The default value is `L2`.	When deploying on VMware, set the network to L2.
Passphrase			
dg-admin Passphrase*	`dg-adminPassword`	The password you have chosen for the dg-admin user. Password must be 8-64 characters.	
dg-oper Passphrase*	`dg-operPassword`	The password you have chosen for the dg-oper user. Password must be 8-64 characters.	

Interfaces

In a 3-NIC deployment, you must provide an IP address for Management Traffic (vNIC0) and Control or Data Traffic (vNIC1). The IP address for Device Access Traffic (vNIC2) is assigned during Data Gateway pool creation as explained in the *Create a Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Note

- Selecting **None** in both the IPv4 Method and the IPv6 Method fields of the vNIC results in a nonfunctional deployment.
- Although the IP address is assigned when the Data Gateway is added to a pool, you must configure network connectivity for the third interface in VMware either during installation or before adding the Data Gateway to the pool.

vNIC Role Assignment

Role assignment allows you to control the traffic that an interface must handle. If your use case is not supported by the preassigned roles, you can manually assign roles to the interfaces.

Each parameter has a predefined role. The parameter accepts the interface value as eth0, eth1, or eth2, which refer to vNIC0, vNIC1, or vNIC2 respectively.

Label	Key	Description	Additional information
Default Gateway*	`NicDefaultGateway`	The interface used as the Default Gateway for processing the DNS and NTP traffic. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	For information on the type of roles that you must assign to the vNICs, refer to Crosswork Data Gateway default vNIC deployment modes.
Administration*	`NicAdministration`	The interface used to access the VM through the SSH access. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
External Logging*	`NicExternalLogging`	The interface used to send logs to an external logging server. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
Management*	`NicManagement`	The interface used to send the enrollment and other management traffic. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
Control*	`NicControl`	The interface used to send the destination, device, and collection configuration. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth1`.	
Northbound System Data*	`NicNBSystemData`	The interface used to send collection data to the system destination. As the system destinations share the same IP as the interface that allows connection to the collection service, the northbound data for system destinations uses the Control role's interface. Options are `eth0`, `eth1`, `eth2` or `eth3`.	
Northbound External Data*	`NicNBExternalData`		

Label	Key	Description	Additional information
		The interface used to send the collection data to the external destinations configured by the user. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth1`.	
Southbound Data*	`NicSBData`	The interface used to collect data from the devices. If the interface only has the NicSBData role, it does not need an IP during the deployment. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth2`.	

Label	Key	Description	Additional information
vNIC IPv4 Address (vNIC0, vNIC1, and vNIC2 based on the number of interfaces you choose to use)			
Note When using two or three NICs both vNIC0 and vNIC1 must be assigned static IPv4 or IPv6 addresses. All unused vNICs (IPv4 or IPv6) should be left set to Method "None" with the other fields left at the default.			

vNIC IPv4 Method*	`Vnic0IPv4Method`
	`Vnic1IPv4Method`
	`Vnic2IPv4Method`
		Method in which the interface is assigned an IPv4 address - `None` or `Static`.
		The default value is `None`.
			• If you are using IPv4, change the value from **none** to **static**, then configure these fields:
			  • vNIC IPv4 Address
			  • vNIC IPv4 Netmask
			  • vNIC IPv4 Skip Gateway
			  • vNIC IPv4 Gateway
			• If you are using IPv6, leave the value set to **none** and retain the default IPv4 settings.

vNIC IPv4 Address	`Vnic0IPv4Address`
	`Vnic1IPv4Address`
	`Vnic2IPv4Address`
		IPv4 address of the interface.	

vNIC IPv4 Netmask	`Vnic0IPv4Netmask`
	`Vnic1IPv4Netmask`
	`Vnic2IPv4Netmask`
		IPv4 netmask of the interface in dotted quad format.	

vNIC IPv4 Skip Gateway	`Vnic0IPv4SkipGateway`
	`Vnic1IPv4SkipGateway`
	`Vnic2IPv4SkipGateway`
		The default value is `False`.
		Setting this to `True` skips configuring a gateway.	

vNIC IPv4 Gateway	`Vnic0IPv4Gateway`
	`Vnic1IPv4Gateway`
	`Vnic2IPv4Gateway`
		IPv4 address of the vNIC gateway.	

vNIC IPv6 Address (vNIC0, vNIC1, and vNIC2 based on the number of interfaces you choose to use)			

Label	Key	Description	Additional information

vNIC IPv6 Method*	`Vnic0IPv6Method`
	`Vnic1IPv6Method`
	`Vnic2IPv6Method`
		Method in which the vNIC interface is assigned an IPv6 address - `None`, `Static`, or `SLAAC`.
		The default value is `None`.
			If you are using IPv6, change the value from **none** to **static**, then configure these fields:
			• vNIC IPv6 Address
			• vNIC IPv6 Netmask
			• vNIC IPv6 Skip Gateway
			• vNIC IPv6 Gateway

vNIC IPv6 Address	`Vnic0IPv6Address`
	`Vnic1IPv6Address`
	`Vnic2IPv6Address`
		IPv6 address of the interface.	

vNIC IPv6 Netmask	`Vnic0IPv6Netmask`
	`Vnic1IPv6Netmask`
	`Vnic2IPv6Netmask`
		IPv6 prefix of the interface.	

vNIC IPv6 Skip Gateway	`Vnic0IPv6SkipGateway`
	`Vnic1IPv6SkipGateway`
	`Vnic2IPv6SkipGateway`
		Options are `True` or `False`.
		Selecting `True` skips configuring a gateway.	

vNIC IPv6 Gateway	`Vnic0IPv6Gateway`
	`Vnic1IPv6Gateway`
	`Vnic2IPv6Gateway`
		IPv6 address of the vNIC gateway.	

DNS Servers			
DNS Address*	`DNS`	Space delimited list of IPv4 or IPv6 addresses of the DNS servers accessible from the management interface.	
DNS Search Domain*	`Domain`	DNS search domain. The default value is `localdomain`.	
DNS Security Extensions*	`DNSSEC`	Options are `False`, `True`, or `Allow-Downgrade`. The default value is `False` Select `True` to use DNS security extensions.	
DNS over TLS*	`DNSTLS`	Options are `False`, `True`, and `Opportunistic`. The default value is `False`. Select `True` to use DNS over TLS.	

Label	Key	Description	Additional information
Multicast DNS*	`mDNS`	Options are `False`, `True`, and `Resolve`. Select `True` to use multicast DNS. The default value is `False`.	If you choose Resolve, only resolution support is enabled. Responding is disabled.
Link-Local Multicast Name Resolution*	`LLMNR`	Options are `False`, `True`, `Opportunistic`, or `Resolve`. The default value is `False`.	If you choose Resolve, only resolution support is enabled. Responding is disabled. Select `True` to use link-local multicast name resolution.
NTPv4 Servers			
NTPv4 Servers*	`NTP`	Space-delimited list of IPv4, IPv6 addresses, or hostnames of the NTPv4 servers accessible in the management interface.	You must enter value here, such as pool.ntp.org. NTP server is critical for time synchronization between the Data Gateway VM, Crosswork, and devices. Using a nonfunctional or dummy address may cause issues when Cisco Crosswork and Crosswork Data Gateway try to communicate with each other. If you are not using an NTP server, ensure that time gap between Crosswork Data Gateway and Crosswork is not more than 10 minutes. Else, Data Gateway fails to connect.
NTPv4 Authentication	`NTPAuth`	Select `True` to use NTPv4 authentication. The default value is `False`.	
NTPv4 Keys	`NTPKey`	Key IDs to map to the server list. Enter a space-delimited list of Key IDs.	
NTPv4 Key File URI	`NTPKeyFile`	SCP URI to the chrony key file.	
NTPv4 Key File Passphrase	`NTPKeyFilePwd`	Password of SCP URI to the chrony key file.	
Remote Syslog Server			

Label	Key	Description	Additional information
Remote Syslog Server*	`UseRemoteSyslog`	Options are `True` and `False`. Select `True` to send Syslog messages to a remote host. The default value is `False`.	Configuring an external syslog server sends service events (CLI/MDT/SNMP/gNMI) to the external syslog server. Otherwise, they are logged only to the Crosswork Data Gateway VM. If you want to use an external syslog server, specify these settings: • Use Remote Syslog Server • Syslog Server Address • Syslog Server Port • Syslog Server Protocol
Syslog Server Addresses	`SyslogAddress`	Hostname, IPv4, or IPv6 address of a syslog server accessible in the management interface.	
Syslog Server Port	`SyslogPort`	Port number of the syslog server. The default port number is 514.	
Syslog Server Protocol	`SyslogProtocol`	Options are `UDP`, `RELP`, or `TCP` to send the syslog. The default value is `UDP`.	
Syslog Multiserver Mode	`SyslogMultiserverMode`	Multiple servers in the failover or simultaneous mode. This parameter is applicable only when the protocol is set to a non-UDP value. UDP must use the simultaneous mode. Options are `Simultaneous` or `Failover`. The default value is `Simultaneous`.	
Syslog over TLS	`SyslogTLS`	Select `True` to use TLS to encrypt syslog traffic. The default value is `False`.	
Syslog TLS Peer Name	`SyslogPeerName`	Syslog server hostname exactly as entered in the server certificate SubjectAltName or subject common name.	
Syslog Root Certificate File URI	`SyslogCertChain`		

Label	Key	Description	Additional information
		PEM formatted root cert of syslog server retrieved using SCP. The host with the URI files must be reachable on the network (from vNIC0 interface via SCP) and the files must be present at the time of install.	
Syslog Certificate File Passphrase	`SyslogCertChainPwd`	Password of SCP user to retrieve Syslog certificate chain.	
Remote Auditd Server			
Remote auditd Server*	`UseRemoteAuditd`	Options are `True` and `False`. The default value is `False`. Select `True` to send auditd messages to a remote host.	If desired, you can configure an external Auditd server. Crosswork Data Gateway sends audit notifications to the Auditd server when configured and present on the network. Specify these three settings to use an external Auditd server.
Auditd Server Address	`AuditdAddress`	Hostname, IPv4, or IPv6 address of an optional Auditd server.	
Auditd Server Port	`AuditdPort`	Port number of an optional Auditd server. The default port is 60.	
Controller and Proxy Settings			
Crosswork Controller IP*	`ControllerIP`	The Virtual IP address or the hostname of the Cisco Crosswork cluster. **Note** If you are using an IPv6 address, it must be surrounded by square brackets ([1::1]). For `ControllerIP` in geo redundancy enabled environment, refer to Geo redundancy parameters.	This is required so that the Data Gateway can enroll with the Crosswork server during the installation and initial start up. Excluding this step requires you to manually ingest the certificate. For more information, refer to Import controller signing certificate file.
Crosswork Controller Port*	`ControllerPort`	Port of the Cisco Crosswork controller. The default port is 30607.	

Label	Key	Description	Additional information
Controller Signing Certificate File URI*	`ControllerSignCertChain`	PEM formatted root cert of Cisco Crosswork to validate signing certs retrieved using SCP. Cisco Crosswork generates the PEM file and is available at [illegible]. The certificate is in the format. [illegible] **Note** If you are using an IPv6 address, it must be surrounded by square brackets ([1::1]). For `ControllerSignCertChain` in geo redundancy enabled environment, refer to Geo redundancy parameters.	Crosswork Data Gateway requires the Controller Signing Certificate File to enroll automatically with Cisco Crosswork. If you specify these parameters during the installation, the certificate file is imported once Data Gateway boots up for the first time. If you do not specify these parameters during installation, then import the certificate file manually by following the procedure Import controller signing certificate file.
Controller SSL/TLS Certificate File URI	`ControllerTlsCertChain`	Cisco Crosswork Controller PEM formatted SSL/TLS certificate file retrieved using SCP.	
Controller Certificate File Passphrase*	`ControllerCertChainPwd`	Password of SCP user (cw-admin) to retrieve Cisco Crosswork certificate chain.	

Label	Key	Description	Additional information
Proxy Server URL	`ProxyURL`	URL of the HTTP proxy server.	The proxy parameters apply to the Crosswork Data Gateway cloud deployment. The Data Gateway must connect to the Internet via TLS, and a proxy server may be required if it is not present in your environment. If you want to use a proxy server, specify these parameters.
Proxy Server Bypass List	`ProxyBypass`	Comma-delimited list of addresses and hostnames that will not use the proxy server.	
Authenticated Proxy Username	`ProxyUsername`	Username for authenticated proxy servers.	
Authenticated Proxy Passphrase	`ProxyPassphrase`	Passphrase for authenticated proxy servers.	
HTTPS Proxy SSL/TLS Certificate File URI	`ProxyCertChain`	HTTPS proxy PEM formatted SSL/TLS certificate file retrieved using SCP.	
HTTPS Proxy SSL/TLS Certificate File Passphrase	`ProxyCertChainPwd`	Password of SCP user to retrieve proxy certificate chain.	

Table 47: Geo redundancy parameters

Label	Key	Description	Additional information
Crosswork Controller IP*	`ControllerIP`	Use FQDN as the unified endpoint. For more information, refer to Unified endpoint requirements.	
Controller Signing Certificate File URI*	`ControllerSignCertChain`	PEM formatted root certificate of Cisco Crosswork to validate signing certs retrieved using SCP. Cisco Crosswork generates the PEM file and is available at `admin@<Crosswork_VM_` [illegible] . The certificate is in the [illegible] format.	Crosswork Data Gateway requires the Controller Signing Certificate File to enroll automatically with Cisco Crosswork. If you specify these parameters during the installation, the certificate file is imported once Data Gateway boots up for the first time. If you do not specify these parameters during installation, then import the certificate file manually by following the procedure Import controller signing certificate file.

Install Crosswork Data Gateway using vCenter vSphere Client

Deploy the Crosswork Data Gateway OVA in vCenter, configure its deployment settings, and power on the virtual machine. The procedure covers profile selection, resource allocation, network selection, and template review.

Use this procedure to install Crosswork Data Gateway using vCenter vSphere Client.

Follow these steps to install Crosswork Data Gateway using vCenter vSphere Client:

Note

- We have included sample images of Crosswork Data Gateway on-premise Standard deployment in the procedure.
- Depending on the vCenter version, the layout and order of fields on the vCenter vSphere screens can vary.
- Leave values blank or use the default unless they are required by the installer (indicated with an asterisk *) or specific to your environment, such as certificates.

Before you begin

Import the required CA certificate chain into the VMware trust store before you deploy the signed OVA image from the VMware vSphere UI. For more information, refer to Import the CA certificate chain on VMware on page 102.

Warning

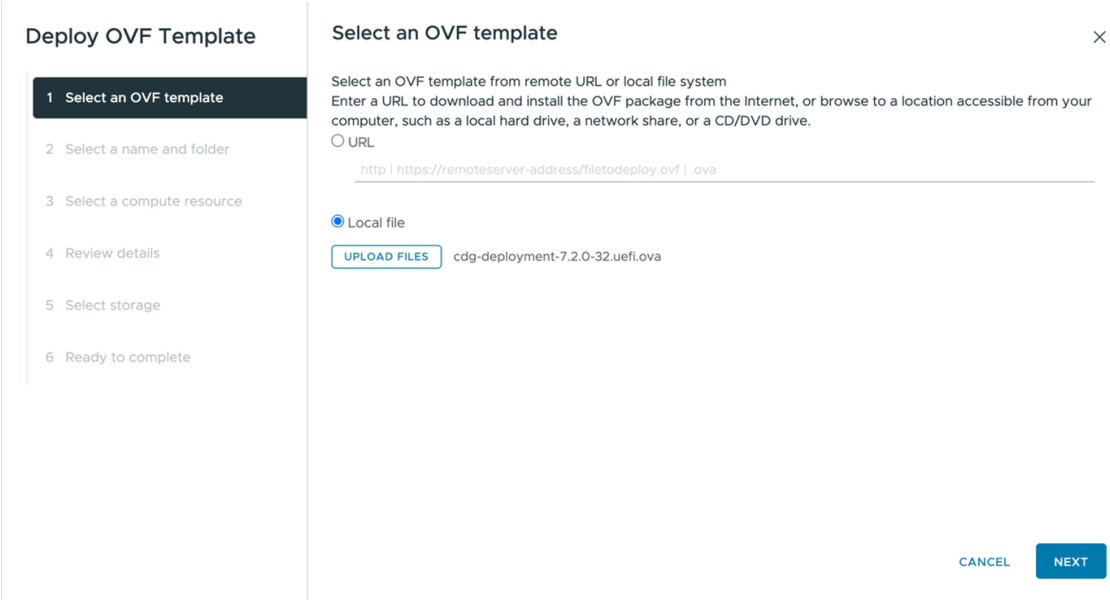
The default VMware vCenter deployment timeout is 15 minutes. If the time taken to fill the OVF template exceeds 15 minutes, vCenter times out and you have to start over again. To prevent this behavior, it is recommended that you plan for the installation by having the necessary parameters and requirements ready. Refer to the Table 46: Crosswork Data Gateway Deployment Parameters and Scenarios on page 132 for list of mandatory and optional parameters.

Procedure

1. Prepare the OVA deployment.
   a) Download the Crosswork Data Gateway UEFI OVA image file from cisco.com (*.ova). For the 7.2.x release, download **cdg-deployment-7.2.1.uefi.ova**.
   b) Connect to vCenter vSphere Client and select **Actions** > **Deploy OVF Template**.
2. The VMware **Deploy OVF Template** wizard appears and highlights the first step, **1 Select template**.
   a) Click **Browse** to navigate to the location where you downloaded the OVA image file and select it.

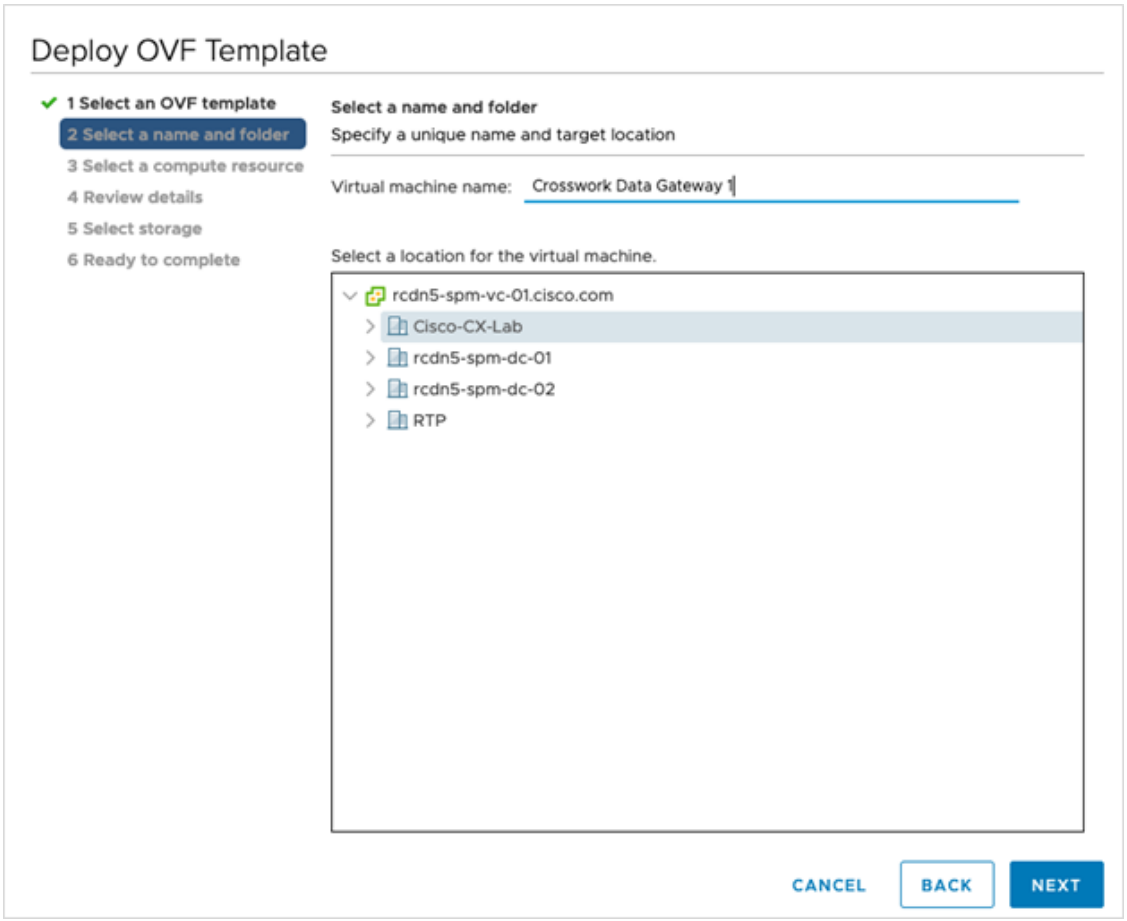
      Once selected, the filename is displayed in the window.

Figure 13: OVF template selection



3. Click **Next** to go to **2 Select a name and folder**, as shown in these figure.
   a) Enter a unique name for the VM that you are creating.
   b) In the **Select a location for the virtual machine** list, choose the data center under which the VM resides.

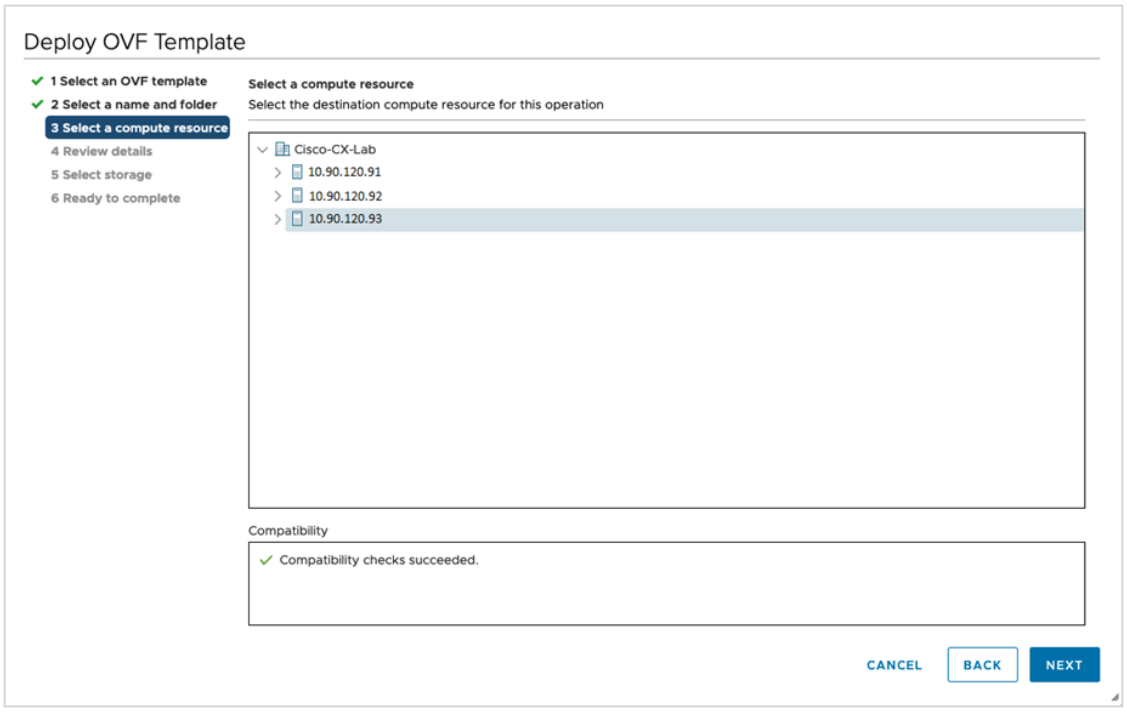
Figure 14: Name and folder selection



4. Choose resources and accept the license terms.

a) Click **Next** to go to **3 Select a computer resource**. Choose the VM's host or cluster.

Figure 15: Computer resource

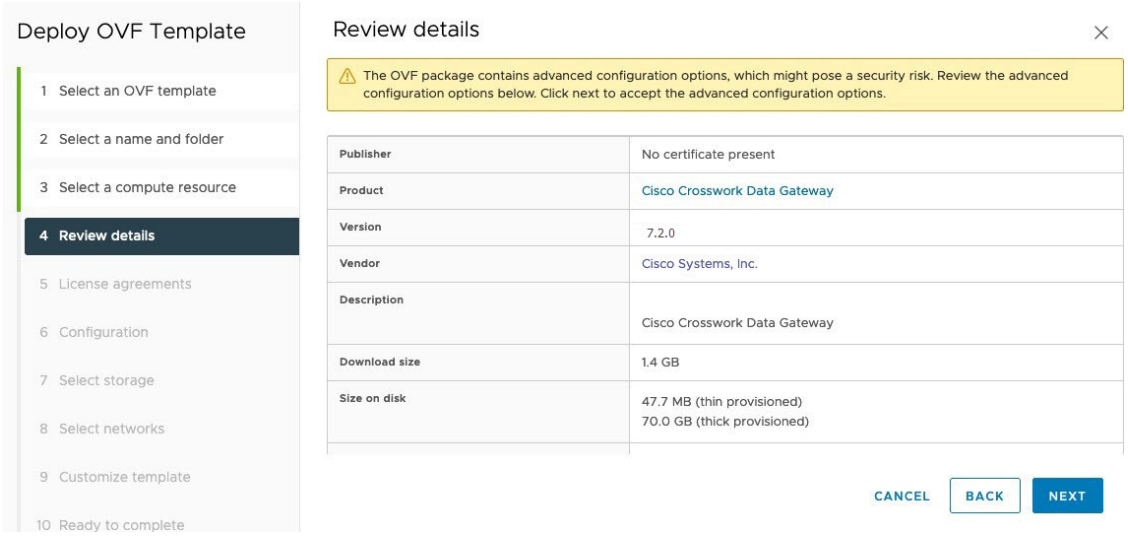


b) Click **Next**. The VMware vCenter Server validates the OVA. Network speed determines how long validation takes. When the validation is complete, the wizard moves to **4 Review details**.

Take a moment to review the OVF template to make sure it matches the version that you want to install and click **Next**.

> **Note**
>
> This information is gathered from the OVF and cannot be modified.

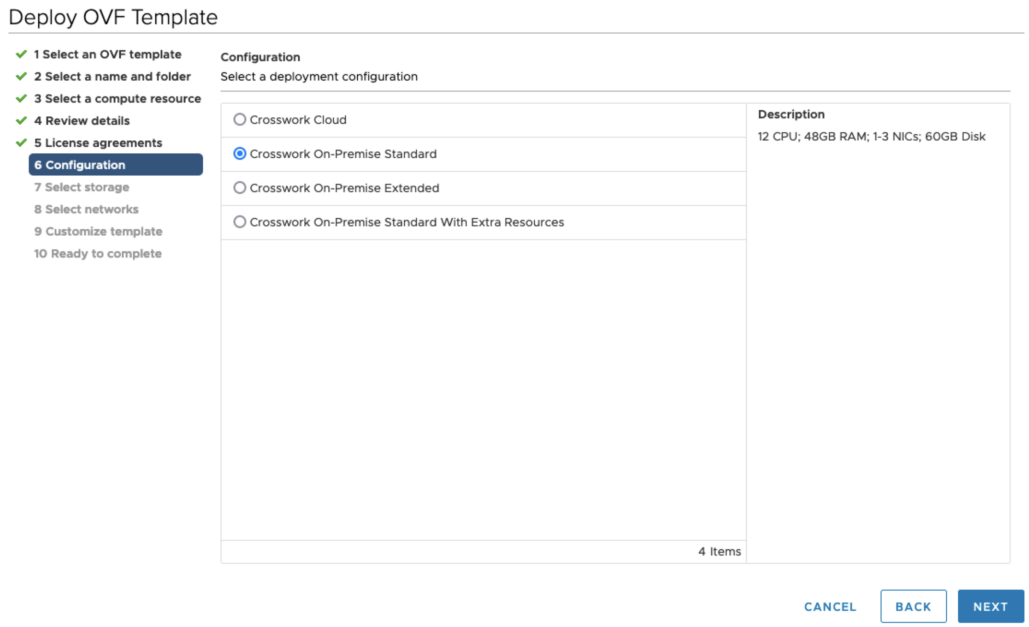
Figure 16: Review details



c) Click **Next** to go to **5 License agreements**. Review the end-user license agreement, and then click **Accept** if you agree with the conditions. Contact your Cisco representative if you do not agree with the conditions.

d) Click **Next** to go to **6 Configuration**, as shown in these figure. Select **Crosswork On-Premise Standard** or **Crosswork On-Premise Extended**. refer to Crosswork Data Gateway deployment type on page 46 for more information.

Figure 17: Deployment configuration



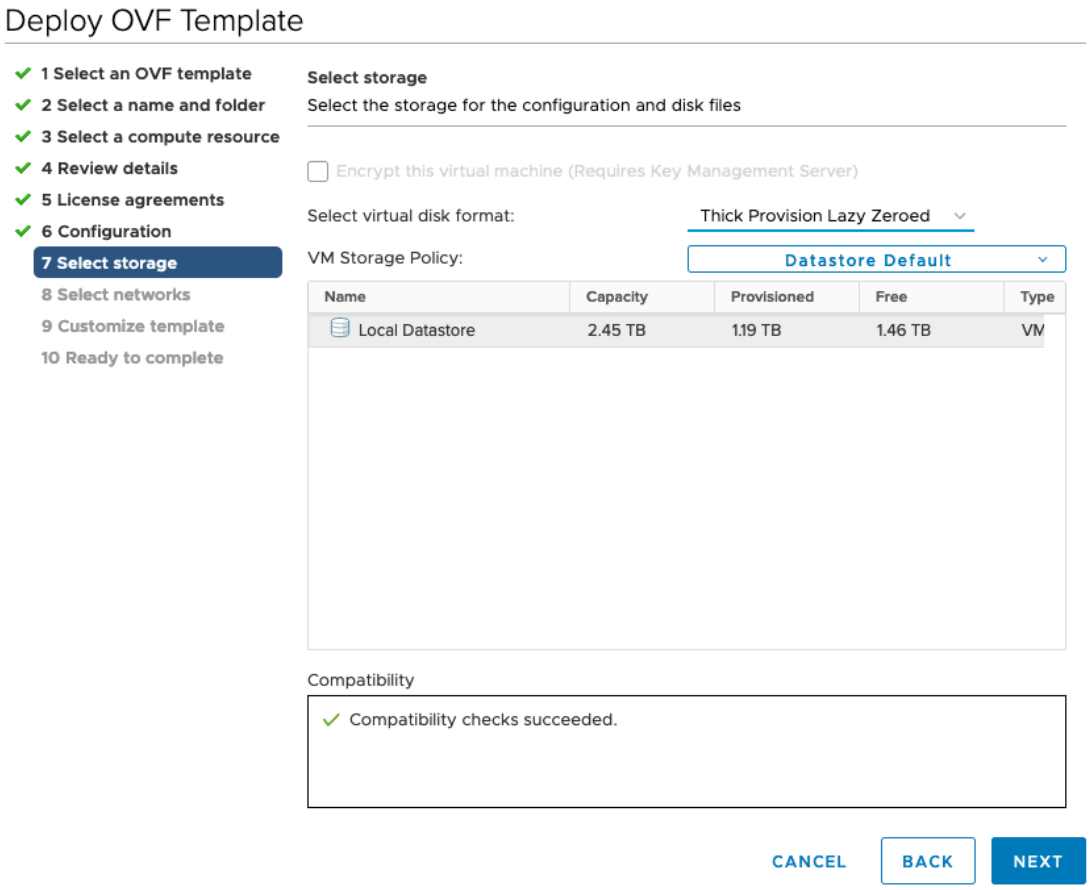
Note

Crosswork Network Controller deployments do not support the Crosswork On-Premise Standard with the extra resources option or any Crosswork Cloud configurations.

5. Click **Next** to go to **7 Select storage**, as shown in these figure.

a) Cisco recommends that you select **Thick provision lazy zeroed** from the **Select virtual disk format** drop-down list.
b) From the **Datastores** table, choose the data store you want to use and review its properties to ensure that there is enough available storage. For Crosswork On-Premise Standard deployment, the storage requirement is 70 GB and for Crosswork On-Premise Extended, it is 570 GB.

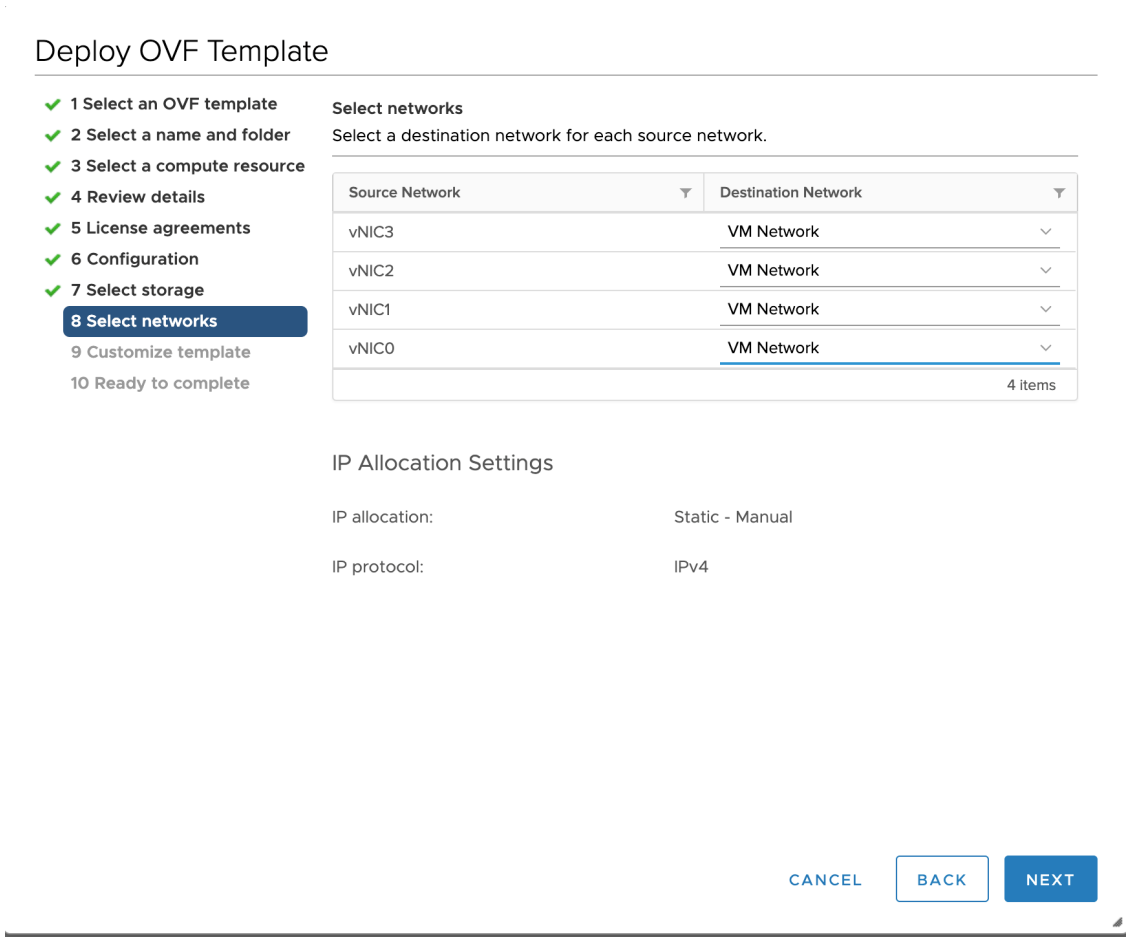
Figure 18: Storage selection



6. Configure the template and complete deployment.
   a) Click **Next** to go to **8 Select networks**, as shown in these figure. From the drop-down, select the network for each vNIC you plan to use. Leave any unused VNICs configured for the default value. For example,
      - 2 NIC: Select the appropriate Destination Network for vNIC0 and vNIC1.
      - 3 NIC: Select the appropriate Destination Network for vNIC0, vNIC1, and vNIC2.
      - Crosswork does not support the **vNIC3** network.

> **Note**
>
> Leave the vNIC3 set to the default value as it is unused.

Figure 19: Network selection



b) Click **Next** to go to **9 Customize template**, with the **Host information** already expanded. Enter the information for the parameters as explained in Table 46: Crosswork Data Gateway Deployment Parameters and Scenarios on page 132.

> **Note**
>
> For larger systems, it is likely that you have more than one Data Gateway VMs. The Data Gateway hostname should, therefore, be unique and created in a way that makes identifying a specific VM easy.

Depending on the NIC deployment, note these:

- For 2 NIC deployments, configure the IP, subnet, and gateway values for vNIC0 and vNIC1. After the Data Gateway pool is created, the VIP address is assigned as a secondary address on vNIC1.
- For the 3 NIC deployments, configure the IP, subnet, and gateway values for vNIC0 and vNIC1. The VIP address will be assigned to vNIC2 once the Data Gateway is added to a pool.

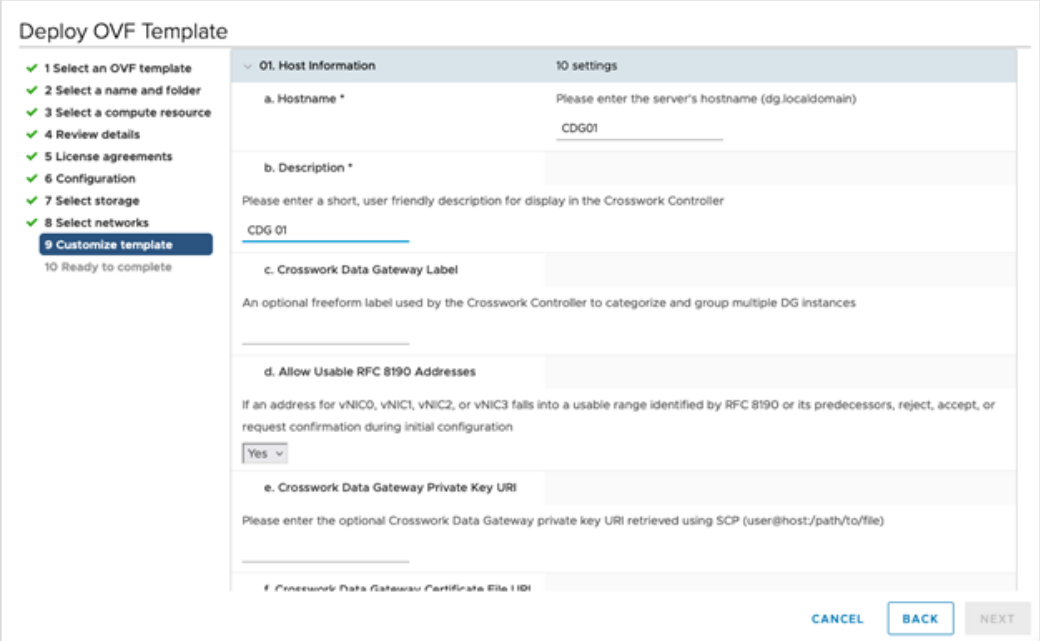
Note

There are three types of configuration settings:

1. Required settings: Mandatory values like IP addresses and hostnames.
2. Optional settings: Use case-dependent fields such as certificates or syslog server details.
3. Reserved settings: Fields that should remain blank unless explicitly instructed by the Customer Experience team.

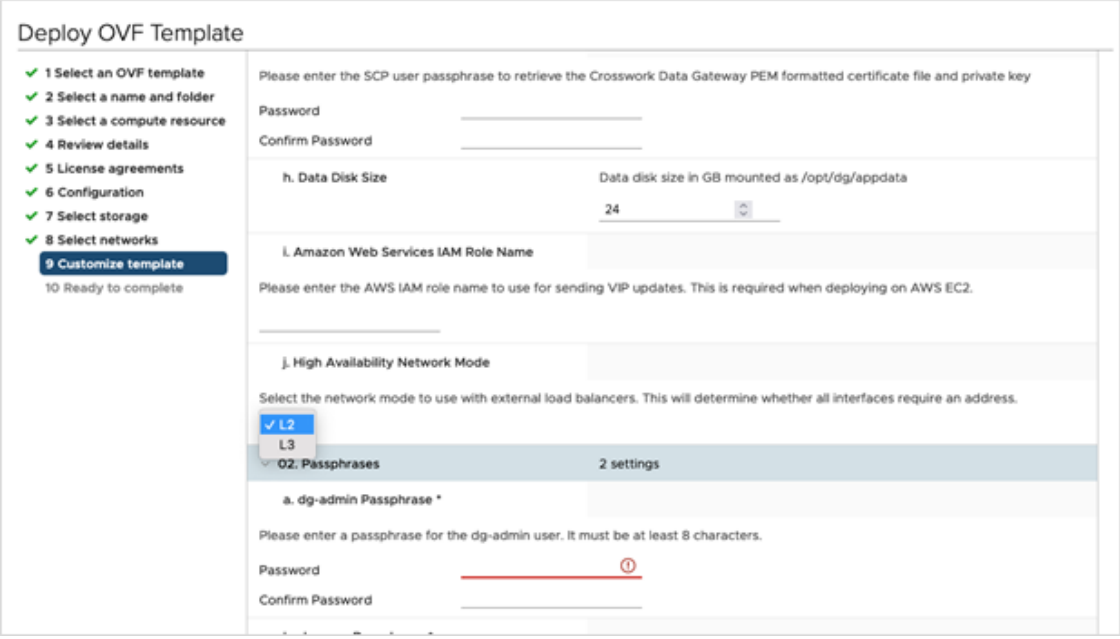
   Ensure only required and relevant optional settings are configured. Do not modify reserved fields unless directed. For more information, refer to Crosswork Data Gateway parameters and deployment scenarios on page 131.

Figure 20: Host information



For creating pools in the VMware environment, select L2 and specify IP addresses for creating the HA pool.

Figure 21: High availability network details



1. Configure the vNIC Role Assignment based on the number of NICs that you have decided to use.

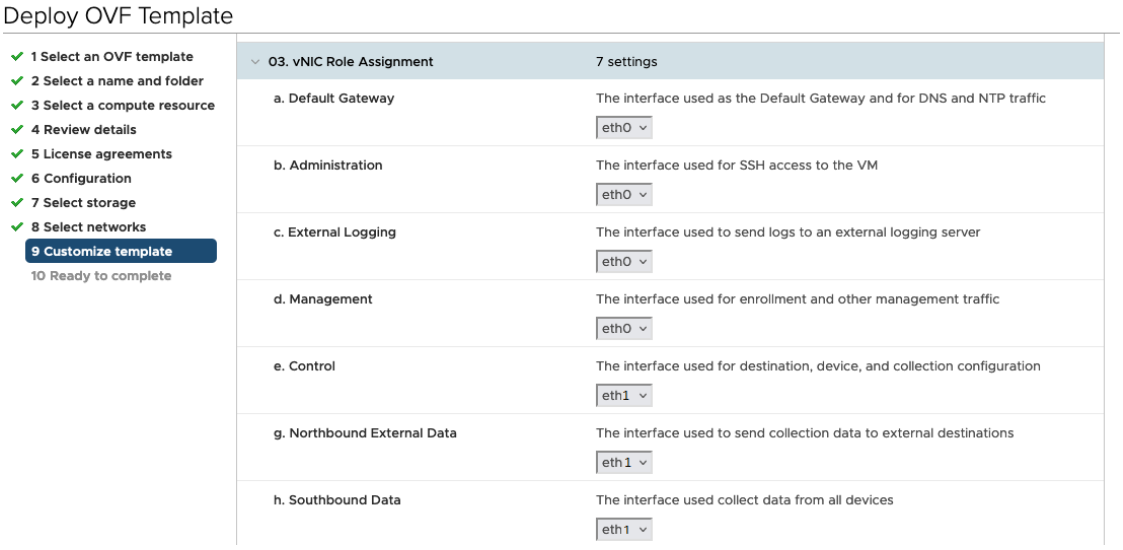
   Based on the number of NICs, refer to these to use the customized template configuration:

   Note

   The default configuration is for 3 NIC deployment.

   - refer to Deploy OVF Template - Customize Template for 2 vNICs deployment.
   - refer to Deploy OVF Template - Customize Template for 3 vNICs deployment.

Figure 22: Customize Template for 2 vNICs deployment



For 3 vNIC deployments, you can leave the settings with the default values.

Figure 23: Customize template for 3 vNICs deployment

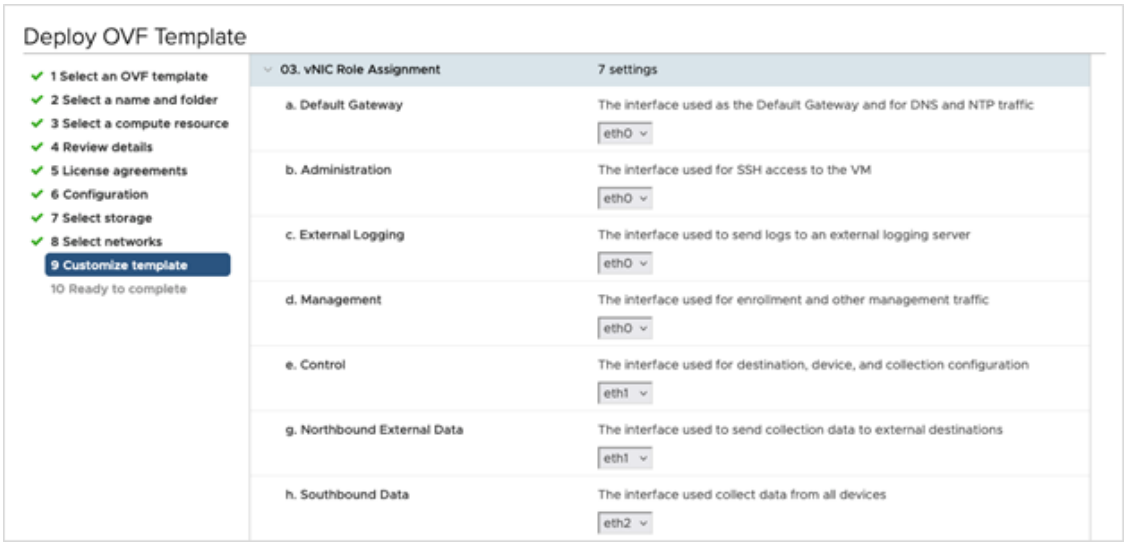
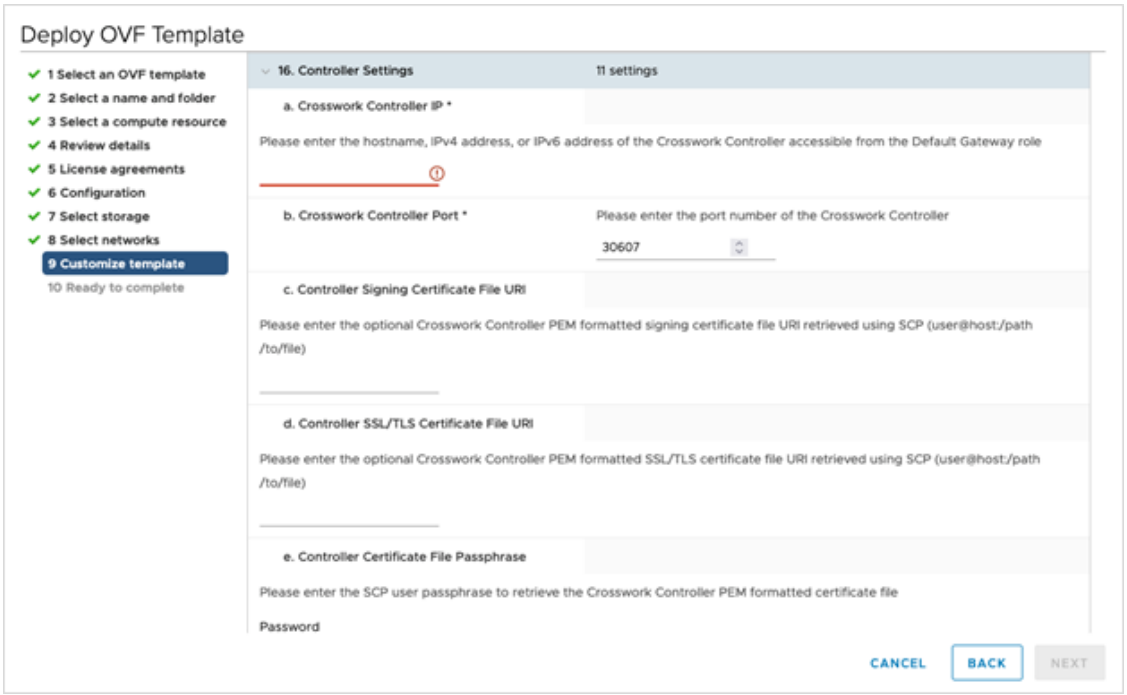
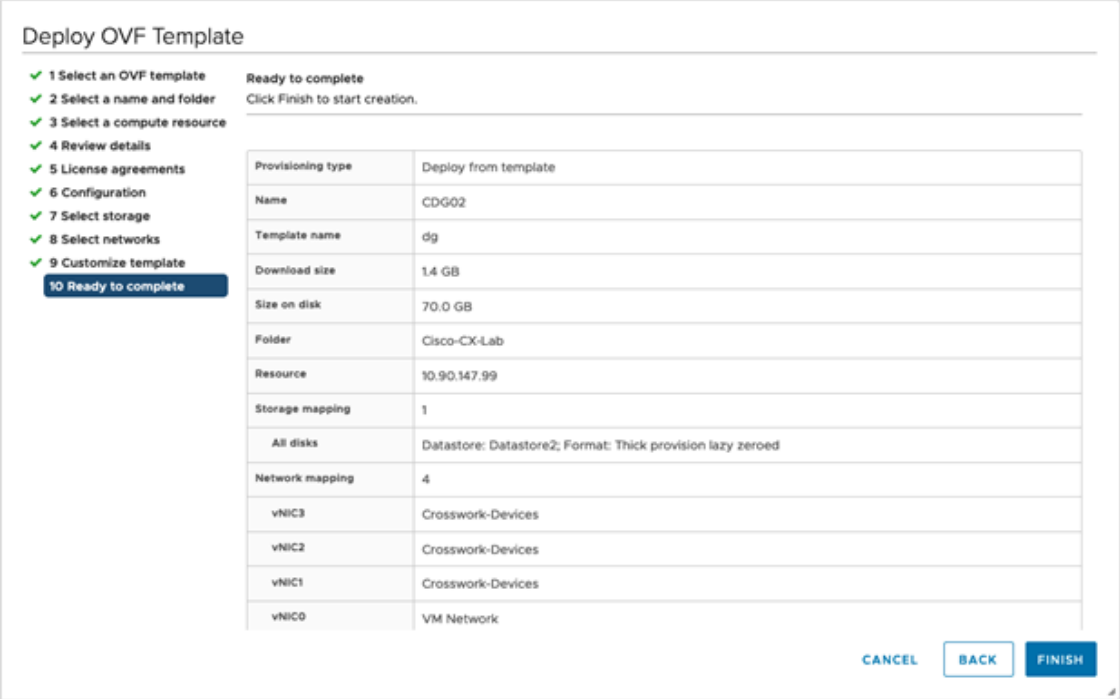


Figure 24: Controller settings



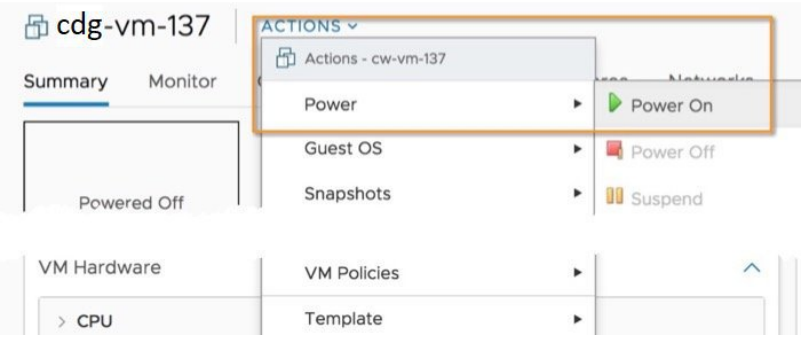
c) Click **Next** to go to **10 Ready to complete**. Review your settings and then click **Finish**.

Figure 25: Review the configuration details



d) Once the deployment status is 100%, power on the VM to complete the deployment process. Expand the host's entry so you can click the VM and then choose **Actions** > **Power** > **Power On**, as shown in these figure:

Figure 26: Power on VM



Wait for at least 5 minutes for the VM to come up and then log in via vCenter or SSH.

> **Warning**
>
> Changing the VM's network settings in vCenter may have significant unintended consequences, including but not limited to the loss of static routes and connectivity. The settings have been validated to provide the best network performance. Change these settings at your own risk.

What to do next

- After you log in, the Data Gateway present you with the welcome screen and options menu indicating that the installation was completed successfully. For information on how to log in, refer to Log in and Log out of the Data Gateway VM on page 156.

  Log out and proceed with the post-installation tasks that are documented in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

- Ensure that each Data Gateway is successfully installed and validated before proceeding to the next. Do not begin creating pools until all Data Gateways are confirmed to be built correctly.

Return to the installation workflow: Install Cisco Crosswork Network Controller on VMware vCenter on page 34

Install Crosswork Data Gateway via the OVF Tool

Use VMware OVF Tool and a deployment script to deploy the Crosswork Data Gateway OVA to vCenter. The topic covers prerequisites, package extraction, script selection, configuration, and password entry.

Use this procedure to install Crosswork Data Gateway via the OVF Tool.

Modify the list of mandatory and optional parameters in the script as per your requirements and run the OVF Tool. Refer to Table 46: Crosswork Data Gateway Deployment Parameters and Scenarios on page 132 for the list of installation parameters and their default values.

> **Note**
>
> The filenames mentioned in this topic are sample names and may differ from the actual filenames on cisco.com.

Follow these steps to log in to the Data Gateway VM from SSH:

Before you begin

- In your vCenter data center, go to **Host** > **Configure** > **Networking** > **Virtual Switches** and select the virtual switch.
- In the virtual switch, select **Edit** > **Security**, and ensure that these DVS port group properties are as shown:
  - Set **Promiscuous mode** as Reject
  - Set **MAC address changes** as Reject

  Confirm the settings and repeat the process for each virtual switch that is used by Crosswork Data Gateway.

Procedure

1. Verify the OVF Tool and obtain the deployment files.
   a) On the machine where you have the OVF tool installed, use these command to confirm that you have OVF tool version 4.4:

   ```
   ovftool --version
   ```

   b) Download the OVA and the sample script files from cisco.com. For these instructions, we use the filename as **cdg-deployment-7.2.1.uefi.ova**.

2. Extract and review the package contents.
   a) Use these command to extract the files from the tar bundle:

   ```
   tar -xvzf cdg-7.2.1-sample-install-scripts.tar.gz
   ```

   The file bundle is extracted. It includes the **cdg-7.2.1-sample-install-scripts.tar.gz** file and scripts for validating the samples install scripts.

   b) Use these command to extract the install scripts from the tar bundle:

   ```
   tar -xvzf cdg-7.2.1-sample-install-scripts.tar.gz
   ```

   c) Review the contents of the README file to understand the components that are in the package and how they are validated.

3. Select and run the deployment script.
   a) Choose the sample script that corresponds to the deployment you plan to use. Cisco provides sample scripts for 1, 2, and 3 vNIC deployments, which you may optimize to meet your needs.

   Sample scripts for 3 vNIC deployments. Customize the script for the type of deployment that you have planned. For more information, refer to the sample scripts:

   - Sample script for deploying Crosswork Data Gateway with IPv4 address on page 553
   - Sample script for deploying Crosswork Data Gateway with IPv6 address on page 551
   - Sample script for Crosswork Data Gateway dual-stack deployment on page 549
   - Sample Script for Crosswork Data Gateway Dual-Stack Deployment with Geo-Redundancy

   b) Use these command to make the script executable:

   ```
   chmod +x {filename}
   ```

   c) Use these command to execute the script from the directory where the OVA and script files are stored:

   ```
   ./{script name} {path and ova file name}
   ```

   For example:

   ```
   ./<script name> <Absolute path to cdg-deployment-7.2.1.uefi.ova>
   ```

   d) If the values provided in the script are valid, provide the vCenter user's password when you are prompted.

   If the script fails due to invalid values, a message like these is displayed:

   ```
   admin@nso-576-tsdn-410-aio:~/CDG_Install$ ./three-nic
   /home/admin/CDG_Install/cdg-deployment-7.2.1.uefi.ova Opening OVA source:
    /home/admin/CDG_Install/cdg-deployment-7.2.1.uefi.ova The manifest does
   not validate Warning: - Line -1: Unsupported value 'firmware' for attribute
    'key' on element 'ExtraConfig'. - Line -1: Unsupported value
   'uefi.secureBoot.enabled' for attribute 'key' on element 'ExtraConfig'.
   Enter login information for target vi://rcdn5-spm-vc-01.cisco.com/ Username:
    johndoe Password: ******
   ```

   After entering the password, monitor the screen or the vCenter console to review the installation progress. For example,

   ```
   Opening VI target:
   vi://johndoe@rcdn5-spm-vc-01.cisco.com:443/Cisco-sample-sample/host/10.10.100.10
   ```

```
 Warning: - Line 146: Unable to parse 'enableMPTSupport' for attribute
'key' on element 'Config'. - Line 229: Unable to parse 'vmxnet3.noOprom'
for attribute 'key' on element 'Config'. Deploying to VI:
vi://johndoe@rcdn5-spm-vc-01.cisco.com:443/Cisco-sample-sample/host/10.10.100.10
 Disk progress: 65%
```

When the installation is complete, the Data Gateway VM is powered on, is automatically configured based on the settings that you have provided in the script, and registers with the Crosswork cluster.

What to do next

Log in to the VM. For more information, refer to Log in and Log out of the Data Gateway VM on page 156. After you log in, the Data Gateway should present you with the welcome screen, and options menu indicating that the installation is complete. Log out and proceed with the postinstallation tasks explained in Crosswork Data Gateway post-installation tasks on page 159.

Log in and Log out of the Data Gateway VM

Explains the available ways to access the Data Gateway VM and directs readers to the login and logout procedures. The child topics describe SSH and vCenter access.

You can log in to the Data Gateway VM in one of these ways:

Access methods

- Access Crosswork Data Gateway VM from SSH on page 156
- Access Crosswork Data Gateway through vCenter on page 157

To log out from the Data Gateway VM, refer to Log out of Crosswork Data Gateway VM on page 157.

Access Crosswork Data Gateway VM from SSH

Connect to the Crosswork Data Gateway VM over SSH with an administrator or operator account and its installation password. The topic also describes the login prompt and a network troubleshooting outcome.

Use this procedure to access Crosswork Data Gateway VM from SSH.

The SSH process is protected from brute force attacks by blocking the client IP after a number of login failures. Failures such as incorrect username or password, connection disconnect, or algorithm mismatch are counted against the IP. Up to 4 failures within a 20 minute window causes the client IP to be blocked for at least 7 minutes. Continuing to accumulate failures cause the blocked time to be increased. Each client IP is tracked separately.

Follow these steps to log in to the Cisco Crosswork Data Gateway VM from SSH:

Procedure

1. From your work station with network access to the Cisco Crosswork Data Gateway management IP, run these command:

   ssh <username>@<ManagementNetworkIP>

   where **ManagementNetworkIP** is the management network IP address.

   For example,

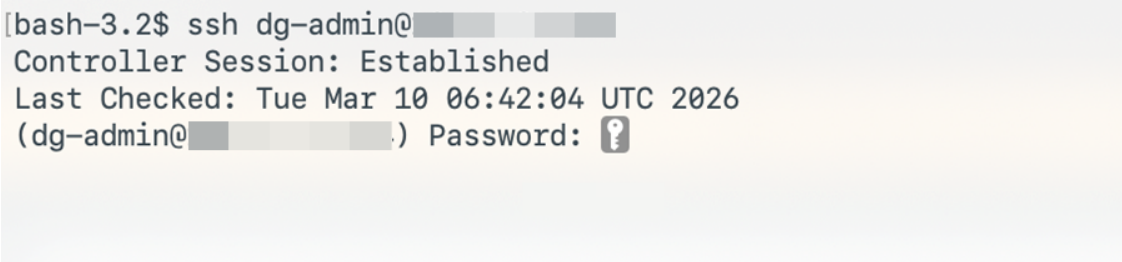
   To login as administrator user: **ssh dg-admin@<ManagementNetworkIP>**

To login as operator user: **ssh dg-oper@<ManagementNetworkIP>**

2. Provide the corresponding password, which was created during installation process, and press **Enter**.

   The Crosswork Data Gateway flash screen opens prompting for password.

   Figure 27: Crosswork screen



If you are unable to access the Cisco Crosswork Data Gateway VM, there is an issue with your network configuration settings. From the console, check the network settings. If they are incorrect, it is best to delete the Cisco Crosswork Data Gateway VM and reinstall with the correct network settings.

Access Crosswork Data Gateway through vCenter

Open the Data Gateway VM console in vCenter and sign in with an administrator or operator account. The procedure describes the console and login prompt.

Use this procedure to access Crosswork Data Gateway through vCenter.

Procedure

1. Locate the VM in vCenter and then right-click and select **Open Console**.

   The Crosswork Data Gateway console comes up.

2. Enter username (`dg-admin` or `dg-oper` as per the role assigned to you) and the corresponding password (the one that you created during the installation process) and press **Enter**.

   The Crosswork Data Gateway flash screen opens prompting for password.

   Figure 28: Crosswork screen



Log out of Crosswork Data Gateway VM

Log out of the Crosswork Data Gateway VM by selecting Logout from its main menu. This action ends the interactive session.

Log out of the Crosswork Data Gateway VM through its main menu.

Procedure

To log out, select option l Logout from the Main Menu and press Enter or click OK.

Crosswork Data Gateway authentication and enrollment

Describes how Crosswork and the Data Gateway VM establish identity, exchange certificates, download configuration, and report enrollment status. It also covers verification steps and status changes during the handshake and image download.

This process begins after the Crosswork Data Gateway VM is installed.

Summary

Cisco Crosswork and the Crosswork Data Gateway VM establish connectivity, confirm each other's identity using signed certificates, and exchange configuration. Cisco Crosswork displays enrollment and operational status for the Data Gateway VM.

Workflow

The authentication and enrollment process follows these stages.

1. Once the Crosswork Data Gateway is installed, it identifies itself and enrolls with Cisco Crosswork automatically. Cisco Crosswork then instantiates a new Data Gateway instance in its database and waits for a "first-sign-of-life" from the Data Gateway VM.
2. After the connectivity is established, the Data Gateway instance confirms the identity of the controller application (Cisco Crosswork) and offers its own proof of identity via signed certificates. Crosswork Data Gateway then downloads the configuration files and functional images (collection profiles) from Cisco Crosswork.
3. To verify if the Data Gateway VM has enrolled successfully with Cisco Crosswork:
   - 1. Log in to the Cisco Crosswork UI. refer to Log into the Cisco Crosswork UI on page 120.
   - 2. Navigate to **Administration** > **Data Gateway Management**.
   - 3. Click on the **Data Gateway Instances** tab.

     All the Data Gateway VMs that have successfully enrolled with Cisco Crosswork are displayed here.
4. The initial **Operational State** of Data Gateway VMs is **Unknown**. During the handshake and image download, the status is **Degraded**. After the handshake is complete, the status is **Not Ready**. While it depends on the bandwidth between the Data Gateway VMs and Cisco Crosswork, this operation typically takes between 5 to 10 minutes. If it takes longer than the stipulated duration, contact Cisco Customer Experience team for assistance.
5. For information about the different operational states of the VMs, refer to the *Crosswork Data Gateways* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

What's next

Return to the installation workflow: Install Cisco Crosswork Network Controller on VMware vCenter on page 34

Crosswork Data Gateway post-installation tasks

Outlines a post-installation configuration task for Crosswork Data Gateway and links to the timezone procedure. The task updates the VM timezone after installation.

Configure the Data Gateway VM timezone after installation.

What to do next

After installing Crosswork Data Gateway, configure the timezone of the Data Gateway VM.

- Configure the timezone of the Crosswork Data Gateway VM on page 159

Return to the installation workflow: Install Cisco Crosswork Network Controller on VMware vCenter on page 34

Configure the timezone of the Crosswork Data Gateway VM

Set the Data Gateway VM timezone through its interactive menu, then reboot the VM to apply the setting. The procedure includes selecting a geographic area and city or region.

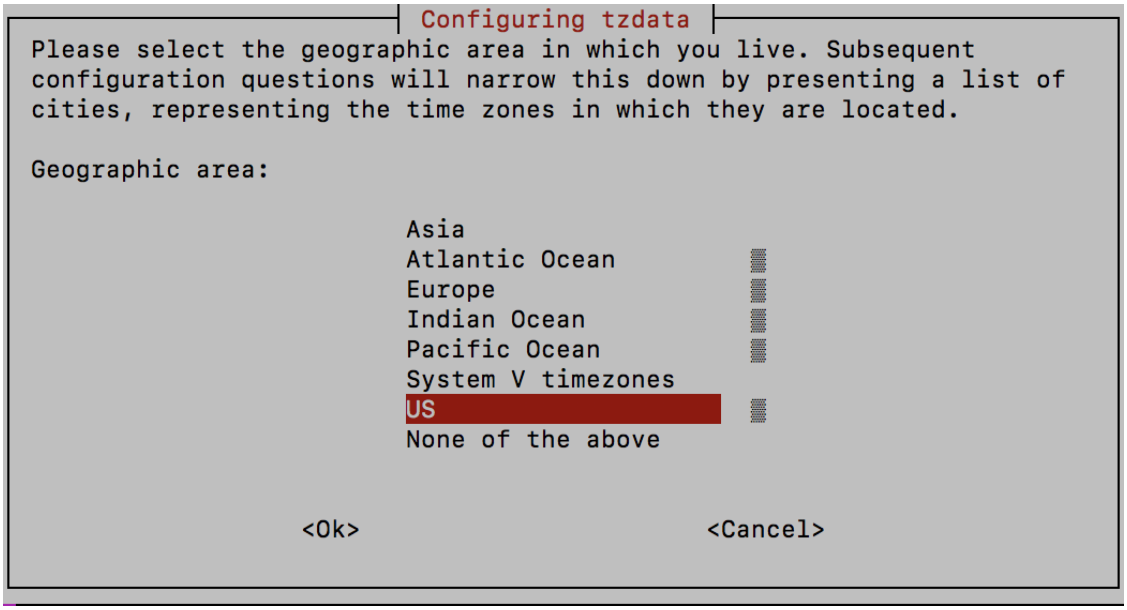
Set the timezone for the Crosswork Data Gateway VM.

By default, the Crosswork Data Gateway VM launches with the UTC timezone. To ensure that all Data Gateway processes (including Showtech logs) reflect the correct timestamp for your location, update the timezone to match your geographical area.

Procedure

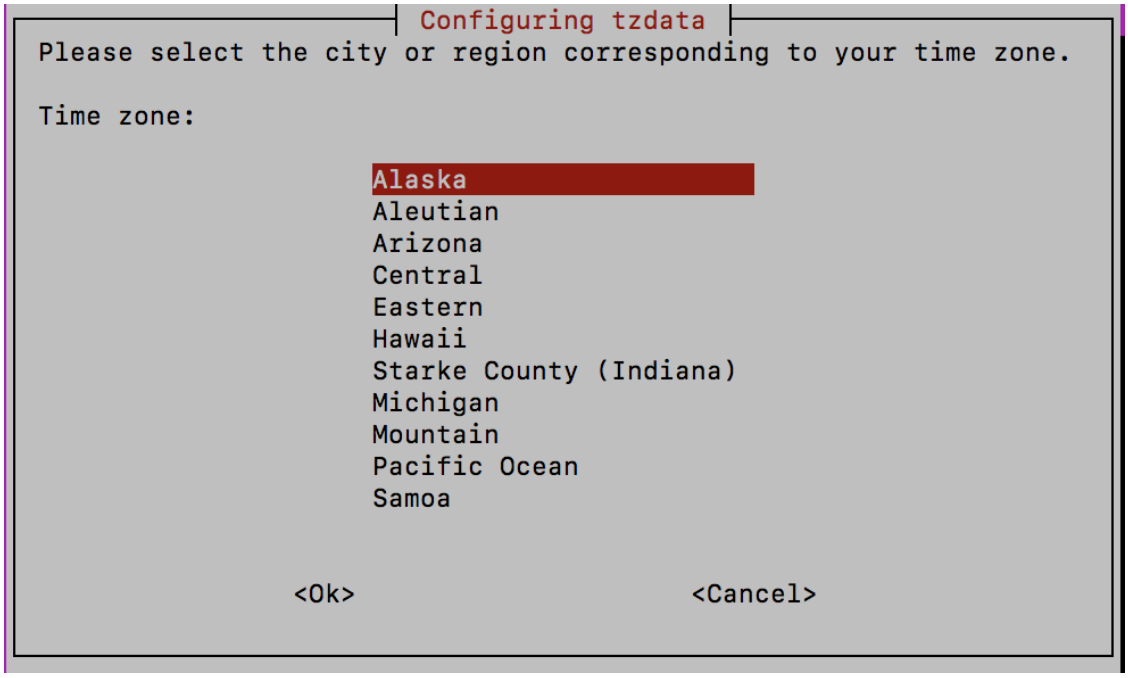
1. Log in to the Crosswork Data Gateway VM.
2. In the Crosswork Data Gateway VM interactive menu, select **3 Change Current System Settings**.
3. From the menu, select **9 Timezone**.
4. Select the geographical area in which you are located.

Figure 29: Geographic area selection



5. Select the city or region that corresponds to your timezone.

Figure 30: Region selection



6. Select **OK** to save the settings.
7. Reboot the Crosswork Data Gateway VM to apply the new timezone to all processes.
8. Log out of the Crosswork Data Gateway VM.

Troubleshoot installation and enrollment

Lists common installation and enrollment problems, diagnostic information, and source-backed actions for resolving them. Entries address timing, certificate, network, and VM state issues.

Use the troubleshooting entries to diagnose installation and enrollment problems.

If Crosswork Data Gateway fails to autoenroll with Crosswork Network Controller, you can collect Crosswork Data Gateway show-tech (**Main menu** > **5 Troubleshooting** > **2 Run show-tech**) and check for the reason in `controller-gateway` logs.

For more information on how to collect showtech logs, refer to the *Download the showtech logs* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. If there are session establishment or certificate-related issues, ensure that the `controller.pem` certificate is uploaded using the Interactive Console.

> **Note**
>
> When using an IPv6 address, it must be surrounded by square brackets ([1::1]).

these table lists common problems that might be experienced while installing or enrolling a Data Gateway, and provides approaches to identifying the source of the problem and solving it.

Table 48: Troubleshooting the Installation or Enrollment

Issue	Action

Crosswork Data Gateway cannot be enrolled with Crosswork Network Controller due to an NTP issue, that is, there is a clock-drift between the two.
 	
The clock-drift might be with either Crosswork Data Gateway or Crosswork Network Controller.
 	
Also, on the NTP servers for Crosswork Network Controller and Crosswork Data Gateway, the initial time is set to the ESXi server. For this reason, the ESXi server must also have NTP configured.
 	
Sync the clock time on the host and retry.
	1. Log in to the Data Gateway VM.
	
	2. From the main menu, select **5 Troubleshooting** > **2 Run show-tech**.
	
	Enter the destination to save the tarball containing logs and vitals and click **OK**.
	
	The show-tech is now encrypted with a file extension ending with .tar.xz.
	
	3. Run these command to decrypt the show-tech file.
	
	`openssl enc -d -aes-256-ctr -pbkdf2 -md sha512 -iter 100000 -in <showtech file> -out <decrypted filename> -pass pass:<password>`
	
	Example,
	
	`openssl enc -d -aes-256-ctr -pbkdf2 -md sha3-512 -iter 100000 -in show-tech-file.tar.xz.enc -out show-tech-file.tar.xz -pass pass:myPassword`
	
	In the show-tech logs (in file `session.log` at location `/opt/dg/log/controller-gateway/session.log`), if you refer to the error `UNAUTHENTICATED:invalid certificate. reason: x509: certificate has expired or is not yet valid`, then there is a clock-drift between Crosswork Data Gateway and Crosswork Network Controller.
	
	3. From the main menu, go to **3 Change Current System Settings** > **1 Configure NTP**.
	
	Configure NTP to sync with the clock time on the Crosswork Network Controller server and try reenrolling Crosswork Data Gateway.

Issue	Action
Crosswork Data Gateway remains in degraded state for more than 10 minutes with reason stated as "Could not collect vitals" due to certificate errors.	1. Log in to the Data Gateway VM. 2. From the main menu, select **5 Troubleshooting** > **2 Run show-tech**. Enter the destination to save the tarball containing logs and vitals and click **OK**. The show-tech is now encrypted with a file extension ending with .tar.xz. 3. Run these command to decrypt the show-tech file. `openssl enc -d -aes-256-ctr -pbkdf2 -md sha3-512 -iter 100000 -in show-tech-file.tar.xz.enc -out show-tech-file.tar.xz -pass pass:myPassword` Example, `openssl enc -d -aes-256-ctr -pbkdf2 -md sha3-512 -iter 100000 -in show-tech-file.tar.xz.enc -out show-tech-file.tar.xz -pass pass:myPassword` In the show-tech logs (in file `gateway.log` at location `/opt/dg/log/controller-gateway/gateway.log`), if you refer to certificate errors, then reupload the Controller Signing Certificate, as explained in these steps: 1. From the main menu, select **3 Change Current System Settings** > **7 Import Certificate**. 2. From the **Import Certificates** menu, select **1 Controller Signing Certificate File** and click **OK**. 3. Enter the SCP URI for the certificate file and click **OK**.
Crosswork Data Gateway remains in degraded state for more than 10 minutes with reason stated as "gRPC connection cannot be established" due to certificate errors.	1. Reupload the certificate file using these steps: a. From the main menu, select **3 Change Current System Settings** > **7 Import Certificate**. b. From the **Import Certificates** menu, select **1 Controller Signing Certificate File** and click **OK**. c. Enter the SCP URI for the certificate file and click **OK**. 2. Reboot the Data Gateway VM: a. From the main menu, select **5 Troubleshooting** and click **OK**. b. From the Troubleshooting menu, select **4 Reboot VM** and click **OK**. c. When the reboot is complete, check if the Data Gateway's operational status is **Up**.

Issue	Action
Crosswork Data Gateway goes into Error state	• For vCenter deployment, check the vNIC values in the OVF template and correct them if required. • Before deploying a new Data Gateway instance with the same configuration, delete the existing errored instance. For details on how to delete the instance and its effects on collection jobs, refer to *Delete the Data Gateway instance from Crosswork Network Controller* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. • If you are redeploying the Data Gateway instance, you must first delete the existing instance that is currently in an error state. For more information about • deleting the instance and its impact on the collection jobs, refer to *Delete the Data Gateway instance from Crosswork Network Controller* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. • understanding the impact of redeployment on enrollment, refer to *Redeploy a Crosswork Data Gateway Instance* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
Crosswork Data Gateway deploys Standard profile instead of Extended profile	Check the `Deployment` parameter in the OVF template in case of vCenter. If `Deployment` parameter mismatches or does not exist for an Extended profile, then Crosswork Data Gateway deploys the Standard profile by default.
During a Crosswork upgrade, some of the Crosswork Data Gateways may not get upgraded or reenrolled leading to logging multiple error messages in the dg-manager logs.	Reenroll or redeploy the Crosswork Data Gateways. For more information, refer to the *Redeploy a Crosswork Data Gateway Instance* and *Re-enroll Crosswork Data Gateway* sections in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
If a Data Gateway is redeployed without moving the gateway to the Maintenance mode, Crosswork enrollment will be unsuccessful and errors will be logged in the dg-manager and controller-gateway logs.	Move the Data Gateway to the **Maintenance** mode or manually reenroll the gateway. For more information, refer to the *Re-enroll Crosswork Data Gateway* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
Crosswork Data Gateway fails to enroll with Cisco Crosswork when the Data Gateway is deployed on a single stack while Crosswork is deployed on dual stack.	Redeploy the Data Gateway on dual stack. For more information, refer to the *Redeploy a Crosswork Data Gateway Instance* and *Re-enroll Crosswork Data Gateway* sections in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Import controller signing certificate file

Import the Controller Signing Certificate File from the Data Gateway VM menu and verify that installation succeeds. The topic also lists conditions that can prevent automatic import.

Import the Controller Signing Certificate File into the Data Gateway VM.

The Controller Certificate file is automatically imported after the VM boots. If there is an import failure, the Data Gateway VM makes several attempts to import the certificate while giving you the option to manually import it.

- You have not specified the **Controller Signing Certificate File URI** under the **Controller Settings** during installation.
- Cisco Crosswork was upgraded or reinstalled and you need to authenticate and enroll Crosswork Data Gateway with Cisco Crosswork.
- Cisco Crosswork configuration is in-progress when Crosswork Data Gateway tries to import the Controller Certificate file.
- The Cisco Crosswork Controller IP address in unreachable or incorrect.
- The Cisco Crosswork username or password is incorrect.

Follow these steps to import the controller signing certificate file:

Note

Procedure

1. From the Data Gateway VM's Interactive Menu, select **3 Change Current System Settings**.

   The **Change System Settings** menu opens.
2. Select **7 Import Certificate**.
3. From the **Import Certificates** menu, select **1 Controller Signing Certificate File**.
4. Enter the SCP URI for the certificate file.

   An example URI is given below:

   ```
   cw-admin@{server  ip }:/home/ cw -admin/ controller.pem
   ```
5. Enter the SCP passphrase (the SCP user password).

   The certificate file is imported.
6. Verify that the certificate was installed successfully. refer to View the controller signing certificate file on page 164.

View the controller signing certificate file

Display the Controller Signing Certificate File through the Data Gateway VM interactive menu. The topic explains that the menu displays the default or newly imported certificate.

Display the Controller Signing Certificate File on the Data Gateway VM.

Procedure

1. From the Data Gateway VM's interactive menu, select **2 Show System Settings**.
2. From the **Show Current System Settings** menu, select **7 Certificates**.
3. Select **2 Controller Signing Certificate File**.

   Crosswork Data Gateway displays the default certificate if no new certificate has been imported. Otherwise, it displays the new certificate if it was successfully imported.

Part III

Install Cisco Crosswork Network Controller on KVM

Topics:

- Install Crosswork Cluster on KVM
- Install Crosswork Data Gateway on KVM

7 Install Crosswork Cluster on KVM

Topics:

- Installation overview
- KVM host bare metal requirements
- Installation parameters
- Install Crosswork Network Controller cluster using CLI
- Monitor cluster activation
- Log into the Cisco Crosswork UI
- Troubleshoot the cluster
- Installation workflow for deploying SR-PCE (Cisco IOS XRv 9000) on RHEL 8.10 KVM

Introduces the KVM installation, host requirements, and deployment workflows covered in this chapter.

This chapter describes installing Crosswork Network Controller on KVM.

It covers host and network requirements, cluster deployment and activation, and SR-PCE deployment on RHEL KVM.

Installation overview

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

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.

KVM host bare metal requirements

these requirements are mandatory if you are planning to install Crosswork Network Controller on RHEL KVM.

these requirements are mandatory if you are planning to install Crosswork Network Controller on RHEL KVM.

these requirements are mandatory if you are planning to install Crosswork Network Controller on RHEL KVM.

Table 49: Host bare metal requirements

Component	Minimum requirement per host
Processor	Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz or latest
NIC	2 x 10 Gbps NICs
OS	Red Hat Enterprise Linux 8.10 Red Hat Enterprise Linux 9.4

Resource requirements

For information on resource requirements per VM and per host, please refer to Resource footprint for KVM on page 26 and Physical server requirements on page 29, respectively.

Best practices

- **Resource allocation**: Refer to the resource requirements outlined in Physical server requirements on page 29, for each host.

  > **Note**
  >
  > Crosswork Network Controller cluster nodes place high demands on the VMs. Ensure that CPU and memory resources on the machines hosting the nodes are not oversubscribed.

- **Node distribution**: Ensure all Crosswork Hybrid nodes are spread across multiple RHEL bare metals. While it is possible to deploy all cluster nodes on a single RHEL bare metal (provided it meets the requirements), it is recommended to distribute the nodes across multiple RHEL bare metals. This approach prevents the host from being a single point of failure and enhances solution resilience.

- **Network configuration**: Ensure the networks required for the Crosswork Management and Data networks are built and configured in the data centers. These networks must allow low-latency L2 communication with a round-trip time (RTT) of 10 ms or less.

  Note

  The same network names must be used and configured on all RHEL bare metal host machines that are hosting the Crosswork VMs.

Note

Crosswork Network Controller cluster VMs (Hybrid and Worker nodes) must run on hardware with Hyper Threading disabled to ensure consistent, real-time performance for CPU-intensive workloads, as Hyper Threading may cause resource contention and unpredictable performance

Installation parameters

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

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

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 on page 43.

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

Note

- 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 50: General parameters

Parameter name	Description
`CwInstaller`	Set as "False".
`Deployment`	Enter the deployment type (IPv4, IPv6, or DUALSTACK.

Parameter name	Description
`Disclaimer`	Enter the disclaimer message.
`ManagementIPv4Address` `ManagementIPv6Address`	The Management IP address of the VM (IPv4 and/or IPv6). **Note** You cannot reuse the bare metal IP addresses for Crosswork VM.
`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) of the cluster. By default, this field is set as empty.
`DataIPv4Address` `DataIPv6Address`	The Data IP address of the VM (IPv4 and/or IPv6). **Note** You cannot reuse the bare metal IP addresses for Crosswork VM.
`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.
`DataVIP`	The Data Virtual IP address for the cluster.

Parameter name	Description
`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) of the cluster. By default, this field is set as empty.
`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.
`Domain`	The domain name used for the cluster.
`CWPassword`	Password to log into Cisco Crosswork. When setting up a VM, ensure the password is strong and meets these criteria: • It must be at least 8 characters long and include uppercase and lowercase letters, numbers, and at least one special character. • these special characters are not allowed: backslash (\), single quote ('), or double quote ("). • Avoid using passwords that resemble dictionary words (for example, "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.
`VMSize`	Sets the VM size for the cluster. For cluster deployments, the only supported option is `"Large"`. **Note** • If you leave this field blank, the default value (`"Large"`) is automatically selected. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`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.

Parameter name	Description
`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. **Note** 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. **Note** The final backup partition size will be calculated dynamically. This parameter defines the minimum.
`ddatafs`	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 **false** for production deployments.

Parameter name	Description
`EnableHardReservations`	Determines the enforcement of VM CPU and Memory profile reservations. This is an optional parameter and the default value is `"True"`, if not explicitly specified. **Note** This parameter accepts a string value, so be sure to enclose the value in double quotes. 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. 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.
`ramdisk`	Size of the Ram disk. **Warning** Use the `RamDiskSize` option only for temporary lab testing or product demonstrations. This option is intended for environments where disk performance does not meet the application requirements and where rebuilding the VMs after a restart is an acceptable trade-off. Do not use this option for deployments that require resiliency or reliability. The minimum value for `RamDiskSize` is `2`. When you set `RamDiskSize` to a non-zero value, leave `HSDatastore` unset as `""`. If the VM is powered down or restarted after installation with `RamDiskSize` configured, the system may be unable to launch applications, and Kubernetes may report errors.
`OP_Status`	This optional parameter is used (uncommented) to import inventory post manual deployment of Crosswork cluster. The parameter refers to the state for this VM. To indicate a running status, the value must be 2 (`#OP_Status = 2`).

Parameter name	Description
`SchemaVersion`	The configuration Manifest schema version. This indicates the version of the installer to use with this template. Schema version should map to the version packaged with the sample template on cisco.com. 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.
`logfs`	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.
`corefs`	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.

Parameter name	Description
`Timezone`	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. This is an optional parameter. **Note** 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 `Timezone` parameter in Crosswork.
`EnableSkipAutoInstallFeature`	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"`. The recommended value for cluster deployment is `"False"`. **Note** • If left blank, the default value is automatically selected. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`EnforcePodReservations`	Enforces minimum resource reservations for the pod. If left blank, the default value (`"True"`) is selected. This parameter accepts a string value, so be sure to enclose the value in double quotes.
`K8Orch`	Enforces minimum resource reservations for the pod. If left blank, the default value (`"True"`) is selected. This parameter accepts a string value, so be sure to enclose the value in double quotes.
`K8sServiceNetwork`	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.

Parameter name	Description
`K8sPodNetwork`	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.
`bootOptions.efiSecureBootEnabled`	Default value is `"True"`. This parameter accepts a string value, so be sure to enclose the value in double quotes.
`IgnoreDiagnosticsCheckFailure`	Used to set the system response in case of a diagnostic check failure. 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. You are recommended to select the default value. This parameter accepts a string value, so be sure to enclose the value in double quotes. **Note** • The log files (`diagnostic_stdout.log` and `diagnostic_stderr.log`) can be found at **/var/log**. The result from each diagnostic execution is kept in a file at **/home/cw-admin/diagnosis_report.txt**. • Use **diagnostic all** command to invoke the diagnostic manually on day N. • Use **diagnostic history** command to view previous test report.

Install Crosswork Network Controller cluster using CLI

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

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

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

Table 51: Installation workflow

Step	Action
1. Ensure you have performed the preliminary checks.	refer to Preliminary checks on page 179 for details.
2. Set up and validate the KVM environment.	refer to Set up and validate KVM on RHEL on page 179.

Step	Action
3. Configure network bridges and SRIOV	refer to Configure network bridges or SRIOV on page 180.
4. Install Crosswork Network Controller cluster on KVM.	refer to Install Crosswork Network Controller cluster on KVM using CLI on page 184.

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 (that is, `/etc/sudoers`).

Preliminary checks

Virtualization: Ensure that your system supports virtualization. This is typically enabled in the BIOS.

Virtualization: Ensure that your system supports virtualization. This is typically enabled in the BIOS. To check, use these commands:

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

Outlines how to install and validate KVM on RHEL before deploying the application.

This topic describes how to install KVM on RHEL.

Before you begin

Confirm you have administrator (sudo) privileges.

Procedure

1. Ensure you have the latest packages (RHEL version 8.10 or 9.4) needed for KVM installation.
2. Install the required virtualization tools.
   a) 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.

   b) Install libvirt virtualization daemon, which is necessary for managing VMs.

   ```
   sudo dnf install -y libvirt
   ```

c) Install virt-manager, a graphical interface for managing VMs.

```
sudo dnf install virt-manager -y
```

d) Install additional virtualization tools for managing VMs.

```
sudo dnf install -y virt-top libguestfs-tools
```

3. Start and enable the libvirtd virtualization daemon.
   a) Start the libvirtd daemon.

```
sudo systemctl start libvirtd
```

   b) Enable the libvirtd daemon.

```
sudo systemctl enable libvirtd
```

   c) Verify that the daemon is running.

```
sudo systemctl status libvirtd
```

4. Add users to the required groups, such as libvirt and qemu.

   In these commands, replace *your_username* with the actual username.

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

5. 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
```

6. 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 these commands to enable it.

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

KVM is successfully installed on RHEL and validated.

Configure network bridges or SRIOV

This section describes the procedures for configuring network bridges or SRIOV.

Crosswork requires a 10G interface for all the data layer communications to operate at a scale. You can choose any networking configuration that provides 10G throughput.

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

For detailed instructions, refer to these topics:

- Configure network bridges on page 181
- Configure SRIOV on page 183

Configure network bridges

Outlines how to configure network bridges so virtual interfaces can communicate across the required networks.

Configure network bridges.

A network bridge, such as Linux bridge and Open vSwitch (OVS), acts like a virtual network switch, allowing multiple network interfaces to communicate as if they are on the same physical network. For detailed information, refer to RHEL 8.x or RHEL 9.x documentation.

Complete these steps to configure network bridges.

Procedure

1. Create network bridge connection for management interface:
   a) 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
   ```

   b) 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> master intMgmt
   ```

   Example

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

   c) Assign IP address, netmask, gateway, DNS, and domain to `intMgmt`.

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

Example

```
nmcli con modify intMgmt ipv4.addresses <hostmgmtIp/mask>
```

```
nmcli con modify intMgmt ipv4.gateway <mgmtgw>
```

```
nmcli con modify intMgmt ipv4.dns <dnsIp>
```

```
nmcli con modify intMgmt ipv4.dns-search 'cisco.com'
```

```
nmcli con modify intMgmt ipv4.method manual
```

```
nmcli con modify intMgmt ipv4.route-metric 50
```

d) Bring up the `intMgmt` network connection.

```
nmcli connection up intMgmt
```

Example

```
nmcli con up intMgmt
```

2. Repeat the above steps and create network bridge connection for data interface:
   a) Create a 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
```

   b) 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> master intData
```

Example

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

   c) Assign IP address and other details to `intData`.

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

Example

```
nmcli con modify intData ipv4.addresses <hostdataIp/mask>
```

```
nmcli con modify intData ipv4.method manual
```

```
nmcli con modify intData ipv4.gateway <datagw>
```

```
nmcli con modify intData ipv4.route-metric 90
```

d) Bring up the `intData` network connection.

```
nmcli connection up intData
```

Example

```
nmcli con up intData
```

Both network bridges, intMgmt and intData, are configured and active, enabling communication across associated network interfaces as if connected to the same physical network.

Configure SRIOV

Describes how to configure SRIOV so a physical network interface can be shared among multiple VMs.

Configure SRIOV.

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

Complete these steps to configure SRIOV.

Procedure

---

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

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

2. Set the number of VFs for the network interfaces according to your requirement. By default, two VFs are configured for each interface. You can 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
   ```

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

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

4. If any of the interfaces are configured for VLAN, assign VLAN IDs to the interfaces.

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

5. Save the changes and reboot the system.
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 184).

Install Crosswork Network Controller cluster on KVM using CLI

Ensure that

Install Crosswork Network Controller cluster on KVM using CLI.

Before you begin

Ensure that

- the KVM host meets all the requirements specified in KVM host bare metal requirements on page 170.
- the checks specified in Preliminary checks on page 179 is completed.
- your KVM environment is set up (refer to Set up and validate KVM on RHEL on page 179), and
- the network bridge or SRIOV is configured (refer to Configure network bridges or SRIOV on page 180).

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.

Procedure

1. As a first step, prepare the config ISO files for the Crosswork Network Controller cluster. You must create a separate config ISO file using xml (`ovf-env.xml`) for each VM in your cluster. For more information, refer to Crosswork Network Controller deployment templates for KVM on page 535.

   > **Note**
   >
   > 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, refer to Installation parameters on page 171.

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

   b) Generate the ISO 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.

2. Download the Crosswork Network Controller cluster qcow2 tar file and extract it twice (the file is double compressed).

   ```
   tar -xvf cnc-platform-cluster-deployment-7.2.1-qcow2.tar.gz
   ```

   This command extracts the tar file (`cnc-platform-cluster-deployment-7.2.1-123-qcow2.tar.gz`) and related checksum files.

   ```
   tar xvf cnc-platform-cluster-deployment-7.2.1-qcow2.tar.gz
   ./
   ./cnc-platform-cluster-deployment-7.2.1-123-qcow2.tar.gz
   ./cnc-platform-cluster-deployment-7.2.1-123-qcow2.tar.gz.signature
   ./CW_RELEASE.cer
   ./cisco_x509_verify_release.py3
   ./cisco_x509_verify_readme_gen.py3
   ./README
   ```

Extract it again: `tar -tvf cnc-platform-cluster-deployment-7.2.1-123-qcow2.tar.gz`

This command creates three qcow2 files:

```
cnc-platform-cluster-deployment-7.2.1-123_rootfs.qcow2
cnc-platform-cluster-deployment-7.2.1-123_dockerfs.qcow2
cnc-platform-cluster-deployment-7.2.1-123_extrafs.qcow2
```

3. Create an installation folder, preferably named after the VM hostname (for example, "node-1-Hybrid").
   a) Navigate to this folder.
      `cd node-1-Hybrid/`
   b) Copy the three extracted qcow2 files from the previous step into this folder.
   c) Create three disks using these commands:

```
qemu-img create -f qcow2 disk3 20G
qemu-img create -f qcow2 disk4 485G
qemu-img create -f qcow2 disk6 15G
```

   After these steps, the folder will contain the three qcow2 files and the three newly created disks.

```
ls -l

cnc-platform-cluster-deployment-7.2.1-123_dockerfs.qcow2
cnc-platform-cluster-deployment-7.2.1-123_extrafs.qcow2
cnc-platform-cluster-deployment-7.2.1-123_rootfs.qcow2
disk3
disk4
disk6
```

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=cnc-platform-cluster-deployment-7.2.1-123_rootfs.qcow2,format=qcow2,bus=scsi
 --disk
path=cnc-platform-cluster-deployment-7.2.1-123_dockerfs.qcow2,format=qcow2,bus=scsi
 --disk path=disk3,format=qcow2,bus=scsi --disk
path=disk4,format=qcow2,bus=scsi --disk
path=cnc-platform-cluster-deployment-7.2.1-123_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.qcow2,format=qcow2,bus=scsi --disk
 path=cw_dockerfs.qcow2,format=qcow2,bus=scsi --disk
path=disk3,format=qcow2,bus=scsi --disk path=disk4,format=qcow2,bus=scsi
--disk path=cw_extrafs.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.

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.

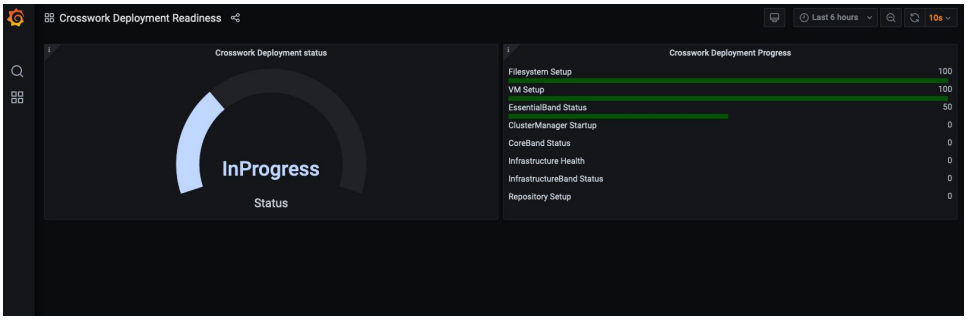
these is a list of critical steps in the process that you can watch for to be certain that things are progressing as expected:

1. The installer uploads the crosswork image file (.tar.gz file) to the data center.
2. The installer creates the VMs, and displays a success message (for example, "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 registed by Kubernetes to form the cluster.
4. Once the cluster is created and becomes accessible, a success message (for example, "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 these 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/NK1bwVxGk/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 31: 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 refer to if the nodes are ready) and `kubectl get pods` (to refer to the list of active running pods) commands.
  5. Repeat the `kubectl get pods` command until you refer to `robot-ui` in the list of active pods. For example, you may refer to `robot-ui-0` and `robot-ui-1` as active pods.
  6. At this point, you can try to access the Cisco Crosswork UI.

Failure Scenario

In the event of a failue 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.

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, refer to the *Compatibility Information* section in the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes*.

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

Procedure

1. Launch one of the supported browsers.
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, refer to the *Manage Certificates* topic in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

3. Log into the Cisco Crosswork using these steps:
   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 password that is at least eight characters long and includes uppercase and lowercase letters, numbers, and at least one special character. Although allowed, avoid using passwords similar to dictionary words (for example, "Pa55w0rd!") or other easily guessable terms.

The **Crosswork Manager** window is displayed.

4. Click on the **Crosswork Health** tab, and click the **Platform Infrastructure** tile to view the health status of the microservices running on Cisco Crosswork.
5. **Import inventory:** 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, refer to Crosswork Network Controller deployment templates for KVM on page 535. Then, import the file back into the Crosswork UI.

Troubleshoot the cluster

This section discusses the common troubleshooting scenarios for the cluster installation.

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.

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 52: 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	This happens when the password is not complex enough. Create a strong password, update the configuration manifest and redeploy. Use a strong VM Password (8 characters long, including upper & 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.
Deployment fails with: *Failed to validate Crosswork cluster initialization.*	The clusters' seed VM is either unreachable or one or more of the cluster VMs have failed to get properly configured. 1. Check whether the VM is reachable, and collect logs from `/var/log/firstBoot.log` and `/var/log/vm_setup.log` 2. Check the status of the other cluster nodes.

The VMs are deployed but the Crosswork cluster is not being formed.	A successful deployment allows the operator logging in to the VIP or any cluster hybrid node IP address to run these command to get the status of the cluster:
	
	`sudo kubectl get nodes`
	
	Note
	
	This command will not work on worker nodes.
	
	A healthy output for a 3-node cluster is:
	
	`NAME STATUS ROLES AGE VERSION`
	`172-25-87-2-hybrid.cisco.com Ready master 41d v1.33.0`
	`172-25-87-3-hybrid.cisco.com Ready master 41d v1.33.0`
	`172-25-87-4-hybrid.cisco.com Ready master 41d v1.33.0`
	
	In case of a different output, collect these logs: `/var/log/firstBoot.log` and `/var/log/vm_setup.log`
	
	In addition, for any cluster nodes not displaying the Ready state, collect:
	
	`sudo kubectl describe node <name of node>`

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

Dual stack scenarios

Table 53: Dual stack scenarios

Scenario	Possible Resolution
During deployment, these error message is displayed: *ERROR: No valid IPv6 address detected for IPv6 deployment.*	If you intend to use a dual stack configuration for your deployment, make sure that the host machine running the Docker installer meets these requirements: • 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 use the host's IPv6 network, use the parameter `--network host` when running the Docker installer. • Confirm that the provided IPv6 network CIDR and gateway are valid and reachable.
During deployment, these error message is displayed: *ERROR: seed v4 host empty*	Ensure you use the approved version of Docker installer (19 or higher) to run the deployment.
During deployment, these error message is displayed: *ERROR: Installation failed. Check installer and the VMs' log by accessing via console and viewing /var/log/firstBoot.log*	Common reasons for failed installation are: • Incorrect IPv4 or IPv6 Gateway IP for either Management or Data interfaces. • Unreachable IPv4 or IPv6 Gateway IP for either Management or Data interfaces. • Errors in mapping the datacenter networks in MgmtNetworkName and DataNetworkName parameters in .tfvars file Check the `firstBoot.log` file for more information, and contact Cisco Customer Experience team for any assistance.

Installation workflow for deploying SR-PCE (Cisco IOS XRv 9000) on RHEL 8.10 KVM

Lists the high-level workflow steps required to deploy and access the SR-PCE (Cisco IOS XRv 9000) virtual router on RHEL 8.10 KVM environments.

Crosswork Network Controller supports the deployment of SR-PCE running on the Cisco IOS XRv 9000 (XRv9k) platform in RHEL 8.10 KVM environments.

This section lists the necessary steps to launch and manage the XRv9k virtual router on RHEL 8.10 using KVM/QEMU and the libvirt virtualization API.

> **Note**
>
> Crosswork Network Controller deployment is supported on both RHEL 8.10 and RHEL 9.4 KVM.

Table 54: Workflow for deploying SR-PCE (Cisco IOS XRv 9000) on RHEL 8.10 KVM

Step	Action
1. Ensure that the host meets all the prerequisites.	Refer to Prerequisites and setup for Cisco IOS XRv 9000 virtualization on page 193.
2. Configure network bridges.	Refer to Configure network bridges for SR-PCE on page 194.
3. Prepare the VM storage.	Refer to Prepare VM storage on page 195.
4. Modify the `virsh` XML template to match your requirements	Refer to Update the VM configuration on page 195.
5. Deploy the VM and verify the status.	Refer to Deploy and verify the VM on page 196.
6. Connect to the router console.	Refer to Connect to the router console on page 196.

Prerequisites and setup for Cisco IOS XRv 9000 virtualization

Describes the prerequisites and host setup required to deploy SR-PCE on Cisco IOS XRv 9000 in RHEL KVM environments.

This topic describes the prerequisites, requirements, and procedures for enabling virtualization and preparing a host machine to deploy SR-PCE running on the Cisco IOS XRv 9000 (XRv9k) platform in RHEL 8.10 KVM environments.

Host machine requirements

Ensure that the host machine meets these requirements:

- Operating System: RHEL 8.10 with virtualization enabled in the BIOS/UEFI.
- Required Packages:
  - qemu-kvm
  - libvirt

- virt-install
- virt-manager
- bridge-utils

Enabling KVM

Verify that the host supports virtualization and that the necessary services are active.

1. Verify CPU virtualization support.

```
egrep -c '(vmx|svm)' /proc/cpuinfo
```

2. Check if the KVM modules are loaded.

```
lsmod | grep kvm
```

3. Install virtualization packages.

```
sudo dnf install -y qemu-kvm libvirt virt-install virt-manager bridge-utils
```

4. Start and enable the libvirt daemon.

```
sudo systemctl enable --now libvirtd
```

Required files for deployment

Obtain these files from the Cisco software portal:

- Cisco IOS XRv 9000 ISO Image (for example, `xrv9k-fullk9-x.vrr-25.3.1.iso`): The bootable image required for installation.
- Virsh XML Template: The configuration file used to define the VM properties.

Configure network bridges for SR-PCE

Explains how to create network bridges that let SR-PCE communicate with the management network and IGP topology.

You must create network bridges to allow the SR-PCE to communicate with the management network and the IGP topology.

This example creates `intData` for management and `br501` for IGP traffic (using VLAN 501 on interface `ens1f3`).

Procedure

1. Create the VLAN interface.

```
sudo nmcli connection add type vlan con-name ens1f3.501 ifname ens1f3.501 dev
 ens1f3 id 501 ipv4.method disabled ipv6.method ignore
```

2. Create the bridge.

```
sudo nmcli connection add type bridge con-name br501 ifname br501 ipv4.method
 disabled ipv6.method ignore
```

3. Add the VLAN interface as a bridge port.

```
sudo nmcli connection add type bridge-slave con-name ens1f3.501-port ifname
ens1f3.501 master br501
```

4. Bring the interfaces up.

```
sudo nmcli connection up br501
sudo nmcli connection up ens1f3.501
sudo nmcli connection up ens1f3.501-port
```

5. Verify the bridge mapping.

```
bridge link
ens1f3.501@ens1f3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 master br501
state forwarding priority 32 cost 2
```

Network bridges for SR-PCE are successfully created, enabling connectivity to both management and IGP networks.

Prepare VM storage

Outlines how to prepare VM storage, installation files, and permissions required for the hypervisor.

The XRv9k requires a dedicated virtual disk and access to the installation ISO. This topic describes creating a virtual disk, organizing the installation files within the host system's libvirt directory, and setting the appropriate ownership permissions for the Hypervisor to access the files.

Procedure

1. Create the virtual disk.

```
qemu-img create -f qcow2 pcedisk1.qcow2 64G
```

2. Organize files in the libvirt directory.

```
sudo mkdir -p /var/lib/libvirt/images/xrv9k
sudo cp pcedisk1.qcow2 /var/lib/libvirt/images/xrv9k/
sudo cp xrv9k-fullk9-x.vrr-25.3.1.iso /var/lib/libvirt/images/xrv9k/
```

3. Set permissions.

```
sudo chown qemu:qemu /var/lib/libvirt/images/xrv9k/pcedisk1.qcow2
sudo chown qemu:qemu
/var/lib/libvirt/images/xrv9k/xrv9k-fullk9-x.vrr-25.3.1.iso
```

Update the VM configuration

Explains how to update the virsh XML template so it matches your environment and hardware requirements.

Modify the `virsh` XML template to match your environment's file paths and hardware requirements.

Procedure

1. Set the source file path for the HDA disk and the CDROM.

```
<!-- HDA Disk -->
<disk type='file' device='disk'>
  <driver name='qemu' type='qcow2'/>
  <source file='/var/lib/libvirt/images/xrv9k/pcedisk1.qcow2'/>
  <target dev='vda' bus='virtio'/>
</disk>

<!-- CDROM -->
<disk type='file' device='cdrom'>
  <driver name='qemu' type='raw'/>
  <source file='/var/lib/libvirt/images/xrv9k/xrv9k-fullk9-x.vrr-25.3.1.iso'/>

  <target dev='hdc' bus='ide'/>
</disk>
```

2. Update the `<os>` section to enable UEFI boot.

```
<os>
  <type arch='x86_64' machine='pc'>hvm</type>
  <boot dev='hd'/>
<loader readonly='yes' secure='no'
type='pflash'>/usr/share/edk2/ovmf/OVMF_CODE.fd</loader>
  <nvram>/usr/share/edk2/ovmf/OVMF_VARS.fd</nvram>
  <bootmenu enable='yes'/>
</os>
```

3. Map the virtual interfaces to the bridges created in Configure network bridges for SR-PCE on page 194.

Deploy and verify the VM

Outlines how to deploy the VM by using an XML configuration file and verify its status.

This topic describes how to deploy the VM using a specified XML configuration file and verify the status of the VM.

Procedure

1. Create the VM.

```
virsh create xrv9k-config.xml
```

2. Verify the status of the VM.

```
virsh list --all
```

Connect to the router console

Explains how to connect to the IOS XR router console after the VM is running.

Once the VM is running, you can access the IOS XR console via telnet.

Procedure

1. Identify the console port by examining the serial section in your XML.

```
<serial type='tcp'>
   <source mode="bind" host="0.0.0.0" service="13914"/>
   <protocol type="telnet"/>
   <target port="0"/>
</serial>
```

2. Wait approximately 10-15 minutes for the router to boot. Then, run this command to connect to the console.

```
telnet localhost 13914
```

3. Configure the initial username, password, and system configurations as prompted.

8 Install Crosswork Data Gateway on KVM

Topics:

- Crosswork Data Gateway on KVM
- Installation workflow of Crosswork Data Gateway
- Parameters required for Crosswork Data Gateway installation
- Install Crosswork Data Gateway using the CLI
- Access and manage the Crosswork Data Gateways on KVM

Introduces the Crosswork Data Gateway installation and management guidance for KVM.

This chapter covers the installation workflow, required parameters, and configuration examples for Crosswork Data Gateway on KVM.

It also describes installing the gateway from the CLI and managing its virtual machines through Cockpit.

Crosswork Data Gateway on KVM

Crosswork Data Gateway collects network data and forwards telemetry and operational data to Crosswork Network Controller.

Crosswork Data Gateway enables seamless data collection and forwards telemetry and operational data from network devices to the Crosswork Network Controller. To activate these data collection capabilities when deploying the Crosswork Network Controller solution on a KVM hypervisor, you must install Crosswork Data Gateway as a critical initial step in your deployment workflow.

Installation workflow of Crosswork Data Gateway

Lists the KVM installation workflow, from environment checks through installation and monitoring.

Use this workflow to install Crosswork Data Gateway on KVM and monitor the installation.

Installation workflow

Follow these steps to complete the installation process.

Table 55: Installation workflow

Step	Action
1. Verify that your installation environment meets the prerequisites.	• Carefully review the prerequisites and confirm that your environment meets all the required specifications. Refer to Installation Prerequisites on page 43. • Review the installation parameters and ensure you have all relevant information readily available for use during the installation process. Refer to Parameters required for Crosswork Data Gateway installation on page 201.

Step	Action
2. Install Crosswork Data Gateway on the bare metal server.	Install Crosswork Data Gateway using a CLI-based approach. Refer to Install Crosswork Data Gateway using the CLI on page 216. **Note** If you plan to install multiple Data Gateway VMs due to load or scale requirements, or you wish to leverage Crosswork Data Gateway High Availability, we recommended to install all the Data Gateway VMs first and then proceed with adding them to a Data Gateway pool. refer to the *Create a Crosswork Data Gateway Pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
3. Monitor the health of the VM.	Once installation is complete, access Cockpit, a web-based management interface to monitor the health and performance of your VM. Refer to Access and manage the Crosswork Data Gateways on KVM on page 220.

Parameters required for Crosswork Data Gateway installation

Lists the parameters used in the Crosswork Data Gateway KVM installation configuration file.

This section outlines the parameters needed for installing Crosswork Data Gateway on KVM. These parameters will be included in the `config.txt` file used during the installation process.

Table 56: Parameters and description

Key	Description	Additional Information
`Deployment`	Parameter conveys the VM resource profile. For an on-premise installation, choose Crosswork On-Premise.	
Host Information		
`Description`	A detailed description of the Data Gateway.	
`Label`	Label used by Cisco Crosswork to categorize and group multiple Data Gateway VMs.	

Key	Description	Additional Information
`AllowRFC8190`	Choose how to validate interface addresses that fall in a usable RFC 8190 range. Options are: `Yes`, `No`, or `Ask`, where the initial configuration script prompts for confirmation. The default value is `Yes` to automatically allow interface addresses in an RFC 8190 range.	
`DGCertKey`	SCP URI to private key file for session key signing. You can retrieve this using SCP (`user@host:path/to/file`).	Cisco Crosswork uses self-signed certificates for handshake with Crosswork Data Gateway. These certificates are generated at installation.
`DGCertChainPwd`	Passphrase of the SCP user to retrieve the Crosswork Data Gateway PEM formatted certificate file and private key.	However, if you want to use third party or your own certificate files, then enter these parameters. Certificate chains override any preset or generated certificates in the Data Gateway VM and are given as an SCP URI (user:host:/path/to/file). The host with the URI files must be reachable on the network (from the vNIC0 interface via SCP) and the files must be present at the time of install.
`DGAppdataDisk`	Indicates the size in GB of a second data disk. The default value of this parameter in each profile is: • 20 GB for Standard. • 520 GB for Extended. Do not change the default value without consulting a Cisco representative.	
`HANetworkMode`	Indicates the mode for the high-availability network. Options are: • L2 • L3 The default value is `L2`.	

Passphrase

Key	Description	Additional Information
`dg-adminPassword`	The password you have chosen for the dg-admin user. Password must be 8-64 characters.	
`dg-operPassword`	The password you have chosen for the dg-oper user. Password must be 8-64 characters.	

Key	Description	Additional Information
`NicDefaultGateway`	The interface used as the default Data Gateway for processing the DNS and NTP traffic. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	For information on the type of roles that you must assign to the vNICs, refer to Table 36: Crosswork Data Gateway default vNIC deployment modes on page 76.
`NicAdministration`	The interface used to access the VM through the SSH access. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
`NicExternalLogging`	The interface used to send logs to an external logging server. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
`NicManagement`	The interface used to send the enrollment and other management traffic. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth0`.	
`NicControl`	The interface used to send the destination, device, and collection configuration. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth1`.	
`NicNBSystemData`	The interface used to send collection data to the system destination. As the system destinations share the same IP as the interface that allows connection to the collection service, the northbound data for system destinations uses the Control role's interface. Options are `eth0`, `eth1`, `eth2`or `eth3`.	
`NicNBExternalData`	The interface used to send the collection data to the external destinations configured by the user. Options are `eth0`, `eth1`, or `eth2`. In a 2-NIC deployment, the default interface is `eth1`; in a 3-NIC deployment, it is `eth2`.	
`NicSBData`		

Key	Description	Additional Information
	The interface used to collect data from the devices. If the interface only has the NicSBData role, it does not need an IP during the deployment. Options are `eth0`, `eth1`, or `eth2`. The default value is `eth2`.	
vNIC IPv4 address[8]		

`Vnic0IPv4Method`
`Vnic1IPv4Method`
`Vnic2IPv4Method`
	Method in which the interface is assigned an IPv4 address - `None` or `Static`.
	The default value is `None`.
		• If you are using IPv4, change the value from **none** to **static**, then configure these fields:
		  • vNIC IPv4 Address
		  • vNIC IPv4 Netmask
		  • vNIC IPv4 Skip Gateway
		  • vNIC IPv4 Gateway

`Vnic0IPv4Address`
`Vnic1IPv4Address`
`Vnic2IPv4Address`
	IPv4 address of the interface.	

`Vnic0IPv4Netmask`
`Vnic1IPv4Netmask`
`Vnic2IPv4Netmask`
	IPv4 netmask of the interface in dotted quad format.	• If you are using IPv6, leave the value set to **none** and retain the default IPv4 settings.

`Vnic0IPv4SkipGateway`
`Vnic1IPv4SkipGateway`
`Vnic2IPv4SkipGateway`
	The default value is `False`.
	Setting this to `True` skips configuring a gateway.	

`Vnic0IPv4Gateway`
`Vnic1IPv4Gateway`
`Vnic2IPv4Gateway`
	IPv4 address of the vNIC gateway.	

vNIC IPv6 address [9]		

[8] vNIC0, vNIC1, and vNIC2 based on the number of interfaces you choose to use.
[9] vNIC0, vNIC1, and vNIC2 based on the number of interfaces you choose to use.

Key	Description	Additional Information

`Vnic0IPv6Method`
`Vnic1IPv6Method`
`Vnic2IPv6Method`
	Method in which the vNIC interface is assigned an IPv6 address - `None`, `Static`, or `SLAAC`.
	The default value is `None`.
		If you are using IPv6, change the value from **none** to **static**, then configure these fields:
		• vNIC IPv6 Address
		• vNIC IPv6 Netmask
		• vNIC IPv6 Skip Gateway
		• vNIC IPv6 Gateway

`Vnic0IPv6Address`
`Vnic1IPv6Address`
`Vnic2IPv6Address`
	IPv6 address of the interface.	

`Vnic0IPv6Netmask`
`Vnic1IPv6Netmask`
`Vnic2IPv6Netmask`
	IPv6 prefix of the interface.	

`Vnic0IPv6SkipGateway`
`Vnic1IPv6SkipGateway`
`Vnic2IPv6SkipGateway`
	Options are `True` or `False`.
	Selecting `True` skips configuring a gateway.	

`Vnic0IPv6Gateway`
`Vnic1IPv6Gateway`
`Vnic2IPv6Gateway`
	IPv6 address of the vNIC gateway.	

DNS servers		
`DNSSEC`	Options are `False`, `True`, or `Allow-Downgrade`. The default value is `False` Select `True` to use DNS security extensions.	
`DNSTLS`	Options are `False`, `True`, and `Opportunistic`. The default value is `False`. Select `True` to use DNS over TLS.	
`mDNS`	Options are `False`, `True`, and `Resolve`. Select `True` to use multicast DNS. The default value is `False`.	If you choose Resolve, only resolution support is enabled. Responding is disabled.
`LLMNR`	Options are `False`, `True`, `Opportunistic`, or `Resolve`. The default value is `False`.	If you choose Resolve, only resolution support is enabled. Responding is disabled. Select `True` to use link-local multicast name resolution.
NTPv4 servers		

Key	Description	Additional Information
`NTPAuth`	Select `True` to use NTPv4 authentication. The default value is `False`.	
`NTPKey`	Key IDs to map to the server list. Enter a space-delimited list of Key IDs.	
`NTPKeyFile`	SCP URI to the chrony key file.	
`NTPKeyFilePwd`	Password of SCP URI to the chrony key file.	

Remote syslog server

Key	Description	Additional Information
`UseRemoteSyslog`	Options are `True` and `False`. Select `True` to send Syslog messages to a remote host. The default value is `False`.	Configuring an external syslog server sends service events (CLI/MDT/SNMP/gNMI) to the external syslog server. Otherwise, they are logged only to the Crosswork Data Gateway VM. If you want to use an external syslog server, specify these settings: • Use Remote Syslog Server • Syslog Server Address • Syslog Server Port • Syslog Server Protocol
`SyslogAddress`	Hostname, IPv4, or IPv6 address of a syslog server accessible in the management interface.	
`SyslogPort`	Port number of the syslog server. The default port number is 514.	
`SyslogProtocol`	Options are `UDP`, `RELP`, or `TCP` to send the syslog. The default value is `UDP`.	
`SyslogMultiserverMode`	Multiple servers in the failover or simultaneous mode. This parameter is applicable only when the protocol is set to a non-UDP value. UDP must use the simultaneous mode. Options are `Simultaneous` or `Failover`. The default value is `Simultaneous`.	
`SyslogTLS`	Select `True` to use TLS to encrypt syslog traffic. The default value is `False`.	
`SyslogPeerName`	Syslog server hostname exactly as entered in the server certificate SubjectAltName or subject common name.	
`SyslogCertChain`	PEM formatted root cert of syslog server retrieved using SCP. The host with the URI files must be reachable on the network (from vNIC0 interface via SCP) and the files must be present at the time of install.	
`SyslogCertChainPwd`	Password of SCP user to retrieve Syslog certificate chain.	

Remote auditd server

Key	Description	Additional Information
`UseRemoteAuditd`	Options are `True` and `False`. The default value is `False`. Select `True` to send auditd messages to a remote host.	If desired, you can configure an external Auditd server. Crosswork Data Gateway sends audit notifications to the Auditd server when it is configured and present on the network. Specify these three settings to use an external Auditd server.
`AuditdAddress`	Hostname, IPv4, or IPv6 address of an optional Auditd server.	
`AuditdPort`	Port number of an optional Auditd server. The default port is 60.	
Controller and proxy settings		
`ControllerIP`	The Virtual IP address or the hostname of the Cisco Crosswork cluster. **Note** If you are using an IPv6 address, it must be surrounded by square brackets ([1::1]). If geo redundancy is enabled, use the unified endpoint. For more information, refer to Unified endpoint requirements on page 433.	This is required so that the Data Gateway can enroll with the Crosswork server during the installation and initial start up. Excluding this step requires you to manually ingest the certificate. For more information, refer to Import controller signing certificate file on page 164.
`ControllerPort`	Port of the Cisco Crosswork controller. The default port is 30607.	

Key	Description	Additional Information
`ControllerSignCertChain`	PEM formatted root cert of Cisco Crosswork to validate signing certs retrieved using SCP. Cisco Crosswork generates the PEM file and is available at this location: `cw-admin@<Crosswork_VM_Management_VIP_Address>:/home/cw-admin/controller.pem` **Note** If you are using an IPv6 address, it must be surrounded by square brackets ([1::1]).	Crosswork Data Gateway requires the Controller Signing Certificate File to enroll automatically with Cisco Crosswork. If you specify these parameters during the installation, the certificate file is imported once Data Gateway boots up for the first time. If you do not specify these parameters during installation, then import the certificate file manually by following the procedure Import controller signing certificate file on page 164.
`ControllerTlsCertChain`	Cisco Crosswork Controller PEM formatted SSL/TLS certificate file retrieved using SCP.	
`ControllerCertChainPwd`	Password of SCP user (cw-admin) to retrieve Cisco Crosswork certificate chain.	
`ProxyURL`	URL of the HTTP proxy server.	The proxy parameters apply to the Crosswork Data Gateway cloud deployment. The Data Gateway must connect to the Internet via TLS, and a proxy server may be required if it is not present in your environment. If you want to use a proxy server, specify these parameters.
`ProxyBypass`	Comma-delimited list of addresses and hostnames that will not use the proxy server.	
`ProxyUsername`	Username for authenticated proxy servers.	
`ProxyPassphrase`	Passphrase for authenticated proxy servers.	
`ProxyCertChain`	HTTPS proxy PEM formatted SSL/TLS certificate file retrieved using SCP.	
`ProxyCertChainPwd`	Password of SCP user to retrieve proxy certificate chain.	
Geo redundancy settings		
`az_id`	The physical location of Availability Zone 1 and 2.	
`region_id`	The physical location of the Data Gateway VM.	

Key	Description	Additional Information
`site_location`	The location of the primary and second Crosswork sites. During enrollment, Crosswork sends this value to cdg-manager to preset the cluster affiliation of the instance.	

Note

- vNIC IPv4 or IPv6 addresses:
  - When using two or three NICs both vNIC0 and vNIC1 must be assigned static IPv4 or IPv6 addresses.
  - All unused vNICs (IPv4 or IPv6) should be left set to Method "None" with the other fields left at the default.
- Interfaces: In a 3-NIC deployment, you must provide an IP address for Management Traffic (vNIC0) and Control or Data Traffic (vNIC1). The IP address for Device Access Traffic (vNIC2) is assigned during Data Gateway pool creation as explained in the *Create a Crosswork Data Gateway Pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
- Although the IP address is assigned when the Data Gateway is added to a pool, you must configure network connectivity for the third interface either during installation or before adding the Data Gateway to the pool.
- The vNIC role assignment allows you to control the traffic that an interface must handle. If your use case is not supported by the preassigned roles, you can manually assign roles to the interfaces. Each parameter has a predefined role. The parameter accepts the interface value as eth0, eth1, or eth2, which refer to vNIC0, vNIC1, or vNIC2 respectively.

Sample configuration file for IPv4

Provides an IPv4 sample configuration for deploying Crosswork Data Gateway with three NICs on KVM.

The configuration file config.txt contains all parameters required to deploy Crosswork Data Gateway with 3 NICs on KVM.

Note

- The IPv4 and IPv6 configurations use almost the same parameters and structure. Some parameter values differ based on the IP protocol in use.
- In a 2-NIC deployment, `NicSBData` is set to `eth1`. In a 3-NIC deployment, it is set to `eth2`.
- The sample configuration shows mandatory parameters in **bold**. All IP addresses and host names are represented as placeholders. Replace them with valid IPv4 or IPv6 addresses as required by your environment.
- Values enclosed in angle brackets ("< parameter_description>") are placeholders. Replace them with values specific to your environment. Note that these are not mandatory parameters for deploying this configuration.

```
AllowRFC8190=Yes
AuditdAddress="<Auditd-server-address>"
AuditdPort=60
AwsIamRole="<Auditd-server-address>"
CloudEnrollmentToken="<Cloud-enrollment-token>"
ControllerCertChainPwd=controller_pwd
ControllerIP=controller_ip
ControllerPort=30607
ControllerSignCertChain=cw-admin@{controller_ip}:/home/cw-admin/controller.pem
ControllerTlsCertChain="<Controller-SSL/TLS-certificate-file-URI>"
Deployment=Crosswork On-Premise
Description=CDG Base VM for Automation
DGAppdataDisk=5
DGCertChain="<Data-gateway-certificate>"
DGCertChainPwd="<Data-gateway-certificate-password>"
DGCertKey="<Crosswork-data-gateway-private-key-URI>"
DNS=dns
DNSSEC=False
DNSTLS=False
Domain=domain
EnrollmentPassphrase="<Enrollment-passphrase >"
EnrollmentURI="<Enrollment-uri>"
HANetworkMode=L2
Hostname=host_name
Label="<Crosswork-data-label>"
LLMNR=False
mDNS=False
NicAdministration=eth0
NicControl=eth0
NicDefaultGateway=eth0
NicExternalLogging=eth0
NicManagement=eth0
NicNBExternalData=eth1
NicNBSystemData=data_interface
NicSBData=eth2
NTP=ntp
NTPAuth=False
NTPKey="<NTPv4-keys>"
NTPKeyFile="<NTPv4-key-file-uri>"
NTPKeyFilePwd="<NTPv4-key-file-passphrase>"
Profile=profile
ProxyBypass="<Proxy-server-bypass-list>"=
ProxyCertChain="<HTTPS-proxy-SSL/TLS-certificate-file-uri>"
```

```
ProxyCertChainPwd="<HTTPS-proxy-SSL/TLS-certificate-file-passphrase>"
ProxyPassphrase="<Authenticated-proxy-passphrase>"
ProxyURL="<Proxy-server-url>"
ProxyUsername="<Authenticated-proxy-username>"
SyslogAddress="<Syslog-server-address>"
SyslogCertChain="<Syslog-root-syslogCertChain-certificate-file-uri>"
SyslogCertChainPwd="<Syslog-certificate-file-phrase>"
SyslogMultiserverMode=simultaneous
SyslogPeerName="<Syslog-TLS-peer-name>"
SyslogPort=514
SyslogProtocol=UDP
SyslogTLS=False
UseRemoteAuditd=False
UseRemoteSyslog=False
Vnic0IPv4Address=mgmt_ip
Vnic0IPv4Gateway=mgmt_gateway
Vnic0IPv4Method=Static
Vnic0IPv4Netmask=IPv4_netmask
Vnic0IPv4SkipGateway=False
Vnic0IPv6Address=::0
Vnic0IPv6Gateway=::1
Vnic0IPv6Method=None
Vnic0IPv6Netmask=64
Vnic0IPv6SkipGateway=False
Vnic1IPv4Address=data_ip
Vnic1IPv4Gateway=data_gateway
Vnic1IPv4Method=Static
Vnic1IPv4Netmask=IPv4_netmask
Vnic1IPv4SkipGateway=False
Vnic1IPv6Address=IPv6_address
Vnic1IPv6Gateway=IPv6_gateway_address
Vnic1IPv6Method=None
Vnic1IPv6Netmask=64
Vnic1IPv6SkipGateway=False
Vnic2IPv4Address=IPv4_address
Vnic2IPv4Gateway=IPv4_gateway_address
Vnic2IPv4Method=None
Vnic2IPv4Netmask=IPv4_netmask_address
Vnic2IPv4SkipGateway=False
Vnic2IPv6Address=IPv4_address
Vnic2IPv6Gateway=IPv4_gateway_address
Vnic2IPv6Method=None
Vnic2IPv6Netmask=IPv4_netmask
Vnic2IPv6SkipGateway=False
Vnic3IPv4Address=IPv4_address
Vnic3IPv4Gateway=IPv4_gateway_address
Vnic3IPv4Method=None
Vnic3IPv4Netmask=IPv4_netmask_address
Vnic3IPv4SkipGateway=False
Vnic3IPv6Address=IPv4_address
Vnic3IPv6Gateway=IPv4_gateway_address
Vnic3IPv6Method=None
Vnic3IPv6Netmask=64
Vnic3IPv6SkipGateway=False
dg-adminPassword=admin_password
dg-operPassword=operator_password
```

Sample configuration file for IPv6

Provides an IPv6 sample configuration for deploying Crosswork Data Gateway with three NICs on KVM.

The configuration file config.txt contains all parameters required to deploy Crosswork Data Gateway with 3 NICs on KVM.

Note

- The IPv4 and IPv6 configurations use almost the same parameters and structure. Some parameter values differ based on the IP protocol in use.
- In a 2-NIC deployment, `NicSBData` is set to `eth1`. In a 3-NIC deployment, it is set to `eth2`.
- The sample configuration shows mandatory parameters in **bold**. All IP addresses and host names are represented as placeholders. Replace them with valid IPv4 or IPv6 addresses as required by your environment.
- Values enclosed in angle brackets ("< parameter_description>") are placeholders. Replace them with values specific to your environment. Note that these are not mandatory parameters for deploying this configuration.

```
AllowRFC8190=Yes
AuditdAddress="<Auditd-server-address>"
AuditdPort=60
AwsIamRole="<Amazon-web-services-IAM-role-name>"
CloudEnrollmentToken="<Cloud-enrollment-token>"
ControllerCertChainPwd=controller_pwd
ControllerIP=[controller_ip]
ControllerPort=30607
ControllerSignCertChain=cw-admin@[controller_ip]:/home/cw-admin/controller.pem
ControllerTlsCertChain="<Controller-SSL/TLS-certificate-file-URI>"
Deployment=Crosswork On-Premise
Description=cdg 165 ipv6 only
DGAppdataDisk=5
DGCertChain="<Data-gateway-certificate>"
DGCertChainPwd="<Data-gateway-certificate-password>"
DGCertKey="<Crosswork-data-gateway-private-key-URI>"
DNS=dns
DNSSEC=False
DNSTLS=False
Domain=domain
EnrollmentPassphrase="<Enrollment-passphrase >"
EnrollmentURI="<Enrollment-uri>"
HANetworkMode=L2
Hostname=host_name
Label="<Crosswork-data-label>"
LLMNR=False
mDNS=False
NicAdministration=eth0
NicControl=eth1
NicDefaultGateway=eth0
NicExternalLogging=eth0
NicManagement=eth0
NicNBExternalData=eth1
NicNBSystemData=eth1
NicSBData=eth2
NTP=ntp
NTPAuth=False
NTPKey="<NTPv4-keys>"
NTPKeyFile="<NTPv4-key-file-uri>"
NTPKeyFilePwd="<NTPv4-key-file-passphrase>"
Profile=Standard
```

```
ProxyBypass="<Proxy-server-bypass-list>"
ProxyCertChain="<HTTPS-proxy-SSL/TLS-certificate-file-uri>"
ProxyCertChainPwd="<HTTPS-proxy-SSL/TLS-certificate-file-passphrase>"
ProxyPassphrase="<Authenticated-proxy-passphrase>"
ProxyURL="<Proxy-server-url>"
ProxyUsername="<Authenticated-proxy-username>"
SyslogAddress="<Syslog-server-address>"
SyslogCertChain="<Syslog-root-syslogCertChain-certificate-file-uri>"
SyslogCertChainPwd="<Syslog-certificate-file-phrase>"
SyslogMultiserverMode=simultaneous
SyslogPeerName="<Syslog-TLS-peer-name>"
SyslogPort=514
SyslogProtocol=UDP
SyslogTLS=False
UseRemoteAuditd=False
UseRemoteSyslog=False
Vnic0IPv4Address=IPv4_address
Vnic0IPv4Gateway=IPv4_gateway_address
Vnic0IPv4Method=None
Vnic0IPv4Netmask=IPv4_netmask_address
Vnic0IPv4SkipGateway=False
Vnic0IPv6Address=mgmt_ip
Vnic0IPv6Gateway=mgmt_gateway
Vnic0IPv6Method=Static
Vnic0IPv6Netmask=64
Vnic0IPv6SkipGateway=False
Vnic1IPv4Address=IPv4_address
Vnic1IPv4Gateway=IPv4_gateway_address
Vnic1IPv4Method=None
Vnic1IPv4Netmask=IPv4_netmask_address
Vnic1IPv4SkipGateway=False
Vnic1IPv6Address=data_ip
Vnic1IPv6Gateway=data_gateway
Vnic1IPv6Method=Static
Vnic1IPv6Netmask=64
Vnic1IPv6SkipGateway=False
Vnic2IPv4Address=IPv4_address
Vnic2IPv4Gateway=IPv4_gateway_address
Vnic2IPv4Method=None
Vnic2IPv4Netmask=IPv4_netmask_address
Vnic2IPv4SkipGateway=False
Vnic2IPv6Address=::0
Vnic2IPv6Gateway=::1
Vnic2IPv6Method=None
Vnic2IPv6Netmask=64
Vnic2IPv6SkipGateway=False
Vnic3IPv4Address=IPv4_address
Vnic3IPv4Gateway=IPv4_gateway_address
Vnic3IPv4Method=None
Vnic3IPv4Netmask=IPv4_netmask_address
Vnic3IPv4SkipGateway=False
Vnic3IPv6Address=::0
Vnic3IPv6Gateway=::1
Vnic3IPv6Method=None
Vnic3IPv6Netmask=64
Vnic3IPv6SkipGateway=False
dg-adminPassword=Data_gateway_admin_password
dg-operPassword=Data_gateway_operator_password
```

Install Crosswork Data Gateway using the CLI

Provides the CLI procedure for installing Crosswork Data Gateway on a KVM host.

Install Crosswork Data Gateway on a KVM host so it can collect and transport network data to Crosswork Network Controller.

You must install Crosswork Data Gateway on KVM to enable it to collect and transport network data to the Crosswork Network Controller. Repeat the installation process for each Data Gateway that you must set up.

Guideline to configure high availability in Crosswork Data Gateway

If you plan to deploy multiple Crosswork Data Gateway VMs either to meet load and scale requirements or to enable HA, follow this sequence:

1. Install all Data Gateway VMs.
2. Add the VMs to a Data Gateway pool.

Refer to the *Create a Crosswork Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Before you begin

Ensure you are familiar with the requirements and have the necessary information ready:

- Review the required parameters and gather all necessary information. Refer to Parameters required for Crosswork Data Gateway installation on page 201.
- The KVM environment has specific configuration requirements, important considerations, and limitations for installation. Refer to Preliminary checks on page 179.

Procedure

1. On the KVM server, create a user directory, such as cdg_deploy, in the `/home` folder. This directory acts as the central location for storing all deployment-related files.

```
mkdir -p /home/cdg-deploy
```

2. Navigate to the newly created directory.

```
cd /home/cdg-deploy/
```

3. Download the Crosswork Data Gateway installation file from cisco.com, then extract it.

```
tar -zxvf cdg-deployment-7.2.1.bios.qcow2.tar.gz
cdg-deployment-7.2.1.bios.qcow2
config.txt
config-onprem.txt
config-cloud.txt
```

   The first command extracts the release verification file and the QCOW2 tar files. The second command extracts the QCOW2 images and the config.txt file. Modify the config.txt file to specify the required values. Refer to Sample configuration file for IPv4 on page 211 and Sample configuration file for IPv6 on page 213.

4. Create a folder named **cdg_name**.

```
mkdir -p cdg_name
```

5. Copy the config.txt file, extracted in Step 3, to the cdg_deploy folder.

```
cp cdg-cloud-deployment-7.2.1.bios.qcow2 cdg_name
vi config.txt
```

a) If you want to deploy multiple Crosswork Data Gateways in a single deployment, run:

```
cp cdg-cloud-deployment-7.2.1.bios.qcow2 cdg_name
cp config.txt $cdg_name
cd $cdg_name
vi config.txt
```

> **Note**
>
> In `config.txt`, modify only the values required for your deployment. Do not modify the default values provided for the remaining parameters. If changes are necessary, contact the Cisco Customer Experience team for guidance.

6. Create a data disk file with the required size (in GB) to allocate disk space for Crosswork Data Gateway. The required size depends on the deployment profile. Refer to Resource footprint for KVM on page 26.

```
qemu-img create -f qcow2 {cdg_name}-datadisk {data_disk}
```

Example

```
qemu-img create -f qcow2 sample_datagateway-datadisk 24
```

7. Create the ISO file.

```
mkisofs -R -relaxed-filenames -joliet-long -iso-level 3 -l -o {cdg_name}.iso
 config.txt
```

8. Replace the parameter values in the config.txt file as needed.

   The parameters include:

   - `ram_value`: The required RAM for installing Crosswork Data Gateway depends on the deployment profile you select.
   - `no_of_cpu`: The number of CPUs required for installing Crosswork Data Gateway depends on the deployment profile you choose.
   - `qcow2_file_name`: The name of the `qcow2` file extracted in Step 3.
   - `cdg_name`: The name of Crosswork Data Gateway you want to install.

   For a complete list of parameters, refer to Parameters required for Crosswork Data Gateway installation on page 201.

9. Install Crosswork Data Gateway.

   - To use network bridges, perform these steps:
     - 2-NIC deployment:

```
virt-install --boot hd,cdrom --connect qemu:///system --virt-type kvm
--name cdg33 --ram ${ram_value} --vcpus {no_of_cpu} --cpu
```

```
host-passthrough --disk path={qcow2_file_name},format=qcow2 --disk
path={cdg_name}-datadisk,format=qcow2 --disk={cdg_name}.iso,device=cdrom
  --os-variant ubuntu-lts-latest --import --network
bridge={intMgmt},model=virtio --network bridge={intData},model=virtio
&
```

Example:

```
virt-install \
  --boot hd,cdrom \
  --connect qemu:///system \
  --virt-type kvm \
  --name cdg33 \
  --ram 8192 \
  --vcpus 4 \
  --cpu host-passthrough \
  --disk path=cdg33.qcow2,format=qcow2 \
  --disk path=cdg33-datadisk,format=qcow2 \
  --disk=cdg33.iso,device=cdrom \
  --os-variant ubuntu-lts-latest \
  --import \
  --network bridge=br-int-mgmt,model=virtio \
  --network bridge=br-int-data,model=virtio &
```

- 3-NIC deployment:

```
virt-install --boot hd,cdrom --connect qemu:///system --virt-type kvm
--name cdg33 --ram ${ram_value} --vcpus {no_of_cpu} --cpu
host-passthrough --disk path={qcow2_file_name},format=qcow2 --disk
path={cdg_name}-datadisk,format=qcow2 --disk={cdg_name}.iso,device=cdrom
  --os-variant ubuntu-lts-latest --import --network
bridge={intMgmt},model=virtio --network bridge={intData},model=virtio
--network bridge={intDevices},model=virtio &
```

Example:

```
virt-install \
  --boot hd,cdrom \
  --connect qemu:///system \
  --virt-type kvm \
  --name cdg33 \
  --ram 8192 \
  --vcpus 4 \
  --cpu host-passthrough \
  --disk path=cdg33.qcow2,format=qcow2 \
  --disk path=cdg33-datadisk.qcow2,format=qcow2 \
  --disk path=cdg33.iso,device=cdrom \
  --os-variant ubuntu-lts-latest \
  --import \
  --network bridge=br-int-mgmt,model=virtio \
  --network bridge=br-int-data,model=virtio \
  --network bridge=br-int-devices,model=virtio &
```

- To use SRIOV, perform these steps:
  - 2-NIC deployment:

```
virt-install --boot hd,cdrom --connect qemu:///system --virt-type kvm
--name cdg33 --ram
${ram_value} --vcpus {no_of_cpu} --cpu host-passthrough --disk
```

```
path={qcow2_file_name},format=qcow2 --disk
path={cdg_name}-datadisk,format=qcow2
--disk={cdg_name}.iso,device=cdrom --os-variant ubuntu-lts-latest
--import --os-variant ubuntu-lts-latest –import
--host-device=pci_0000_17_0e_1 --host-device=pci_0000_17_06_1 &
```

Example:

```
virt-install \
  --boot hd,cdrom \
  --connect qemu:///system \
  --virt-type kvm \
  --name cdg33 \
  --ram 8192 \
  --vcpus 4 \
  --cpu host-passthrough \
  --disk path=cdg33.qcow2,format=qcow2 \
  --disk path=cdg33-datadisk.qcow2,format=qcow2 \
  --disk path=cdg33.iso,device=cdrom \
  --os-variant ubuntu-lts-latest \
  --import \
  --host-device=pci_0000_17_0e_1 \
  --host-device=pci_0000_17_06_1 &
```

- 3-NIC deployment:

```
virt-install --boot hd,cdrom --connect qemu:///system --virt-type kvm
--name cdg33 --ram
${ram_value} --vcpus {no_of_cpu} --cpu host-passthrough --disk
path={qcow2_file_name},format=qcow2 --disk
path={cdg_name}-datadisk,format=qcow2
--disk={cdg_name}.iso,device=cdrom --os-variant ubuntu-lts-latest
--import --os-variant ubuntu-lts-latest –import
--host-device=pci_0000_17_0e_1 --host-device=pci_0000_17_06_1
--host-device=pci_0000_17_02_1 &
```

Example:

```
virt-install \
  --boot hd,cdrom \
  --connect qemu:///system \
  --virt-type kvm \
  --name cdg33 \
  --ram 8192 \
  --vcpus 4 \
  --cpu host-passthrough \
  --disk path=cdg33.qcow2,format=qcow2 \
  --disk path=cdg33-datadisk.qcow2,format=qcow2 \
  --disk path=cdg33.iso,device=cdrom \
  --os-variant ubuntu-lts-latest \
  --import \
  --host-device=pci_0000_17_0e_1 \
  --host-device=pci_0000_17_06_1 \
  --host-device=pci_0000_17_02_1 &
```

Note

If the VNC console for the Data Gateway is not available on the KVM server after deploying Data Gateway, append these parameters to the end of the command and redeploy the Data Gateway.

```
--graphics vnc,listen=0.0.0.0 --video virtio --noautoconsole
```

For example,

```
virt-install --boot hd,cdrom --connect qemu:///system --virt-type kvm
--name railtel-cdg-2 --ram 114688 --vcpus 20 --cpu host-passthrough --disk
path=cw-na-dg-7.2.1.bios.qcow2,format=qcow2 --disk
path=railtel-cdg-2-datadisk,format=qcow2
--disk=railtel-cdg-2.iso,device=cdrom --os-variant ubuntu-lts-latest
--import --network bridge=intMgmt,model=virtio --network
bridge=intData,model=virtio --network bridge=intSouthbound,model=virtio
--graphics vnc,listen=0.0.0.0 --video virtio --noautoconsole
```

What to do next

Access the Cockpit UI to view the health details of the VM. Refer to Access and manage the Crosswork Data Gateways on KVM on page 220.

Access and manage the Crosswork Data Gateways on KVM

Use the Cockpit UI to view installed Crosswork Data Gateways and monitor VM health and performance.

Use the Cockpit UI to view installed Crosswork Data Gateway VMs and monitor their health and performance.

After installing Crosswork Data Gateway on KVM, monitor the health and performance of your VM from the Cockpit UI. This UI allows you to manage VMs and perform administrative tasks.

Procedure

1. Install Cockpit on the server.

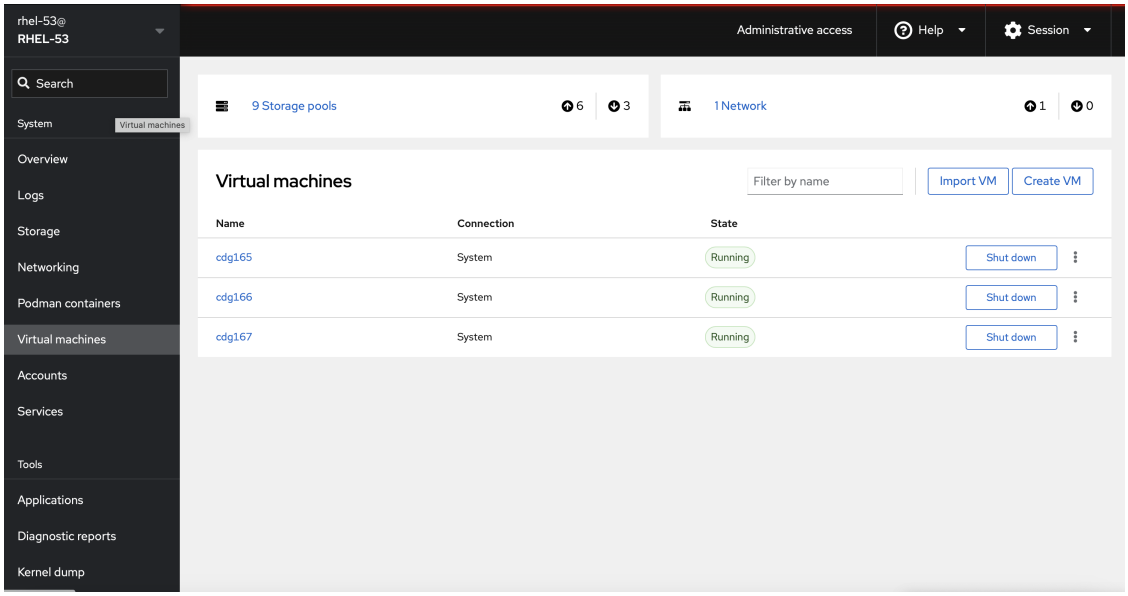
   ```
   sudo dnf install cockpit
   ```

2. Enable the Cockpit socket.

   ```
   systemctl enable --now cockpit.socket
   ```

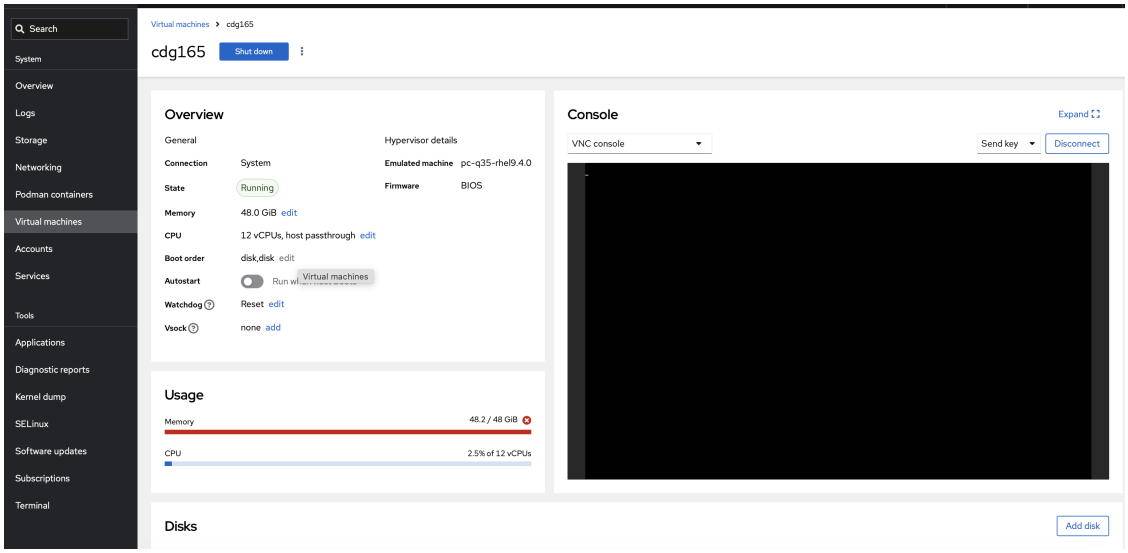
3. Access the Cockpit UI at `https://<bare-metal-ip>:9090`.
4. In the Cockpit UI, select **Virtual machines** from the left menu. The **Virtual machines** pane displays the virtual machines where Crosswork Data Gateway is installed.

   Figure 32: Cockpit UI



5. Select the VM name to view its connection, state, memory, and other health details.

Figure 33: Virtual machines details



For more information about Cockpit, refer to the Cockpit documentation.

The Cockpit UI displays the installed Crosswork Data Gateway VMs. Selecting a VM displays its connection, state, memory, and other health details.

Part IV

Install Cisco Crosswork Network Controller on AWS EC2

Topics:

- Install Cisco Crosswork Network Controller on AWS EC2

9 Install Cisco Crosswork Network Controller on AWS EC2

Topics:

- Installation overview
- Amazon EC2 requirements for Crosswork
- Extract CF template image
- Roles and policy permissions
- CloudFormation (CF) template parameters
- Module-based deployment method
- CloudFormation template deployment and management
- Access the Crosswork UI
- Crosswork Data Gateway post-installation tasks

Introduces the AWS EC2 installation requirements and deployment process for Cisco Crosswork Network Controller.

This chapter describes the AWS EC2 requirements, CloudFormation parameters, deployment tasks, and post-installation guidance for Cisco Crosswork Network Controller.

It also includes the Data Gateway post-installation and enrollment guidance relevant to the AWS installation.

Installation overview

Explains installation overview, including the source-described role, behavior, and relationship to AWS EC2 installation and deployment.

This section provides an overview of how Cisco Crosswork is installed on Amazon EC2

Cisco Crosswork uses the CloudFormation (CF) templates to deploy the cluster stacks. The CF process is faster and less error-prone than the manual procedure to build the cluster, however you must have the necessary skills to prepare a CF template with details of the cluster deployment.

Note

The terms 'stack' and 'instance' refers to cluster and VM respectively.

Note

The CF templates (.yaml file) provided are samples that must be customized according to your production preferences and executed as per the steps mentioned in this chapter.

Amazon EC2 requirements for Crosswork

Provides reference information about amazon EC2 requirements for Crosswork, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section describes the settings that must be configured to install Crosswork Network Controller on Amazon EC2

Crosswork can be deployed in Amazon Elastic Compute Cloud (EC2) Amazon EC2 is a web service that provides compute resources in the cloud to host your Crosswork applications.

Crosswork is deployed in Amazon EC2 using CloudFormation (CF) templates. The CloudFormation process is faster and less error-prone than the manual procedure to build the cluster, however you must have the necessary skills to prepare a CloudFormation template with details of the cluster deployment.

Installing Crosswork and its components in the AWS environment requires you to review and meet these prerequisites:

> **Note**
>
> Most of the requirements discussed in this section are AWS concepts and not imposed exclusively by Crosswork.

Table 57: AWS Prerequisites and Settings

Requirement	Description
VPC and Subnets	Virtual Private Cloud (VPC) is created and configured with dedicated subnets for Crosswork interfaces (Management and Data) and Crosswork Data Gateway (Management, Data, and Device) interfaces Direct IP connectivity is required between all subnets.
Endpoints	An endpoint is created in your VPC with these parameters: • **Service name:** EC2 • **Private DNS names:** Enabled • **Endpoint type:** Interface • Under **Subnets**, specify the management subnet that you intend to use for the installation. If you are using different management subnets for the Crosswork VM and the data gateway VM **Note** For information on how to configure the endpoints, refer to the AWS documentation.

Requirement	Description
IAM role	A role is created in Identity and Access Management (IAM) with relevant permission policies. An IAM role is an identity that has specific permissions with credentials that are valid for short durations. Roles can be assumed by entities that you trust. **Note** • The minimum permissions required for a Crosswork role are **ec2:DescribeNetworkInterfaces**, **ec2:AssignPrivateIpAddresses** and **ec2:UnassignPrivateIpAddresses**. • The trust policy for your role must have the **"Action": "sts:AssumeRole"** condition.
Key pairs	Key pairs (private keys used to log into the VMs) are created and configured.
Placement Groups	A placement group of *Cluster* strategy is created. In a *cluster* placement group, instances are logically grouped in a single availability zone that benefit from low network latency and high network throughput. This requirement is required only for launching the Crosswork cluster instances.
IP addresses	**Crosswork cluster**: When using dual NICs (one for the Management network and one for the Data network), you require a management and data IP address (IPv4 or IPv6) for each node being deployed (Hybrid or Worker) and two additional IP addresses to be used as the management and data Virtual IP (VIP) address. For example, in the case of a 3 VM cluster with dual NIC, you need 8 IP addresses (4 for management network and 4 for data network). **Crosswork Data Gateway**: IP addresses for Management Traffic and Data Traffic only. IP address for Device Access Traffic is assigned during the data gateway pool creation. • The IP addresses must be able to reach the gateway address for the network where Crosswork Data Gateway will be installed, or the installation fails. • At this time, your IP allocation is permanent and cannot be changed without redeployment. For more information, contact the Cisco Customer Experience team.

Requirement	Description
Security group	A security group must be created and configured to specify which ports or traffic are allowed.
Instance type	The resource profile for your instance deployment. The AWS Instance type should be selected to conform with the VM resource and network requirements listed in Plan Your Deployment on page 17. • **Crosswork Cluster**: • Select **m5.4xlarge** for demos or lab deployments. • Select **m5.8xlarge** for production deployments. • **Crosswork Data Gateway** (production and lab deployments): • **Standard** - Select **m5.4xlarge** • **Extended** - Select **m5.8xlarge**
CloudFormation (CF) template	The CF template (.yaml) files for the Crosswork components that must be uploaded during the installation. For more information, see Extract CF template image on page 229.
Route53DomainName	Domain name configured for Route53 DNS hosted zone.
User data	The VM-specific parameters script that must be specified during the manual installation procedure.
Hosted Zone ID	The Hosted Zone ID must be provided with the domain name (Route53DomainName). The Network Load Balancer (NLB) deployments require a predefined Route53 hosted zone.

Extract CF template image

Use this procedure to extract CF template image, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Extract CF template image.

This section explains the procedure to extract and validate the Crosswork Network Controller CF template image.

Note

- The file names mentioned in this topic are sample names and may differ from the actual file names in release version.
- The image files (AMI files) for the deployment are not posted on cisco.com, but are instead copied to the customer's AWS account based on request. For more information, please contact your Cisco Account team.

Procedure

1. Download the Crosswork Network Controller CF template package (**cnc-cluster-cft-7.2.1.tar.gz cnc-cluster-cft-7.2.0-releasecnc720-18.tar.gz**) from cisco.com.
2. Use this command to unzip the package:

   ```
   tar xzvf <package name> -C {specify a new directory name}
   ```

   Example:

   ```
   tar xzvf cnc-cluster-cft-7.2.1.tar.gz -C {specify a new directory name}
   ```

   The contents of the package is unzipped to the specified directory. This directory contains the CF template image and files necessary to validate the image.

   For example:

   ```
   tar xzvf cnc-cluster-cft-7.2.1.tar.gz -C {specify a new directory name}
   x cnc-cluster-cft-7.2.1-releasecnc720-87.tar.gz
   x cnc-cluster-cft-7.2.1-releasecnc720-87.tar.gz.signature
   x README
   x CW-CCO_RELEASE.cer
   x cisco_x509_verify_release.py3
   x cisco_x509_verify_release.py
   ```

3. Review the contents of the README file in order to understand everything that is in the package and how it will be validated in these steps.
4. Navigate to the directory created in the previous step and use this command to verify the signature of the installer image:

   Note

   Use `python --version` to find out the version of Python on your machine.

   If you are using Python 2.x, use this command to validate the file:

   ```
   python cisco_x509_verify_release.py -e <.cer file> -i <.tar.gz file> -s
   <.tar.gz.signature file>
   -v dgst -sha512
   ```

If you are using Python 3.x, use this command to validate the file:

```
python cisco_x509_verify_release.py3 -e <.cer file> -i <.tar.gz file> -s
<.tar.gz.signature file>
-v dgst -sha512
```

For example:

```
python cisco_x509_verify_release.py3 -e CW-CCO_RELEASE.cer -i
cnc-cluster-cft-7.2.1-releasecnc720-87.tar.gz -s
cnc-cluster-cft-7.2.1-releasecnc720-87.tar.gz.signature -v dgst -sha512
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 CW-CCO_RELEASE.cer.
Successfully verified the signature of
cnc-cluster-cft-7.2.1-releasecnc720-87.tar.gz using CW-CCO_RELEASE.cer
```

The contents of the package is extracted and validated successfully.

5. In the directory, locate the install-cnc-templates file and follow the instructions provided within its **Description** section.

   Customize the CF templates in the directory to install Crosswork Network Controller on Amazon EC2

What to do next

Return to the installation workflow: Install Cisco Crosswork Network Controller on AWS EC2 on page 38

Roles and policy permissions

Provides reference information about roles and policy permissions, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

Provides reference information about roles and policy permissions, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section describes the roles and the policy permissions that you must have when deploying the CF template on Amazon. For information on how to create and manage the roles, refer to the Amazon documentation.

Table 58: Amazon EC2 Roles and Actions Assigned to the Roles

Role	Actions
EC2	DescribeInternetGateways, DescribeNetworkInterfaces, DescribeImages, DeleteLaunchTemplate, DescribeSubnets, DescribeAccountAttributes, DescribeSecurityGroups, RunInstances, DescribeVpcs, DescribeInstances, CreateNetworkInterface, CreateTags, DescribeKeyPairs, CreateLaunchTemplate, DeleteNetworkInterface, TerminateInstances

Role	Actions
ELB	DescribeLoadBalancers, CreateLoadBalancer, ModifyLoadBalancerAttributes, AddTags, DeleteLoadBalancer
ELB v2	DescribeLoadBalancers, CreateLoadBalancer, AddTags, DeleteLoadBalancer, CreateTargetGroup, CreateListener, DeleteListener, DescribeTargetGroups, ModifyLoadBalancerAttributes, DescribeListeners, RegisterTargets, DeleteTargetGroup, ModifyTargetGroupAttributes, DescribeTargetHealth
IAM	CreateNodegroup, DescribeNodegroup, DeleteNodegroup

CloudFormation (CF) template parameters

Provides reference information about cloudFormation (CF) template parameters, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

Provides reference information about cloudFormation (CF) template parameters, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section explains the important parameters that must be specified for module deployments.

- CF template parameters for installing Cisco Crosswork cluster VMs on page 232
- Crosswork Data Gateway on page 15
- CF template parameters for installing a single hybrid cluster or worker node on page 241

Note

- The parameters that are mandatory for creating the templates are indicated explicitly. Parameters without this indication are optional and are populated with the default values, which you can alter based on your deployment requirement.
- All IP addresses you enter as parameters should be available.

CF template parameters for installing Cisco Crosswork cluster VMs

Provides reference information about cF template parameters for installing Cisco Crosswork cluster VMs, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

Provides reference information about cF template parameters for installing Cisco Crosswork cluster VMs, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section describes the parameters that are required for deploying Cisco Crosswork Cluster VMs with 3 hybrid VMs on Amazon EC2. It also describes the Management and Data NLB parameters.

Once you have determined the subnet for your cluster nodes and any other virtual machines you are going to deploy, confirm that there are enough available IP addresses to support the number of VMs (and virtual IP addresses) needed.

Table 59: Cisco Crosswork Cluster VMs Deployment Parameters

Parameter	Description
`VpcId`	The virtual private cloud (VPC) ID of your existing VPC. For example, vpc-0f83aac74690101a3.
`SecGroup`	Precreated security group that must be applied to the stack. For example, sg-096ff4bc355af16a0. The group must allow ingress access to ports 22, 30160:31560.
`CwSSHPassword`	The SSH password of the Crosswork Network Controller. **Note** We recommend using an external secret store for the password.
`CwAmiId`	The Crosswork AMI ID. This is a mandatory parameter.
`CwMgmtSubnet1Id`	Management subnet for Crosswork VM 1. This is a mandatory parameter.
`CwMgmtSubnet2Id`	Management subnet for Crosswork VM 2. This is a mandatory parameter.
`CwMgmtSubnet3Id`	Management subnet for Crosswork VM 3. This is a mandatory parameter.
`CwMgmtSubnet1Netmask`	The first management subnet netmask in the dotted-decimal form. For example, 255.255.255.0. This parameter is ignored when deploying on a single interface. This is a mandatory parameter.
`CwMgmtSubnet2Netmask`	The second management subnet netmask in the dotted-decimal form. For example, 255.255.255.0. This parameter is ignored when deploying on a single interface. This is a mandatory parameter.
`CwMgmtSubnet3Netmask`	The third management subnet netmask in the dotted-decimal form. For example, 255.255.255.0. This parameter is ignored when deploying on a single interface. This is a mandatory parameter.
`CwMgmtSubnet1Gateway`	The management default gateway on the selected data subnet. Typically, the first address on the subnet. This parameter is ignored when deployed on single interface mode. This is a mandatory parameter.

Parameter	Description
`CwMgmtSubnet2Gateway`	The management default gateway on the selected data subnet. Typically, the first address on the subnet. This parameter is ignored when deployed on single interface mode. This is a mandatory parameter.
`CwMgmtSubnet3Gateway`	The management default gateway on the selected data subnet. Typically, the first address on the subnet. This parameter is ignored when deployed on single interface mode. This is a mandatory parameter.
`ManagementVIPName`	Crosswork Management VIP name. For example, dev1-cwmgnt. This will be the host name to access the Crosswork cluster.
`DataVIPName`	Crosswork Data VIP name. For example, dev1-cwdata.
`Route53DomainName`	Domain name used for all Route53 objects. This is the DNS domain name for the Crosswork cluster. This is a mandatory parameter.
`HostedZoneId`	The Hosted Zone ID provided with the domain name (Route53DomainName). The Network Load Balancer (NLB) deployments require a predefined Route53 hosted zone. This is a mandatory parameter.
`UseExternalNLB`	Determines whether to use an external NLB for the Crosswork cluster (multi-AZ or subnet) or a Crosswork VIP (only single AZ or subnet). Options are `True` or `False`. This is a mandatory parameter.
`CwClusterPlacementStrategy`	The EC2 instance placement strategy that is valid for single availability zone. Default 'cluster' ensures maximum throughput. Options are: • `cluster` • `partition` • `spread`
`CwNodeType`	The Crosswork Node Type for deployment. Options are `Hybrid` or `Worker`. A replacement Hybrid node must reuse the same IP addresses as the Hybrid node it is replacing. Default value is `Worker`. This is a mandatory parameter.

Parameter	Description
`InterfaceDeploymentMode`	The deployment mode. Options are `1` to deploy the Management interface or `2` to deploy the Management and Data interface.
`CwDataSubnet1Id`	Data subnet of Crosswork VM 1. In a single interface, the deployments happen on the subnet where the Management interface is deployed. This is a mandatory parameter.
`CwDataSubnet2Id`	Data subnet of Crosswork VM 2. In a single interface, the deployments happen on the subnet where the Management interface is deployed. This is a mandatory parameter.
`CwDataSubnet3Id`	Data subnet of Crosswork VM 3. In a single interface, the deployments happen on the subnet where the Management interface is deployed. This is a mandatory parameter.
`CwDataSubnet1Netmask`	The first data subnet netmask in the dotted-decimal form. For example, `255.255.255.0`. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
`CwDataSubnet1Gateway`	The first default data gateway on the selected data subnet. Typically, the value is the first address on the subnet. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
`CwDataSubnet2Netmask`	The second data subnet netmask in the dotted-decimal form. For example, `255.255.255.0`. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
`CwDataSubnet2Gateway`	The second data subnet netmask in the dotted-decimal form. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
`CwDataSubnet3Netmask`	The third data subnet netmask in the dotted-decimal form. For example, `255.255.255.0`. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
`CwDataSubnet3Gateway`	The third data subnet netmask in the dotted-decimal form. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.

Parameter	Description
`CwMgmtVIP`	The current Crosswork Management VIP address.
`CwDataVIP`	The current Crosswork Data VIP address. When using an external NLB, you can leave this parameter empty.
`Cw1MgmtIP`	A free address on the management subnet. If not specified, an address is automatically assigned.
`Cw1DataIP`	A free address on the data subnet. If not specified, an address is automatically assigned.
`Cw2MgmtIP`	A free address on the management subnet. If not specified, an address is automatically assigned .
`Cw2DataIP`	A free address on the data subnet. If not specified, an address is automatically assigned.
`Cw3MgmtIP`	A free address on the management subnet. If not specified, an address is automatically assigned.
`Cw3DataIP`	A free address on the data subnet. If not specified, an address is automatically assigned .
`OtherCwMgmtIP1`	The Management IP address \#1 of the existing Crosswork node. This is used when the deployment happens with an external load balancer.
`OtherCwMgmtIP2`	The Management IP address \#2 of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.
`OtherCwDataIP1`	The Data IP address \#1 of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.
`OtherCwDataIP2`	The Data IP address \#2 of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.

Table 60: Crosswork VM Customization

Parameter	Description
`InstanceType`	The EC2 instance type for the node instances. This is a mandatory parameter.
`RunAsSpotInstance`	A spot instance. Options are: • `True`: to enable the feature • `False`: to disable the feature Default value is `False`. This is a mandatory parameter.

Parameter	Description
DataDiskSize	Crosswork data disk size. The default is 450 GB and should be fine for most deployments. Enter the default unless otherwise directed by Cisco Customer Experience team. This is a mandatory parameter.
K8sServiceNetwork	The network address for the Kubernetes service network. The CIDR range is fixed to '/16'. If not provided, the default will be taken, that is, 10.96.0.0. This is a mandatory parameter.
K8sPodNetwork	The network address for the Kubernetes pod network. The CIDR range is fixed to '/16'." This is a mandatory parameter.
SkipAutoInstall	Configures the Skip Auto Install feature. Options are: • True: to enable the feature • False: to disable the feature Default value is False. This is a mandatory parameter.

Table 61: Cisco Crosswork Cluster Management NLB Deployment Parameters

Parameter	Description
VpcId	The virtual private cloud (VPC) ID of your existing VPC. For example, vpc-0f83aac74690101a3.
CwTargetSubnetIdList	This is a list of the Crosswork management subnets. This is a mandatory parameter.
CwTargetIP1	This is a Crosswork VM management IP. In this template, this is a mandatory parameter.
CwTargetIP2	This is a Crosswork VM management IP. In this template, this is a mandatory parameter.
CwTargetIP3	This is a Crosswork VM management IP. In this template, this is a mandatory parameter.
Route53DomainName	Domain name used for all Route53 objects. This is a mandatory parameter.
HostName	The domain name used for all Route53 objects. This is a mandatory parameter.
HostedZoneId	The hosted zone ID. This is a mandatory parameter.

Table 62: Data NLB Deployment Parameters

Parameter	Description
VpcId	The virtual private cloud (VPC) ID of your existing VPC. For example, vpc-0f83aac74690101a3.
CwTargetSubnetIdList	The first management subnet for the Crosswork VMs. This is a mandatory parameter.
CwTargetIP1	A free address on the management subnet. If not specified, an address is automatically assigned.
CwTargetIP2	A free address on the management subnet. If not specified, an address is automatically assigned.
CwTargetIP3	A free address on the management subnet. If not specified, an address is automatically assigned.
Route53DomainName	Domain name used for all Route53 objects. This is a mandatory parameter.
HostName	The domain name used for all Route53 objects. This is a mandatory parameter.
HostedZoneId	The hosted zone ID. This is a mandatory parameter.

CF template parameters for installing Crosswork Data Gateway

Provides reference information about cF template parameters for installing Crosswork Data Gateway, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

Provides reference information about cF template parameters for installing Crosswork Data Gateway, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section describes the parameters that are required when creating the data gateway control plane, node, pool, and other important containers. It also has parameters that are required for creating EC2 Crosswork Data Gateway NLB stack.

Table 63: Crosswork Data Gateway Deployment Parameters

Parameter	Description
AwsIamRole	The Amazon Web Services IAM role name for the EC2 VIP update.
VpcId	The virtual private cloud (VPC) ID of your existing VPC. For example, vpc-0f83aac74690101a3.
SecGroup	A pre-created security group that must be applied to the stack. For example, sg-096ff4bc355af16a0. The group must allow ingress access to all ports that Crosswork, NSO, Crosswork Data Gateway, and IOS-XR uses.
CDGSSHPassword	The SSH password to be configured on the data gateway node.

Parameter	Description
CDGOperPassword	The password to be configured on data gateway for Dg-Oper user.
CDGAmiId	The Crosswork Data Gateway AMI ID.
InstanceType	The EC2 instance type for the node instances. This is a mandatory parameter.
CNCControllerIP	Host address or name of the data gateway controller. In a multi-AZ deployment, this value must be the name. This is a mandatory parameter.
CNCControllerPassword	The cw-admin user password used to access Crosswork or CNC Controller.
InterfaceDeploymentMode	Crosswork Data Gateway deployment mode. The options are: • 1: to deploy all the interfaces. • 2: to deploy the Management and Data interfaces. • 3: to deploy the Management, Data, and Control interfaces.
CDGInterface0IPAddress	A free IP address on the subnet. If set to 0.0.0.0, the IP address is automatically allocated. This is a mandatory parameter.
CDGInterface0SubnetId	The first interface subnet for the data gateway VM.
CDGInterface0Gateway	The default gateway on the selected subnet. Typically, the first address on the subnet.
CDGInterface0SubnetNetmask	The first interface subnet netmask in the dotted-decimal form. For example, 255.255.255.0. This is a mandatory parameter.
CDGInterface1IPAddress	A free IP address on the first subnet. If set to 0.0.0.0, the IP address is automatically allocated. This is a mandatory parameter.
CDGInterface1SubnetId	The second interface subnet for the data gateway. The subnet must be in the same availability zone as the CDGInterface0SubnetId.
CDGInterface1Gateway	The second interface default gateway on the selected subnet. Typically, the first address on the subnet. This is a mandatory parameter.

Parameter	Description
`CDGInterface1SubnetNetmask`	The second interface subnet netmask in the dotted-decimal form. For example, `255.255.255.0`. This parameter is ignored when dual interface mode is not used. This is a mandatory parameter.
`CDGInterface2IPAddress`	A free IP address on the second subnet. If set to `0.0.0.0`, the IP address is automatically allocated. This is a mandatory parameter.
`CDGInterface2SubnetId`	The third interface subnet for the data gateway VM. The subnet must be in the same availability zone as the `CDGInterface0SubnetId`.
`CDGInterface2Gateway`	The third interface default gateway on the selected subnet. Typically, the first address on the subnet. This is a mandatory parameter.
`CDGInterface2SubnetNetmask`	The third interface subnet netmask in the dotted-decimal form. For example, `255.255.255.0`. This parameter is ignored when triple interface mode is not used. This is a mandatory parameter.
`CNCControllerIP`	Host address of the data gateway controller.
`HANetworkMode`	The data gateway HA mode. The pool mode options are: • `L2`: Use this option when you specify IP addresses for creating the HA pool. • `L3`: Use this option when you specify FQDN for creating the HA pool and for multi-subnet deployment. This is a mandatory parameter.
`DataDiskSize`	Size of the Crosswork data disk. The minimum size is 20. Default size is 50. This is a mandatory parameter.
`CDGProfile`	The deployment profile of data gateway. • `Standard` • `Extended` This is a mandatory parameter.
`CdgInstanceHostname`	The data gateway instance name, for example CDG-01.
`az_id`	The physical location of Availability Zone 1 and 2.
`region_id`	The physical location of the data gateway VM.

Parameter	Description
site_location	The location of the primary and second Crosswork sites. During enrollment, Crosswork sends this value to cdg-manager to preset the cluster affiliation of the instance.

Table 64: Crosswork Data Gateway and Network Load Balancer (NLB) Stack Parameters

Parameter	Description
VpcId	The VPC ID of the worker instances. This is a mandatory parameter.
SubnetId1	The management ID of subnet 1. This is a mandatory parameter.
SubnetId2	The management ID of subnet 2. This is a mandatory parameter.
DomainName	The domain name. This is a mandatory parameter.
HostedZoneId	The hosted zone ID. This is a mandatory parameter.
CdgPoolHostname	Name of the Route53 record. This is a mandatory parameter.
CdgTargetIP1	The IP address 1 of the Management node. In the event of a single data gateway, one target IP must be configured.
CdgTargetIP2	The IP address 2 of the Management node.
LBIPaddress1	The first LB IP address on subnet. This is a mandatory parameter.
LBIPaddress2	The second LB IP address on subnet. This is a mandatory parameter.

CF template parameters for installing a single hybrid cluster or worker node

Provides reference information about cF template parameters for installing a single hybrid cluster or worker node, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

Provides reference information about cF template parameters for installing a single hybrid cluster or worker node, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section describes the parameters that are required for deploying a single cluster node (Hybrid or Worker).

Note

- A replacement hybrid node must reuse the same IP addresses as the hybrid VM it is replacing.
- As you will be adding another node (worker or hybrid) to the existing cluster determine the subnet that is being used and find an additional available IP on that subnet.

Table 65: Single Hybrid Cluster or Worker Cisco Crosswork Nodes Deployment Parameters

Parameter	Description
`VpcId`	The virtual private cloud (VPC) ID of your existing VPC. For example, vpc-0f83aac74690101a3.
`SecGroup`	A pre-created security group that must be applied to the stack. For example, sg-096ff4bc355af16a0. The group must allow ingress access to ports 22, 30160:31560.
`EC2ENIRole`	Existing role name for the Crosswork cluster. The role must permit EC2 access.
`CwAmiId`	The Crosswork AMI ID. This is a mandatory parameter.
`CwSSHPassword`	The SSH password of the Crosswork Network Controller. **Note** We recommend using an external secret store for the password.
`InstanceType`	The EC2 instance type for the node instances. This is a mandatory parameter.
`ManagementVIPName`	Crosswork Management VIP name. For example, dev1-cwmgnt.
`DataVIPName`	Crosswork Data VIP name. For example, dev1-cwdata.
`Route53DomainName`	Domain name used for all Route53 objects. This is a mandatory parameter.
`UseExternalNLB`	Determines whether to use an external NLB for the Crosswork cluster (multi-AZ or subnet) or a Crosswork VIP (only single AZ or subnet). Options are `True` or `False`. This is a mandatory parameter.
`CwMgmtSubnetId`	The management subnet for the Crosswork VMs.

Parameter	Description
CwMgmtSubnetNetmask	The management subnet netmask in dotted decimal form. For example, 255.255.255.0. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
CwDataSubnetGateway	The management default gateway on the selected data subnet. Typically, the first address on the subnet. This parameter is ignored when deployed on single interface mode. This is a mandatory parameter.
CwDataSubnetId	The data subnet for the Crosswork VMs.
CwDataSubnetNetmask	The data subnet netmask in dotted decimal form. For example, 255.255.255.0. This parameter is ignored when deploying in a single interface mode. This is a mandatory parameter.
CwDataSubnetGateway	The data default gateway on the selected data subnet. Typically, the first address on the subnet. This parameter is ignored when deployed on single interface mode. This is a mandatory parameter.
CwNodeType	The Crosswork Node Type for deployment. Options are Hybrid or Worker. A replacement Hybrid node must reuse the same IP addresses as the Hybrid node it is replacing. This is a mandatory parameter.
DataDiskSize	Crosswork data disk size. The default is 450 (in GB) and should be fine for most deployments. Enter the default unless otherwise directed by Cisco Customer Experience team. This is a mandatory parameter.
K8sServiceNetwork	The network address for the Kubernetes service network. The CIDR range is fixed to '/16'. If not provided, the default (10.96.0.0) is taken.
K8sPodNetwork	The network address for the Kubernetes pod network. The CIDR range is fixed to '/16'. If not provided, the default (10.244.0.0) is taken.

Table 66: Optional VM parameters

Parameters	Description
CwMgmtVIP	The current Crosswork Management VIP address.
CwDataVIP	The current Crosswork Data VIP address. When using an external NLB, you can leave this parameter empty.

Parameters	Description
`Cw1MgmtIP`	A free address on the management subnet. If not specified, an address is automatically assigned.
`Cw1DataIP`	A free address on the data subnet. If not specified, an address is automatically assigned.
`OtherCwMgmtIP1`	The first Management IP address of the existing Crosswork node. This is used when the deployment happens with an external load balancer.
`OtherCwMgmtIP2`	The second Management IP address of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.
`OtherCwDataIP1`	The first Data IP address of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.
`OtherCwDataIP2`	The second Data IP address of the existing Crosswork node. This parameter is used when the deployment happens with an external load balancer.

Module-based deployment method

Provides reference information about module-based deployment method, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

The module-based deployment procedures involve deploying each resource separately. Each resource has its own template file, which can be used to deploy them individually. For more information, refer to these topics:

- Cisco Crosswork cluster deployment on Amazon on page 244
- Crosswork Data Gateway deployment on Amazon on page 245
- Additional Crosswork cluster node deployment on page 250

Cisco Crosswork cluster deployment on Amazon

Provides reference information about cisco Crosswork cluster deployment on Amazon, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section provides an overview of how Cisco Crosswork cluster is installed on Amazon EC2.

Cisco Crosswork uses a set of CF templates to deploy Crosswork cluster.

Crosswork Cluster Deployment Workflow

The Crosswork cluster deployment procedure involves deploying various Crosswork resources using the corresponding YAML files.

Table 67: Resources Deployed During Crosswork Cluster Deployment

Resource	Description
EC2 Cluster	The main stack (**cw-cluster.yaml**) which will deploy other nested stacks for creating EC2 CW NLBs.
Management NLB	The **cw-mgmt-nlb.yaml** file creates Network Load Balancer, Target Groups, Listeners and Route53Record for EC2 CW Management Nodes.
Data NLB	The **cw-data-nlb.yaml** file creates Network Load Balancer, Target Groups, Listeners and Route53Record for EC2 CW Data Nodes

Installation Parameters

For list of important parameters that you must specify in the CF templates that are used to deploy Crosswork cluster, see CF template parameters for installing Cisco Crosswork cluster VMs on page 232 EC2

> **Note**
>
> Once you have determined the subnet for your cluster nodes and any other virtual machines you are going to deploy, confirm that there are enough available IP addresses to support the number of VMs (and virtual IP addresses) needed.

Deploy the CF Templates

You can install the Crosswork cluster on Amazon EC2 Crosswork Cluster Deployment Workflow on page 244.

For instructions on how to deploy the CF templates on Amazon EC2 Deploy a CF template on page 251.

Verify the Installation

Verify that the Crosswork cluster installation is successful by following the steps in Monitor the installation on page 252.

Deploy an Additional Crosswork Cluster Node

For instructions on how to deploy an additional worker/hybrid node on the Crosswork cluster, see Additional Crosswork cluster node deployment on page 250.

What to do next

Return to the installation workflow: Install Cisco Crosswork Network Controller on AWS EC2 on page 38

Crosswork Data Gateway deployment on Amazon

Provides reference information about crosswork Data Gateway deployment on Amazon, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section provides an overview of how Crosswork Data Gateway is installed on Amazon EC2

Crosswork Data Gateway Deployment Workflow

Crosswork Data Gateway deployment procedure involves deploying various Crosswork resources using the corresponding YAML files.

The main file **cdg-stack-ec2.yaml** deploys the stacks for a single Data Gateway NLB (**cdg-nlb.yaml**) and two Data Gateways (**cdg.yaml**).

- An additional Data Gateway VM to the Data Gateway high availability pool is deployed using the **cdg.yaml** file. For each additional VM deployment, you must repeat the deployment procedure.
- An additional NLB and Data Gateway high availability pool is deployed using the **cdg-nlb.yaml** file.

This table provides information about the components are installed:

Table 68: Resources Deployed During Crosswork Data Gateway Deployment

Resource	Description
EC2 Crosswork Data Gateway	The resources related to EC2 node are created by deploying the **cdg.yaml** file.
Crosswork Data Gateway Network Load Balancer	The EC2 NLB components (target groups, network load balancer, data listeners, and NLB route 53 record) are created by deploying the **cdg-nlb.yaml** file.

Installation Parameters

For list of important parameters in the Crosswork Data Gateway CF templates, see Crosswork Data Gateway on page 15

The Data Gateway is deployed on Amazon EC2

Deploy the CF Templates

For instructions on how to deploy the CF templates on Amazon EC2 Deploy a CF template on page 251.

> **Note**
>
> Amazon EC2 mandates entering an IP address for the vNIC2 interface when Data Gateway is deployed using 3 NICs. This is an AWS EC2 requirement and not imposed by Crosswork.

Verify the Installation

Verify that the Crosswork Data Gateway installation is successful by following the steps in Monitor the installation on page 252.

What to do next

Return to the installation workflow: Install Cisco Crosswork Network Controller on AWS EC2 on page 38

Auto-configuration for deploying Crosswork Data Gateway

Provides reference information about auto-configuration for deploying Crosswork Data Gateway, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

The auto-configuration procedure discovers the configuration parameters that are missing, and it automatically defines the mandatory parameters to install Base VM. The configuration parameters are passed using the Dynamic Host Configuration Protocol (DHCP) framework. In the Day 0 configuration, the auto-configuration mechanism defines only the essential parameters with the default values.

A default password is provided during the auto-configuration to comply with the security policies. On the initial log in, the dg-admin and dg-oper users must change the default password. The Crosswork Data Gateway services are inactive until the default password is changed.

The auto-configuration process supports the default 3 NIC deployment. In particular, only eth0 is configured for the Management network.

The DHCP interaction takes place over the eth0 interface. The auto-configuration procedure uses the default values stored on the DHCP server. After Base VM is deployed, you can configure or modify the default values using the Interactive Console. For more information about the console, refer to *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

> **Note**
>
> The auto-configuration mechanism is not supported for deploying data gateway on the VMware platform.

How does Auto-configuration Work

1. Install the Base VM.
2. Crosswork Data Gateway searches for the Day 0 `config.txt` file.
3. If the `config.txt` file is not found, the Crosswork Data Gateway starts the auto-configuration process, and the required parameters are set with the default values. In the background, the Crosswork Data Gateway sends a request to the DHCP server to obtain the default parameter values. Once the response is received, the parameters are configured accordingly.

Parameters used during Auto-Configuration

The auto-configuration utility configures these parameters with the default values. For more information about these parameters, see Crosswork Data Gateway parameters and deployment scenarios on page 131.

Table 69: Cisco Crosswork Data Gateway Mandatory Deployment Parameters

Name	Parameter	Default Value
AllowRFC8190	`AllowRFC8190`	The default value is `Yes`.
Auditd Server Port	`AuditdPort`	The default port is `60`.
Crosswork Controller Port	`ControllerPort`	The default port is `30607`.
Description	`Description`	The default value is `CDG auto configure`.

Name	Parameter	Default Value
dg-admin Passphrase	`dg-adminPassword`	The default password is `changeme`. Reset the default value with the password that you have chosen for the dg-admin user. Password must be 8-64 characters.
dg-oper Passphrase	`dg-operPassword`	The default password is `changeme`. Reset the default value with the password you have chosen for the dg-oper user. Password must be 8-64 characters.
Data Disk Size	`DGAppdataDisk`	The default value of this parameter is `5`.

DNS Address	`DNS`	The default values of this parameter are
		`208.67.222.222`
		`208.67.220.220`

DNS Security Extensions	`DNSSEC`	The default value of this parameter is `False`.
DNS over TLS	`DNSTLS`	The default value of this parameter is `False`.
DNS Search Domain	`Domain`	The default value of this parameter is `localdomain`.
Crosswork Data Gateway HA mode	`HANetworkMode`	The default value of this parameter is `L2`.
Hostname	`Hostname`	The default value of this parameter is `dg-<eth0 address>`. Where `<eth0-address>` is the address of vNIC0.
Link-Local Multicast Name Resolution	`LLMNR`	The default value of this parameter is `False`.
Multicast DNS	`mDNS`	The default value of this parameter is `False`.
NicAdministration	`NicAdministration`	The default value of this parameter is `eth0`.
NicControl	`NicControl`	The default value of this parameter is `eth1`.
NicDefaultGateway	`NicDefaultGateway`	The default value of this parameter is `eth0`.

Name	Parameter	Default Value
NicExternalLogging	`NicExternalLogging`	The default value of this parameter is `eth0`.
NicManagement	`NicManagement`	The default value of this parameter is `eth0`.
NicNBExternalData	`NicNBExternalData`	The default value of this parameter is `eth1`.
NicNBSystemData	`NicNBSystemData`	The default value of this parameter is `eth1`.
NicSBData	`NicSBData`	The default value of this parameter is the last active interface such as eth0 for 1-NIC deployment, eth1 for 2-NIC.

NTPv4 Servers	`NTP`	The default values of this parameter are
		`162.159.200.1`
		`65.100.46.164`
		`40.76.132.147`
		`104.131.139.195`

Use NTPv4 Authentication	`NTPAuth`	The default value of this parameter is `False`.
Profile	`Profile`	The default value of this parameter is `Standard`.
Syslog Multiserver Mode	`SyslogMultiserverMode`	The default value of this parameter is `Simultaneous`.
Syslog Server Port	`SyslogPort`	The default value of this parameter is `514`.
Syslog Server Protocol	`SyslogProtocol`	The default value of this parameter is `UDP`.
Use Syslog over TLS	`SyslogTLS`	The default value of this parameter is `False`.
Use Remote Auditd Server	`UseRemoteAuditd`	The default value of this parameter is `False`.
Use Remote Syslog Server	`UseRemoteSyslog`	The default value of this parameter is `False`.
vNIC IPv4 Method	`Vnic0IPv4Method`	The default value of this parameter is `DHCP`.
vNIC IPv4 Skip Gateway	`Vnic0IPv4SkipGateway`	The default value of this parameter is `False`.

Name	Parameter	Default Value
vNIC IPv6 Method	`Vnic0IPv6Method`	The default value is `None`.
vNIC IPv6 Skip Gateway	`Vnic0IPv6SkipGateway`	The default value is `False`.
vNIC IPv4 Method	`Vnic1IPv4Method`	The default value is `DHCP`.
vNIC IPv4 Skip Gateway	`Vnic1IPv4SkipGateway`	The default value is `False`.
vNIC IPv6 Method	`Vnic1IPv6Method`	The default value is `None`.
vNIC IPv6 Skip Gateway	`Vnic1IPv6SkipGateway`	The default value is `False`.
vNIC IPv4 Method	`Vnic2IPv4Method`	The default value is `DHCP`.
vNIC IPv4 Skip Gateway	`Vnic2IPv4SkipGateway`	The default value is `False`.
vNIC IPv6 Method	`Vnic2IPv6Method`	The default value is `None`.
vNIC IPv6 Skip Gateway	`Vnic2IPv6SkipGateway`	The default vale is `False`.

Additional Crosswork cluster node deployment

Provides reference information about additional Crosswork cluster node deployment, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

This section explains how to deploy an additional worker/hybrid node on the Crosswork cluster.

Deploying an additional node on the Crosswork cluster involves deploying the Crosswork network configuration and VM customization resources using the **cw-add-vm.yaml** file.

> **Note**
>
> Before deploying an additional worker node, ensure that the Crosswork cluster and Crosswork application have been created.

> **Note**
>
> A new hybrid node MUST reuse the same IP addresses as the hybrid VM it is replacing, and a maximum of 3 hybrid nodes are allowed.

Installation Parameters

For list of important parameters that you must specify in the CF template that is used to deploy an additional node on the Crosswork cluster, see CF template parameters for installing a single hybrid cluster or worker node on page 241

Deploy the CF Templates

You can install an additional worker/hybrid node on the Crosswork cluster by customizing the CF template.

For instructions on how to deploy the CF templates on Amazon EC2 Deploy a CF template on page 251.

Verify the Installation

Verify that the nodes are attached to the Crosswork cluster. On the EC2 **Compute** section. For more information, see Monitor the installation on page 252.

What to do next

Return to the installation workflow: Install Cisco Crosswork Network Controller on AWS EC2 on page 38

CloudFormation template deployment and management

Explains cloudFormation template deployment and management, including the source-described role, behavior, and relationship to AWS EC2 installation and deployment.

these sections explain how to deploy a CF template on Amazon EC2

- Deploy a CF template on page 251
- Monitor the installation on page 252

Deploy a CF template

Use this procedure to deploy a CF template, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Deploy a CF template.

You can install Crosswork on Amazon EC2

Before you begin

- Make sure that you have met the Resource Requirements and Prerequisites prescribed for installing Crosswork on Amazon EC2
- Ensure that you have access to the CloudFormation templates that are stored in the S3 bucket or on your local machine. If the template is in Amazon S3, keep the URL of the template file copied.

Procedure

1. Log in to the AWS account and navigate to the S3 bucket. If the CF template is on your local computer, you can upload the template.
2. In the AWS CloudFormation console, navigate to the **Stacks** page and choose **Create stack > With new resources (standard)**. The **Create stack** page opens.
3. Enter these details:
   a. Under **Prerequisite – Prepare template**, select **Template is ready**.
   b. Under **Specify template > Template source**, select one of these options:
      - If you have the YAML or JSON file URL directing to the S3 bucket where the CF template is located, select **Amazon S3 URL**. In the **Amazon S3 URL** field, enter the URL and click **Next**.
      - If the CF template is saved on your local computer, select **Upload a template file** and click **Choose File** to select the file that you want to upload. After you have selected the template, Amazon uploads the file and displays the S3 URL. Click **Next**.

> **Note**
>
> (Optional) Click **View in Designer** to view a visual representation of the execution flow in your CF template.

4. In the **Specify stack details** page, enter the relevant values for the stack name and parameter values. Click **Next**.

> **Note**
>
> The parameter field names visible in this window are defined by the parameters in the CF template.

5. Review the parameter values that you have configured.
6. Under the **Capabilities**, select the check boxes next to:
   - **I acknowledge that AWS CloudFormation might create IAM resources with custom names.**
   - **I acknowledge that AWS CloudFormation might require this capability: CAPABILITY_AUTO_EXPAND.**
7. Click **Submit**.

What to do next

The time taken to create the cluster can vary based on the size of your deployment profile and the performance characteristics of your hardware. See Monitor the installation on page 252 to know how you can check the status of the installation.

Monitor the installation

Use this procedure to monitor the installation, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Monitor the installation.

This section describes how to verify if the deployment is complete without errors.

Procedure

---

1. In the CloudFormation console, from the left-hand side **Stacks** pane, select the stack that you have deployed.
2. The stack details are displayed on the right. Click on each tab in this window to view details of the stack. If the stack creation is in progress, the status of the stack in the **Events** tab is CREATE_IN_PROGRESS.
3. After the stack is created:
   - The status of the stack changes to CREATE_COMPLETE and the **Logical ID** displays the stack name.
   - The **Resources** tab displays details of the all the resources that the CF template has created, including the physical IDs.

- The **Outputs** tab has details of the VM's interface IP addresses.

What to do next

After the stack creation is complete, you can access the Crosswork UI and monitor the health of your cluster. For more information on how to log in to the Crosswork UI, see Access the Crosswork UI on page 253.

Access the Crosswork UI

Use this procedure to access the Crosswork UI, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Access the Crosswork UI.

After the stacks are created, you can check if all the nodes are up and running in the cluster from the Cisco Crosswork UI.

Before you begin

- Ensure that you have a spare Network Load Balancer (NLB). To access Crosswork UI, use an external NLB that routes requests to its targets using the protocol DNS and port number that you specify.
- Verify that the Crosswork cluster and pods are in the running state. For information on how to view the status of the cluster, see Monitor the installation on page 252.
- Make sure to keep the IP address of the Management node copied. This IP address is used to access the Crosswork UI. You can copy the IP address from the **Outputs** tab of the CloudFormation console. For information on accessing the **Outputs** tab, see Monitor the installation on page 252.

Procedure

1. Log in to the AWS console and navigate to **Target Groups** to register the targets.
2. Under **Targets**, click **Register targets**. The **Register targets** page opens.
3. In the Specify the port as 30603. Click **IPv4 address**, specify the Management IP address that you copied from the CloudFormation console. **Include as pending below**
4. Click **Register pending targets**.

   To deregister the targets that are no longer in use, select the target and click **Deregister**.
5. After the target is in the healthy state, click on the load balancer name under **Details**. The **Load balancer** page opens.
6. Copy the DNS name from the **DNS name** column.
7. Launch a supported browser and enter this URL in the address bar: `https://<DNS_name>:30603/`

> **Note**
>
> When you access 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 Cisco Crosswork server. After you add a security exception, the browser accepts the server as a trusted site in all future login attempts.

8. Log in to Cisco Crosswork :
   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.

9. (Optional) Click on the **Crosswork Health** tab, and click on the Crosswork Infrastructure tile to view the health status of the microservices running on Cisco Crosswork.

What to do next

Return to the installation workflow: Install Cisco Crosswork Network Controller on AWS EC2 on page 38

Crosswork Data Gateway post-installation tasks

Explains crosswork Data Gateway post-installation tasks, including the source-described role, behavior, and relationship to AWS EC2 installation and deployment.

This section lists the steps that you can complete after you have deployed Crosswork Data Gateway.

Configure the time zone of the Data Gateway VM

Use this procedure to configure the time zone of the Data Gateway VM, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

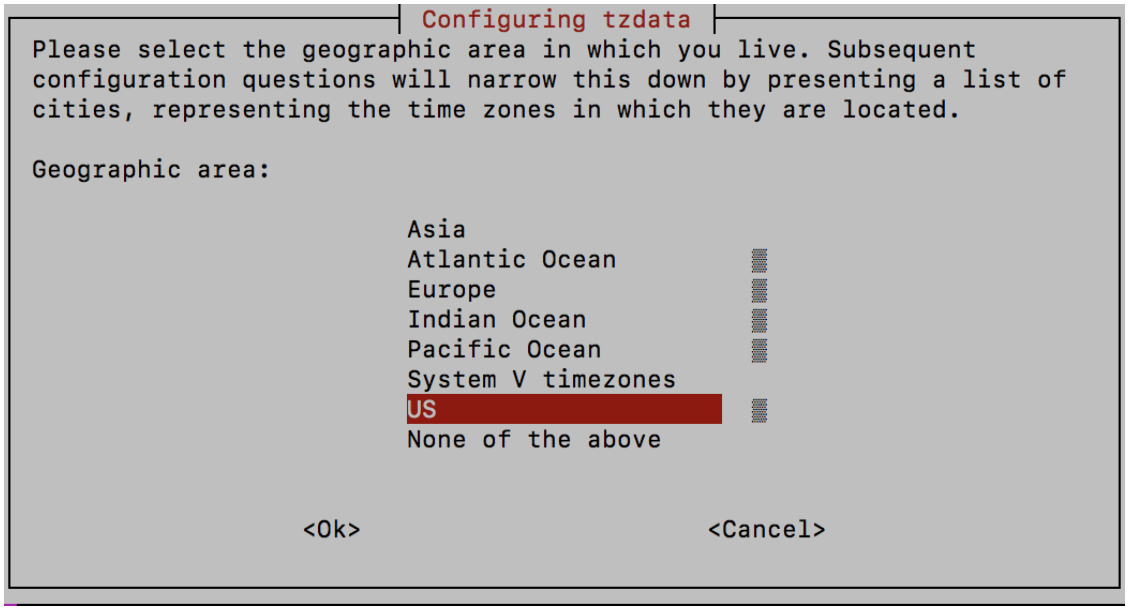
Configure the time zone of the Data Gateway VM.

In general, the data gateway VM launches with the default timezone as UTC. Cisco recommends that you configure the timezone to match your geographical area. With this configuration, all the Crosswork Data Gateway processes including the Showtech logs use the same configured timezone.

Procedure

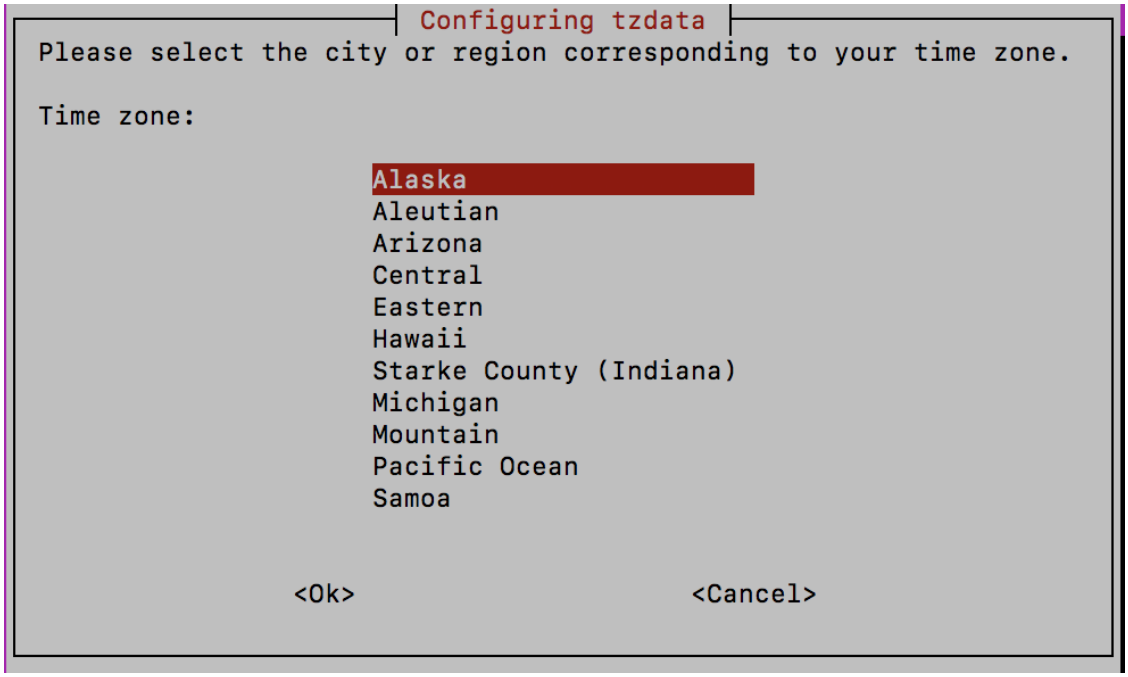
1. In the data gateway VM interactive menu, select **Change Current System Settings**.
2. Select **9 Timezone**.
3. Select the geographic area in which you live.

   Figure 34: Geographic area selection



4. Select the city or region corresponding to your timezone.

   Figure 35: Region selection



5. Select **OK** to save the settings.
6. Reboot the data gateway VM so that all processes pick up the new timezone.

7. Log out of the data gateway VM.

Data Gateway VM login and logout

Explains data Gateway VM login and logout, including the source-described role, behavior, and relationship to AWS EC2 installation and deployment.

This section describes how to log in and out to the data gateway VM.

Follow these steps to access and log out of the data gateway VM:

- Access Crosswork Data Gateway VM from SSH on page 156
- Log out of Crosswork Data Gateway VM on page 256

Access the Crosswork Data Gateway VM from SSH

Use this procedure to access the Crosswork Data Gateway VM from SSH, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Access the Crosswork Data Gateway VM from SSH.

Secure Shell (SSH) offers a protection from brute force attacks by blocking the client IP after several login failures. Failures such as incorrect username or password, connection disconnect, or algorithm mismatch are counted against the IP. Up to 4 failures within a 20 minute window causes the client IP to be blocked for at least 7 minutes. Continuing to accumulate failures cause the blocked time to be increased. Each client IP is tracked separately.

Follow these steps to log in to the data gateway VM from SSH.

Procedure

1. From your work station with network access to the data gateway management IP, run this command:

   ssh <username>@<ManagementNetworkIP>

   where **ManagementNetworkIP** is the management network IP address.

   For example,

   To login as administrator user: **ssh dg-admin@<ManagementNetworkIP>**

   To log in as operator user: **ssh dg-oper@<ManagementNetworkIP>**

   The Crosswork Data Gateway flash screen opens prompting for password.

2. Input the corresponding password (the one that you created during installation process) and press **Enter**.

If you are unable to access the data gateway VM, there is an issue with your network configuration settings. From the console, check the network settings. If they are incorrect, it is best to delete the data gateway VM and reinstall with the correct network settings.

Log out of Crosswork Data Gateway VM

Use this procedure to log out of Crosswork Data Gateway VM, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Log out of Crosswork Data Gateway VM.

Procedure

To log out of the data gateway VM, from the **Main Menu**, select **l Logout** and press **Enter** or click **OK**.

Installation and enrollment troubleshooting

Provides reference information about installation and enrollment troubleshooting, including the source-specified values, constraints, configuration parameters, and deployment relationships relevant to AWS EC2.

If Crosswork Data Gateway fails to autoenroll with Crosswork Network Controller, you can collect Crosswork Data Gateway show-tech (**Main menu** > **5 Troubleshooting** > **2 Run show-tech**) and check for the reason in `controller-gateway` logs. For more information on how to collect show-tech logs, refer to the *Collect show-tech logs from the Interactive Console* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. If there are session establishment or certificate-related issues, ensure that the `controller.pem` certificate is uploaded using the Interactive Console.

Note

When using an IPv6 address, it must be surrounded by square brackets ([1::1]).

This table lists common problems that might be experienced while installing or enrolling Crosswork Data Gateway, and provides approaches to identifying the source of the problem and solving it.

Table 70: Troubleshooting the Installation/Enrollment

Issue	Action
Crosswork Data Gateway cannot be enrolled with Crosswork Network Controller due to an NTP issue, that is, there is a clock-drift between the two. The clock-drift might be with either Crosswork Data Gateway or Crosswork Network Controller. Also, on the NTP servers for Crosswork Network Controller and Crosswork Data Gateway, the initial time is set to the ESXi server. For this reason, the ESXi server must also have NTP configured. **Sync the clock time on the host and retry.**	1. Log in to the data gateway VM. 2. From the main menu, select **5 Troubleshooting** > **2 Run show-tech**. Enter the destination to save the tarball containing logs and vitals and click **OK**. The show-tech is now encrypted with a file extension ending with .tar.xz. 3. Run this command to decrypt the show-tech file. `openssl enc -d -AES-256-CBC -pbkdf2 -md sha512 -iter 100000 -in <showtech file> -out <decrypted filename> -pass pass:<encrypt string>` In the show-tech logs (in file `session.log` at location `/opt/dg/log/controller-gateway/session.log`), if the logs contain the error `UNAUTHENTICATED:invalid certificate. reason: x509: certificate has expired or is not yet valid`, then there is a clock-drift between Crosswork Data Gateway and Crosswork Network Controller. 3. From the main menu, go to **3 Change Current System Settings** > **1 Configure NTP**. Configure NTP to sync with the clock time on the Crosswork Network Controller server and try reenrolling Crosswork Data Gateway.

Issue	Action
Crosswork Data Gateway remains in degraded state for more than 10 minutes with reason stated as "Could not collect vitals" due to certificate errors.	1. Log in to the data gateway VM. 2. From the main menu, select **5 Troubleshooting** > **2 Run show-tech**. Enter the destination to save the tarball containing logs and vitals and click **OK**. The show-tech is now encrypted with a file extension ending with .tar.xz. 3. Run this command to decrypt the show-tech file. `openssl enc -d -AES-256-CBC -pbkdf2 -md sha512 -iter 100000 -in <showtech file> -out <decrypted filename> -pass pass:<encrypt string>` In the show-tech logs (in file `gateway.log` at location `/opt/dg/log/controller-gateway/gateway.log`), if you see certificate errors, then reupload the Controller Signing Certificate, as explained in the steps below: 1. From the main menu, select **3 Change Current System Settings** > **7 Import Certificate**. 2. From the **Import Certificates** menu, select **1 Controller Signing Certificate File** and click **OK**. 3. Enter the SCP URI for the certificate file and click **OK**.
Crosswork Data Gateway remains in degraded state for more than 10 minutes with reason stated as "gRPC connection cannot be established" due to certificate errors.	1. Reupload the certificate file using these steps: a. From the main menu, select **3 Change Current System Settings** > **7 Import Certificate**. b. From the **Import Certificates** menu, select **1 Controller Signing Certificate File** and click **OK**. c. Enter the SCP URI for the certificate file and click **OK**. 2. Reboot the data gateway VM following the steps below: a. From the main menu, select **5 Troubleshooting** and click **OK**. b. From the Troubleshooting menu, select **4 Reboot VM** and click **OK**. c. Once the reboot is complete, check if the data gateway's operational status is **Up**.
During a Crosswork upgrade, some of the Crosswork Data Gateways may not get upgraded or reenrolled leading to logging multiple error messages in the dg-manager logs.	Reenroll or redeploy the data gateways. For more information, refer to the *Redeploy a Crosswork Data Gateway Instance* and *Reenroll Crosswork Data Gateway* sections in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Issue	Action

If a Crosswork Data Gateway instance that was previously attached to Crosswork is now reattached to a different Crosswork version 4.x or 5.0, the operational state of the instance may be Degraded with the robot-astack-influxdb error.	1. Log in to the Crosswork UI from the SSH.
	2. Run the Docker executive commands to access the **robot-astack-influxdb** pod.
	3. In the pod, navigate to the directory and delete it:
	`/mnt/datafs/influxdb`
	4. Restart the service using this command:
	`supervisorctl restart all`

If Data Gateway is redeployed without moving the gateway to the Maintenance mode, Crosswork enrollment will be unsuccessful and errors will be logged in the dg-manager and controller-gateway logs.	Move the Data Gateway to the **Maintenance** mode or manually reenroll the gateway. For more information, refer to the *Reenroll Crosswork Data Gateway* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Import the controller signing certificate file

Use this procedure to import the controller signing certificate file, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

Import the controller signing certificate file.

- You have not specified the **Controller Signing Certificate File URI** under the **Controller Settings** during installation.
- Cisco Crosswork was upgraded or reinstalled and you need to authenticate and enroll Crosswork Data Gateway with Cisco Crosswork.
- Cisco Crosswork configuration is in-progress when Crosswork Data Gateway tries to import the Controller Certificate file.
- The Cisco Crosswork Controller IP address in unreachable or incorrect.
- The Cisco Crosswork username or password is incorrect.

Follow these steps to import the controller signing certificate file:

> **Note**

Procedure

1. From the Data Gateway VM's Interactive Menu, select **3 Change Current System Settings**.

   The **Change System Settings** menu opens.

2. Select **7 Import Certificate**.
3. From the **Import Certificates** menu, select **1 Controller Signing Certificate File**.
4. Enter the SCP URI for the certificate file.

   An example URI is given below:

```
cw-admin@{server  ip }:/home/ cw -admin/ controller.pem
```

5. Enter the SCP passphrase (the SCP user password).

   The certificate file is imported.

6. Verify that the certificate was installed successfully. See View the controller signing certificate file on page 261.

View the controller signing certificate file

Use this procedure to view the controller signing certificate file, following the source steps and any stated prerequisites or verification actions for the AWS EC2 installation.

View the controller signing certificate file.

Follow these steps to view the signing certificate:

Procedure

---

1. From the Data Gateway VM's interactive menu, select **2 Show System Settings**.
2. From the **Show Current System Settings** menu, select **7 Certificates**.
3. Select **2 Controller Signing Certificate File**.

   Crosswork Data Gateway displays the default certificate if no new certificate has been imported. Otherwise, it displays the new certificate if it was successfully imported.

Part V

Install Crosswork Network Controller Applications

Topics:

- Install Crosswork Applications

10 Install Crosswork Applications

Topics:

- Install Crosswork Network Controller applications
- Add the auto-action file

Provides an overview of Crosswork Network Controller application installation, including CAPP validation, upload, installation, and activation, plus adding an auto-action YAML file.

This chapter explains how to install Crosswork Network Controller applications and add an auto-action YAML file. It covers CAPP validation, upload, installation, activation, and follow-up health checks.

Install Crosswork Network Controller applications

Install signed Crosswork Network Controller tier CAPP files through the UI, validate bundles when needed, and install and activate individual applications.

Install the Crosswork Network Controller applications that you require.

Use the Crosswork Network Controller UI to install applications.

Applications are offered in various tiers. Each tier contains Crosswork applications in a format unique to Crosswork called CAPP (Crosswork APPlication).

Before you begin

Ensure that all requirements of your application are met.

Procedure

1. **Download the CAPP files:** Navigate to cisco.com and download the signed Crosswork Network Controller tier CAPP that you require to a directory in your machine. For the purpose of these instructions, we will use the file name **cnc-premier-7.2.1.tar.gz**.
2. (Optional) CAPP files are automatically validated in the backend when you add them to the Crosswork Network Controller UI. If the upload fails, you can manually extract and validate the CAPP file using these steps:
   a) Extract the CAPP file.

   ```
   tar -xvf <signature file>
   ```

   Example:

   ```
   cd <folder where tar was download>
   tar -xvf cnc-premier-7.2.1.tar.gz
   README
   cnc-premier-7.2.1-88.tar.gz
   cnc-premier-7.2.1-88.tar.gz.tar.gz.signature
   CW-CCO_RELEASE.cer
   cisco_x509_verify_release.py3
   cisco_x509_verify_release.py
   ```

   b) Use the Python script to validate the signature of each file you plan to use.

   > **Note**
   >
   > Use `python --version` to find out the version of Python on your machine.

   To validate the file with Python 2.x, run:

   ```
   python cisco_x509_verify_release.py -e <.cer file> -i <.tar.gz file> -s
   <.tar.gz.signature file> -v dgst -sha512
   ```

   To validate the file with Python 3.x, run:

   ```
   python cisco_x509_verify_release.py3 -e <.cer file> -i <.tar.gz file> -s
   <.tar.gz.signature file> -v dgst -sha512
   ```

Example:

```
python cisco_x509_verify_release.py3 -s
cnc-premier-7.2.1-88.tar.gz.signature -i cnc-premier-7.2.1-88.tar.gz -e
CW-CCO_RELEASE.cer
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 CW-CCO_RELEASE.cer.
Successfully verified the signature of cnc-premier-7.2.1-88.tar.gz using
CW-CCO_RELEASE.cer
```

Note

If you do not have Python installed, go to python.org and download the version of python that is appropriate for your workstation.

3. **Add the CAPP file to Crosswork Network Controller UI:**
   a) Log in to Cisco Crosswork and on the home page, click **Administration** > **Crosswork Manager**. The **Crosswork Summary** page is displayed with Crosswork Cluster and Crosswork Platform Infrastructure tiles.

      You can click the tiles to get more information.

   b) Click **Application Management** and select the **Applications** tab.
   c) Click **Add new file** > **Upload application bundle (.tar.gz)**.

      The **Add Application Bundle (.tar.gz)** dialog box is displayed.

   d) In the dialog box, choose either **URL** or **SCP** as your preferred protocol. Based on your selection, fill in the additional fields with the required information.

      Click **Add** to proceed.

      Note

      When installing a Crosswork Network Controller package, there is no need to untar the package. You can add the package tarball as-is to the Crosswork UI and the applications within are automatically added. You can then install the individual applications as needed.

   e) The add operation progress is displayed on the **Applications** screen. You can also view the installation progress in the **Job History** tab.

> **Note**
>
> When loading a Crosswork Network Controller package, the loading process may stop at 50% for a while depending on the resources your host platform has available.

The newly added application files are displayed as tiles on the **Applications** screen.

4. **Install the Application CAPP file:**

   a) Click the **Install** prompt on the application tile. You can also click ⋮ on the tile, and select the **Install** option from the drop down list.

   > **Note**
   >
   > After you install Crosswork Cluster and Crosswork Data Gateway, install these Crosswork Network Controller applications in this sequence:
   >
   > 1. Element Management Functions (available in Essentials tier)
   > 2. Crosswork Optimization Engine (available in Advantage tier)
   > 3. Crosswork Active Topology (available in Advantage tier)
   > 4. Crosswork Service Health (available in Advantage tier)

   b) After an application is installed, it must be activated to become functional. The first-time installation also activates a CAPP file. However, if the activation fails after a successful installation, you can manually activate the application.

   To manually activate an application, click the ⋮ on the application tile, and select **Activate**.

   > **Note**
   >
   > In a single VM deployment, applications are installed and activated automatically, eliminating the need for separate activation.

5. Repeat step 3 for installing any remaining applications.
6. (Optional) Click ⋮ on the application tile, and select the **View details** option to view details of the installed application.
7. Once an application (or all applications) have been installed, check the health of the environment to make sure all the applications are healthy. It can take up to an hour for all the processes to launch and for the applications to report as healthy. If a newly installed application is not healthy after an hour, contact the Cisco Customer Experience team.

The application is now installed. Its tile icon changes. The Crosswork UI dynamically loads all related resources, UI screens, and menu options.

After installation succeeds, the 90-day evaluation period starts automatically. You can register the Crosswork Network Controller license with your Cisco Smart Account in the **Smart licenses** tab.

What to do next

Return to the installation workflow:

- VMware: Install Cisco Crosswork Network Controller on VMware vCenter on page 34
- AWS EC2: Install Cisco Crosswork Network Controller on AWS EC2 on page 38
- KVM: Install Cisco Crosswork Network Controller on KVM on page 36

Add the auto-action file

Add an auto-action YAML file through the Crosswork Network Controller UI to automate the sequential addition and installation of applications.

Add an auto-action file to automate the installation and activation of Crosswork Network Controller applications.

The auto-action functionality is optional and can automate the installation and activation of applications as needed.

Auto-action is supported for both cluster and single VM deployments.

Before you begin

Ensure you have defined the auto-action file (.yaml) with relevant details to automate the download and installation of Crosswork Network Controller applications. For more information on the sample file, refer to Sample auto action templates on page 555.

Procedure

1. Log in to Cisco Crosswork and on the home page, click **Administration** > **Crosswork Manager**.

   The **Crosswork Summary** page is displayed.
2. Click the **Application Management** tab.
3. Click **Add new file** > **Upload auto-install (.yaml)**.

   The **Upload YAML for Auto-Installation** dialog box is displayed.
4. Click **Browse** and select the auto-action file from your machine.
5. Click **Upload** to proceed.

Once the upload is successful, jobs are automatically triggered sequentially to add and install the applications. If a failure occurs, check the job details for more information. If a job is not scheduled, check the Alarm page for additional details.

Part VI

Integrate Cisco NSO and SR-PCE with Cisco Crosswork Network Controller

Topics:

- Integrate Cisco NSO
- Integrate SR-PCE

11 Integrate Cisco NSO

Topics:

- NSO integration workflow
- Configure standalone NSO
- Configure Cisco NSO for layered service architecture
- Install Cisco NSO function pack bundles from Crosswork UI
- Install NSO function packs manually
- Configure JSON web token authentication for NSO

Introduces Cisco NSO integration options and summarizes provider setup, function-pack management, troubleshooting, and authentication.

This chapter covers integrating standalone and Layered Service Architecture Cisco NSO deployments with Crosswork.

Chapter contents

It includes provider configuration, function-pack management, troubleshooting, manual installation resources, and JSON Web Token authentication.

NSO integration workflow

Summarizes the supported standalone and Layered Service Architecture deployments and lists the integration stages for authentication, provider access, and function packs.

Cisco Crosswork Network Controller supports Cisco Network Services Orchestrator (Cisco NSO) as a standalone deployment or as a Layered Service Architecture (LSA) deployment.

You select the deployment type when you install NSO. Use the standalone NSO workflow for a standalone NSO deployment. Use the LSA NSO workflow for an NSO deployment that uses customer-facing service (CFS) and resource-facing service (RFS) nodes.

NSO integration stages

These stages describe the NSO integration workflow:

1. Install a compatible Cisco NSO version. For compatibility details, refer to the *Release Notes for Crosswork Network Controller, Release 7.2.1 SP* for your Crosswork Network Controller release. For installation instructions, refer to Cisco NSO documentation.
2. Configure the NSO access-control settings that Crosswork uses for role-based access control and Device Access Groups.
3. Add the NSO provider in Crosswork. Use the standalone provider properties for standalone NSO, and use the LSA provider properties for LSA NSO.
   - For standalone NSO, refer to Configure standalone NSO.
   - For LSA NSO, refer to Configure Cisco NSO for Layered Service Architecture.
4. Edit the provider policy details if the default device-addition policy does not match your environment.
5. Install the required Cisco NSO function packs. For more information, refer to Install Cisco NSO function pack bundles from Crosswork UI on page 286 (recommended) or Install NSO function packs manually on page 294.

   Note

   - The NSO core function packs are bundled in the *cnc-function-packs-7.2.1.signed.bin* file.
   - The Cisco Crosswork Network Controller Function Pack SDK Application (*cnc-function-pack-sdk-7.2.1-bin.zip.tar.gz*) is also available for download on cisco.com. The SDK provides tools and source-code examples you can use to develop, build, package and deploy the TSDN function pack on Crosswork Network Controller.

6. Configure JSON Web Token (JWT) authentication between Crosswork and NSO.

Configure standalone NSO

Configure standalone NSO for Crosswork JSON Web Token (JWT) authentication, role-based access control, provider access, and function pack installation.

Use this workflow when your Crosswork deployment uses standalone Cisco NSO.

Before you begin

- Install a compatible Cisco NSO version.
- Create the Crosswork credential profile that you will use for the NSO provider.
- Know the NSO server IP address or fully qualified domain name (FQDN) and the Network Element Driver (NED) device models and driver versions used in your topology.

Procedure

1. Enable service-level NETCONF Access Control Model (NACM) on the standalone NSO server.

   Example

   ```
   ncs_cli -u admin
   configure
   set nacm enforce-nacm-on-service true
   commit dry-run
   commit
   ```

2. Create the NACM group before you create the rule list.

   For an admin user with a name other than `admin`, add the user to the `CNC#ALL-ACCESS` group.

   ```
   ncs_cli -u admin
   configure
   set nacm groups group "CNC#ALL-ACCESS" user-name admin
   commit dry-run
   commit
   ```

   For a non-admin user, create a NACM group that matches the Device Access Group. This example uses `RW-CW` as the user and `DAG-2` as the Device Access Group.

   ```
   ncs_cli -u admin
   configure
   set nacm groups group "CNC#DAG-2" user-name RW-CW
   commit dry-run
   commit
   ```

3. Generate the NACM rules in Crosswork and apply the rules to NSO.

   In Crosswork, choose **Administration** > **Users and Roles** > **Users** > **Generate NACM Rules**. Apply the generated rules to NSO by using Commit Manager or by copying the XML to a file such as `sample-nacm.xml` and loading it with the merge operation.

   ```
   ncs_cli -u admin
   configure
   load merge /home/nso/sample-nacm.xml
   commit
   ```

4. Add the standalone NSO provider in Crosswork.

   Refer to Add a standalone Cisco NSO provider.

5. Install the required Cisco NSO function packs.

   Refer to the NSO function pack installation topics in this chapter.

Restrict sync-to access for the Crosswork NSO user

Restrict sync-to access by creating NSO access-control rules and a Linux user account for Crosswork authentication.

Use this task to create a restricted NSO user for Crosswork authentication when you do not want the Crosswork-authenticated user to execute `sync-to`.

Before you begin

- Log in to the NSO server Linux shell.
- Know the user name and password that you want Crosswork to use when it authenticates to NSO.

Procedure

1. Open the NSO CLI without authentication.

   Example

   ```
   ncs_cli -C -noaaa
   ```

2. Create an NACM rule list that denies `sync-to` access for the restricted Crosswork user group.

   This example uses `restrictedCNC` as the NACM group name.

   Example

   ```
   config
   nacm rule-list CNC
   group restrictedCNC
   rule restrict-syncto action deny
   module-name tailf-ncs
   path /ncs:devices/device/sync-to
   exit
   rule restrict-multi-syncto action deny
   module-name tailf-ncs
   path /ncs:devices/sync-to
   exit
   rule restrict-group-syncto action deny
   module-name tailf-ncs
   path /ncs:devices/device-group/sync-to
   exit
   rule any-access action permit
   exit
   cmdrule any-command action permit
   commit dry-run
   commit
   ```

3. Create the Linux group and user account for Crosswork authentication.

   This example uses `restrictedCNC` as the Linux group name and `CNCuser` as the user name.

   Example

   ```
   sudo groupadd restrictedCNC
   sudo useradd CNCuser
   sudo usermod -a -G restrictedCNC CNCuser
   sudo passwd CNCuser
   ```

4. When prompted, enter the password for the Crosswork-authenticated NSO user.
5. Use the restricted user account when you configure Crosswork to authenticate to NSO.

Crosswork authenticates to NSO with a restricted user account that cannot execute `sync-to`.

Add a standalone Cisco NSO provider

Add a standalone Cisco NSO provider before you configure JSON Web Token (JWT) authentication and NSO access-control settings.

The Cisco NSO provider supplies network services, device configuration services, device management, and configuration maintenance services to Crosswork applications.

> **Note**
>
> Crosswork Network Controller scans NSO for device status changes only after an explicit NSO action. To add devices from NSO to Crosswork Network Controller, perform an NSO action for a device, or update and save the NSO provider policy details.

Before you begin

- Create a credential profile for the Cisco NSO provider.
- Confirm the Cisco NSO device configurations.
- Know the provider name, server IP address or fully qualified domain name (FQDN), protocol, port, Network Element Driver (NED) model, and NED version for the provider. You can find the Cisco NSO version by using the `version` command.
- Make sure that the users who need the NSO cross-launch feature do not have read-only permissions.

Procedure

1. Choose **Administration** > **Manage Provider Access**.
2. Click **Add**.
3. Enter the required Cisco NSO provider field values.

Table 71: Standalone NSO provider fields

Field	Value
Provider Name	Enter the provider name.
Credential Profile	Select the Cisco NSO credential profile.
Family	Select **NSO**.
Protocol	Select **HTTPS**, **SSH**, or both. To use **Backup NSO**, configure SSH connectivity.
Server Details	Enter the IP address or FQDN of the NSO server. For NSO high availability (HA), use the management virtual IP (VIP) address.

Field	Value
Port	For HTTPS, enter the port configured in `/etc/ncs/ncs.conf`. The default NSO HTTPS port is **`8888`**.
Model	Select the model prefix that matches the NED CLI used by Cisco NSO. • Valid values are **`Cisco-IOS-XR`**, **`Cisco-NX-OS`**, and **`Cisco-IOS-XE`**. • Add a model for each device type in the topology. • For telemetry, only **`Cisco-IOS-XR`** is supported.
Version	Enter the NED software version installed for the device model in NSO.
Timeout	Enter the connection timeout in seconds. The default is 30 seconds.

4. Enter the standalone NSO provider properties.

Table 72: Standalone NSO provider properties

Property key	Value
`forward`	Set to **`true`**. This property is required for Crosswork Network Controller provisioning operations in the UI and for the northbound interface to NSO through the Crosswork application programming interface (API) gateway. **Note** The default value is **`false`**. For Crosswork-to-NSO device-addition workflows, set this property to **`true`** and configure the relevant provider policy attributes in **Edit policy details**.
`nso_crosslaunch_url`	Enter the URL for cross-launching standalone NSO in this format: `https://<NSO IP address/FQDN>:<port-number>`

5. Click **Save**.
6. If the default policy values do not match your environment, edit the NSO provider policy details.

   Use policy details to control how Crosswork matches, adds, synchronizes, and manages devices between Crosswork Network Controller and NSO.

a. On the NSO provider, choose **Actions** > **Edit policy details**.

b. Edit the configuration fields to match your environment.

c. Click **Save**.

> **Note**
>
> Set **Onboard from** to **`TRUE`** when Crosswork Network Controller must add devices from NSO and the devices do not have an IP address match. The default value is **`FALSE`**.

Table 73: Editable NSO provider policy attributes

Attribute	Description
Match	Set to **`True`** to match Crosswork Network Controller devices with NSO devices based on IP address.
MatchRule	Enter an expression that defines the subset of devices to match.
Onboard To NSO	Set to **`True`** to add missing devices to NSO.
Onboard To Rule	Enter an expression that defines the subset of devices to add to NSO.
Onboard From	Set to **`True`** to add devices from NSO to Crosswork Network Controller when they are missing.
Onboard From Rule	Enter an expression that defines the subset of devices to add from NSO.
Sync From	Set to **`True`** to run sync-from on the NSO device after Crosswork adds the device.
Sync From Rule	Enter an expression that defines the subset of devices for sync-from.
Check Sync	Set to **`True`** to check the synchronization status of NSO devices.
Check Sync Rule	Enter an expression that defines the subset of devices for check-sync.
NED	Specify the Network Element Driver (NED) to use. By default, NSO uses the latest CLI NED.
Rule	Enter an expression that defines which devices use a specific NED.

7. Confirm with your network administrator whether proxy forwarding must be configured in your network. If proxy forwarding is required for the Crosswork UI, configure proxy forwarding on the NSO server.

Add the proxy-forwarding settings in the `ncs.conf` file on the NSO server.

Replace the example subnet and mask values with the values for your Crosswork Network Controller environment. If you have multiple subnets, repeat the `<allowed-proxy-ip-prefix>` blocks for each subnet.

Use this example to configure proxy-forwarding settings in `ncs.conf`.

```
<webui>
...
  <use-forwarded-client-ip>
    <proxy-headers>X-Forwarded-For</proxy-headers>
    <proxy-headers>X-REAl-IP</proxy-headers>
    <allowed-proxy-ip-prefix>10.195.72.0/24</allowed-proxy-ip-prefix>

<allowed-proxy-ip-prefix>2001:0db8:85a3:0000:0000:8a2e:0370::/112</allowed-proxy-ip-prefix>

  </use-forwarded-client-ip>
...
</webui>
```

Configure Cisco NSO for layered service architecture

Configure Cisco NSO for Layered Service Architecture (LSA), including Crosswork JSON Web Token (JWT) authentication, customer-facing service (CFS) and resource-facing service (RFS) access control, provider access, and function pack installation.

Use this workflow when your Crosswork deployment uses Cisco NSO Layered Service Architecture (LSA).

An LSA deployment uses customer-facing service (CFS) and resource-facing service (RFS) NSO nodes. Crosswork identifies the NSO provider as CFS or RFS, and only one CFS is allowed.

Before you begin

- Install a compatible Cisco NSO version and set up NSO LSA.
- Create the Crosswork credential profiles that you will use for the CFS and RFS providers.
- Know the CFS and RFS provider names, IP addresses or fully qualified domain names (FQDNs), protocols, ports, Network Element Driver (NED) models, and NED versions.

Procedure

1. On each NSO RFS node, enable local authentication in the `ncs.conf` file under the `<aaa>` section.

   Example

   ```
   <local-authentication>
     <enabled>true</enabled>
   </local-authentication>
   ```

   If you are using the CFS node, skip this step.

2. Restart NSO on each RFS node.

   Example

   ```
   sudo /etc/init.d/ncs restart
   ```

3. Enable service-level NETCONF Access Control Model (NACM) on both the CFS and RFS nodes.

   Example

   ```
   ncs_cli -u admin
   configure
   set nacm enforce-nacm-on-service true
   commit dry-run
   commit
   ```

4. Configure user association for LSA role-based authentication.
   - If NSO is used exclusively with Crosswork, remove the `auth-group default` map. If there are non-Crosswork NSO users, keep the default map and create an `auth-group umap` entry for every Crosswork user.
   - Define each Crosswork user under `aaa:aaa` as an authentication user on every RFS node. The username must match the username in Crosswork, but the password can differ.
   - Add every Crosswork user as a `umap` entry under the device authentication group in the CFS so that the CFS can pass user requests to the RFS node as the corresponding user.

5. Define a generic NACM group and NACM rule with all permissions on the CFS.

   When you create a user in Crosswork, add that user to the `CNC#ALL-ACCESS` NACM group in CFS.

   ```
   group "CNC#ALL-ACCESS" {
     user-name [ RW-CW admin rw-user ];
   }
   ```

   Copy the NACM rules from Crosswork and apply them to the CFS and RFS nodes.

   ```
   <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
     <!-- NACM rules for NSO - CFS -->
     <rule-list>
       <name>CNC#ALL-ACCESS#rule-list</name>
       <group>CNC#ALL-ACCESS</group>
       <rule>
         <name>CNC#ALL-ACCESS#rule-list#task-level-access</name>
         <action>permit</action>
         <log-if-permit xmlns="http://tail-f.com/yang/acm"/>
       </rule>
       <cmdrule xmlns="http://tail-f.com/yang/acm">
         <name>any-access</name>
         <action>permit</action>
       </cmdrule>
     </rule-list>
   </nacm>
   <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
     <!-- NACM rules for NSO - RFS -->
     <rule-list>
       <name>CNC#ALL-ACCESS#rule-list</name>
       <group>CNC#ALL-ACCESS</group>
       <rule>
         <name>CNC#ALL-ACCESS#rule-list#task-level-access</name>
         <action>permit</action>
         <log-if-permit xmlns="http://tail-f.com/yang/acm"/>
       </rule>
       <cmdrule xmlns="http://tail-f.com/yang/acm">
         <name>any-access</name>
         <action>permit</action>
       </cmdrule>
   ```

```
  </rule-list>
</nacm>
```

6. Add the Device Access Groups and NACM rules on the RFS node.

   Crosswork automatically creates Device Access Groups, so you do not need additional NSO steps for this task. Apply the generated NACM rules to NSO by using Commit Manager or by loading the XML file with the merge operation.

7. Add the LSA NSO providers in Crosswork.

   Refer to Add Cisco NSO providers for Layered Service Architecture.

8. If Geo-HA is set up and you encounter a 503 error, add the proxy settings to the `/etc/environment` file on the CFS node.

```
https_proxy="http://proxy.esl.cisco.com:80"
http_proxy="http://proxy.esl.cisco.com:80"
no_proxy="localhost,127.0.0.1,10.0.0.0/8,192.168.0.0/16,172.16.0.0/12,cisco.com,<az1
 mgmt vip>,<az2 mgmt vip>,<fqdn of CW geo-mgmt VIP>"
```

   Source the file, and then reboot the CFS nodes for the proxy settings to take effect.

9. Install the required Cisco NSO function packs.

   Refer to the NSO function pack installation topics in this chapter.

Add Cisco NSO providers for layered service architecture

Add Cisco NSO providers for a Layered Service Architecture (LSA) deployment after you configure access-control settings for customer-facing service (CFS) and resource-facing service (RFS) nodes.

Use this task to add the CFS and RFS NSO providers that Crosswork uses in an LSA deployment.

Crosswork supports Cisco NSO LSA deployments that use multiple NSO providers. The CFS NSO contains the services, and the RFS NSO contains the devices. Crosswork identifies the NSO provider as CFS or RFS, and only one CFS is allowed.

> **Note**
>
> Crosswork Network Controller scans NSO for device status changes only after an explicit NSO action. To add devices from NSO to Crosswork Network Controller, perform an NSO action for a device, or update and save the NSO provider policy details.

Before you begin

- Configure the CFS and RFS NETCONF Access Control Model (NACM) settings.
- Enable LSA settings in Crosswork before you add an NSO LSA provider.
- Make sure that RFS node IP addresses configured on the CFS match the RFS node IP addresses shown in the Crosswork UI. If the IP addresses do not match, Crosswork reports that the LSA cluster is missing RFS providers.
- Create the credential profiles for the LSA NSO providers.
- Confirm the Cisco NSO device configurations.

Procedure

1. If LSA is not enabled, enable LSA settings in Crosswork.
   a. Choose **Administration** > **Settings** > **System settings** > **Layered service architecture**.
   b. Select **Enable**.
   c. Select the device distribution method: **Round Robin**, **Capacity**, or **User Defined**.
   d. Click **Save**.

   > **Note**
   >
   > After you save the LSA setting, you cannot disable LSA without first removing all NSO providers.

2. Choose **Administration** > **Manage Provider Access**.
3. Click **Add**.
4. Enter the required Cisco NSO provider field values for the CFS or RFS provider.

Table 74: LSA NSO provider fields

Field	Value
Provider Name	Enter the provider name.
Credential Profile	Select the Cisco NSO credential profile for this provider.
Family	Select **NSO**.
Protocol	Select **HTTPS**, **SSH**, or both. To use **Backup NSO**, configure SSH connectivity.
Server Details	Enter the IP address or fully qualified domain name (FQDN) of the CFS or RFS server. For NSO high availability (HA), use the management virtual IP (VIP) address.
Port	For HTTPS, enter the port configured in `/etc/ncs/ncs.conf`. The default NSO HTTPS port is **8888**.
Model	Select the model prefix that matches the Network Element Driver (NED) CLI used by Cisco NSO. • Valid values are **Cisco-IOS-XR**, **Cisco-NX-OS**, and **Cisco-IOS-XE**. • Add a model for each device type in the topology. • For telemetry, only **Cisco-IOS-XR** is supported.
Version	Enter the NED software version installed for the device model in NSO.

Field	Value
Timeout	Enter the connection timeout in seconds. The default is 30 seconds.

5. Enter the LSA NSO provider properties.

Table 75: LSA NSO provider properties

Property key	Value
`forward`	Set to `true`. This property is required for Crosswork Network Controller provisioning operations in the UI and for the northbound interface to NSO through the Crosswork application programming interface (API) gateway. **Note** For a CFS node, only the `forward` property key is used. The default value is `false`. For Crosswork-to-NSO device-addition workflows, set this property to `true` and configure the relevant provider policy attributes in **Edit policy details**.

`input_url_prefix`	For an LSA provider, enter the RFS ID in this format, where `x` is the RFS node number:
	
	`/rfc-x`
	
	Example for RFS node 1:
	
	`input_url_prefix: /rfc-1`

6. Click **Save**.
7. If the default policy values do not match your environment, edit the NSO provider policy details.

   Use policy details to control how Crosswork matches, adds, synchronizes, and manages devices between Crosswork Network Controller and NSO.

   a. On the NSO provider, choose **Actions** > **Edit policy details**.

   b. Edit the configuration fields to match your environment.

   c. Click **Save**.

Note

Set **Onboard from** to `TRUE` when Crosswork Network Controller must add devices from NSO and the devices do not have an IP address match. The default value is `FALSE`.

Table 76: Editable NSO provider policy attributes

Attribute	Description
Match	Set to `True` to match Crosswork Network Controller devices with NSO devices based on IP address.
MatchRule	Enter an expression that defines the subset of devices to match.
Onboard To NSO	Set to `True` to add missing devices to NSO.
Onboard To Rule	Enter an expression that defines the subset of devices to add to NSO.
Onboard From	Set to `True` to add devices from NSO to Crosswork Network Controller when they are missing.
Onboard From Rule	Enter an expression that defines the subset of devices to add from NSO.
Sync From	Set to `True` to run sync-from on the NSO device after Crosswork adds the device.
Sync From Rule	Enter an expression that defines the subset of devices for sync-from.
Check Sync	Set to `True` to check the synchronization status of NSO devices.
Check Sync Rule	Enter an expression that defines the subset of devices for check-sync.
NED	Specify the Network Element Driver (NED) to use. By default, NSO uses the latest CLI NED.
Rule	Enter an expression that defines which devices use a specific NED.

8. Repeat this task for each required CFS and RFS provider.

What to do next

Install the required NSO function packs.

Install Cisco NSO function pack bundles from Crosswork UI

Describes the Cisco NSO function pack bundle format and identifies the Crosswork UI tabs used to view, add, and deploy bundles.

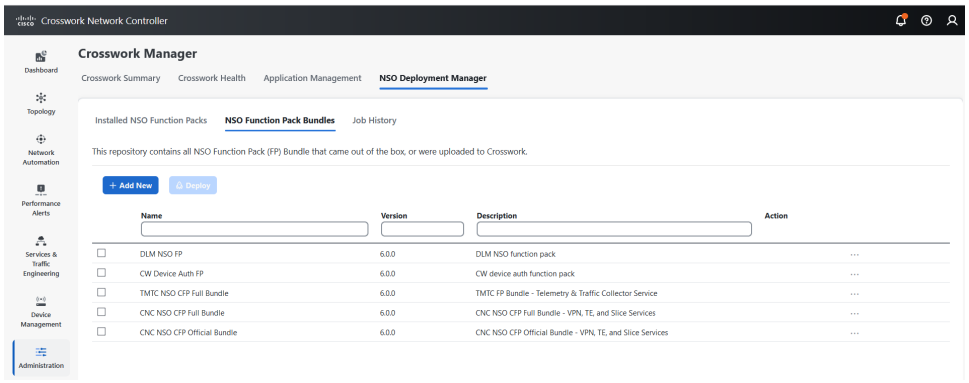
In the Cisco NSO function pack bundles, the NSO function pack files are bundled as tar.gz files. To ensure interoperability with Crosswork, Cisco NSO requires the installation of the essential function packs.

Crosswork UI function pack tabs

In the Crosswork UI, the **NSO Deployment Manager** tab lets you manage the function pack bundles using this tabs:

- **Installed NSO Function Packs**: Provides the list of NSO function packs deployed on the configured NSO server. Refer to View NSO function pack bundles on page 286 for more information.
- **NSO Function Pack Bundles**: Allows you to add and deploy the function pack bundles. Use this tab, to view the artifacts in the function pack bundle, download, and delete the function pack bundles. Refer to Manage NSO function pack bundles on page 287 for more information.
- **Job History**: The **Job History** tab displays a summary of the jobs, job ID, time when the job is started and completed, job description, and target. Refer to View NSO function pack job history on page 293 for more information.

Figure 36: NSO deployment manager window



View NSO function pack bundles

View NSO function pack bundles installed on the available NSO servers in the Crosswork UI.

View installed Cisco NSO function pack bundles in the Crosswork UI, including their operational state, description, and version. If an NSO service provider is unreachable, the UI cannot display its installed function packs.

Operators can retrieve a list of all of the function pack bundles installed on each available NSO server via the Crosswork UI or using the API.

Note

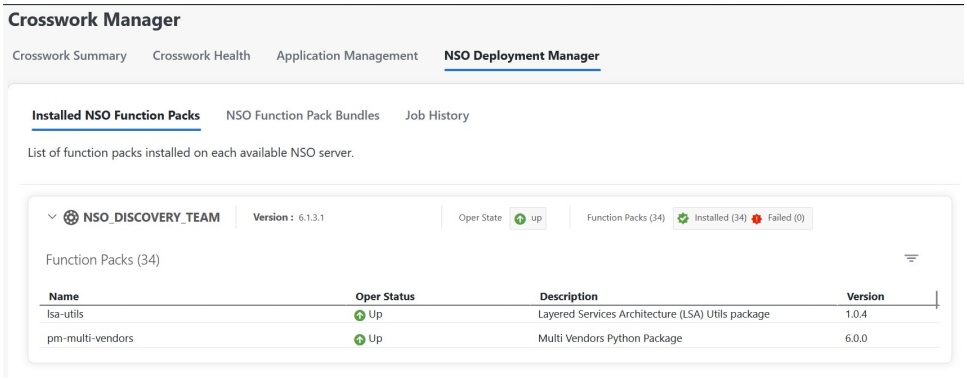
If any of the NSO service providers is unreachable, you cannot view the installed NSO function packs. An error "*Server is temporarily unavailable, try to relogin*" is displayed.

Follow the steps below to view the installed NSO function pack bundles through the UI.

Procedure

1. From the main menu, choose **Administration** > **Crosswork Manager**, click the **NSO Deployment Manager** tab.
2. Click the **Installed NSO Function Packs** tab.
3. Select the NSO provider and expand the NSO server entry to view the list of function packs and their details (function pack name, operational state as **Up** or **Down**, description, and version).

   Figure 37: Installed NSO function packs window



Manage NSO function pack bundles

Add a custom NSO function pack bundle to Crosswork after verifying the stated NSO server prerequisites.

Add custom NSO function pack bundles to Crosswork after verifying the required Java, Python, and open-file settings on the NSO server. Monitor the upload status in Job History.

Before you begin

Ensure that these prerequisites are met:

- Java version 17.0.0 or higher is installed on the NSO server.
- Python version 3.11.0 or higher is installed on the NSO server.
- NSO is configured to allow 64,000 openFileDescriptors.

You can add and deploy custom NSO function packs in addition to the function packs that are added by default to the Crosswork UI. The preinstalled bundles include this packs:

Table 77: Default NSO core function pack bundles

Package Name	Contents
DLM NSO FP	Cisco NSO DLM Service Pack
Device Auth NSO FP	Cisco Crosswork Change Automation NSO Function Pack
TMTC NSO FP	Cisco NSO Telemetry Traffic Collector Function Pack
CNC NSO FPs Plus Sample FPs	Crosswork Network Controller NSO function packs for VPN, TE, and Slice services, with sample FPs included
CNC NSO FPs	Crosswork Network Controller NSO function packs for VPN and TE services

Package Name	Contents
Cisco NSO Traffic Engineering manager	Crosswork Network Controller NSO TE Manager for implicit service and transport provisioning.
CWM SOLUTION FPS	NSO Function Pack for Crosswork Workflow Manager Solutions

> **Note**
>
> Only available after deployment of Premier tier.

Procedure

1. Ensure that your NSO setup meets all of the prerequisites.

   Check the python and java versions using the `--version` command.

   ```
   python --version

   Python 3.11.0

   java --version

   openjdk 17.0.9 2023-10-17
   OpenJDK Runtime Environment (build 17.0.9+9-Ubuntu-120.04)
   OpenJDK 64-Bit Server VM (build 17.0.9+9-Ubuntu-120.04, mixed mode, sharing)
   ```

   To verify or configure the required file descriptor limits on the NSO server, check the `/etc/security/limits.conf` file. If the required values are not already configured, add these lines before the commented line #End of file:

   ```
   $ grep hard /etc/security/limits.conf
   $ grep soft /etc/security/limits.conf
   * soft nproc 65535
   * hard nproc 65535
   * soft nofile 65535
   * hard nofile 65535
   * soft memlock 65535
   * hard memlock 65535
   ```

2. From the main menu, choose **Administration > Crosswork Manager**, click the **NSO Deployment Manager** tab.
3. Click the **NSO Function Pack Bundles** tab.

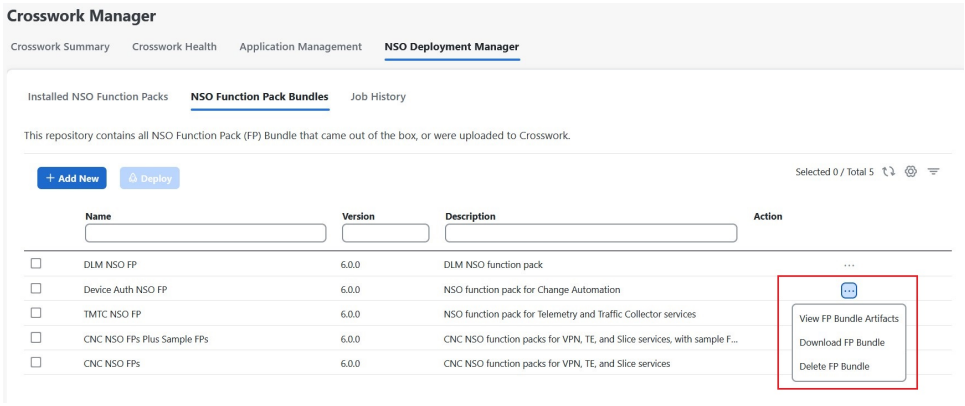
   All the installed NSO function pack bundles get displayed with the bundle name, version, and description information.

   To manage the bundles, select one or more bundles and click the **Action** menu to perform this:

   - **View FP Bundle Artifacts**: View the hierarchy of the artifacts that are bundled in the selected package.
   - **Download FP Bundle**: Download the function pack bundle.

- **Delete FP Bundle**: Delete the function pack bundle.

Figure 38: Action Menu



4. Click **Add New** to install the new function pack bundle. The **Add New NSO Function Pack Bundle** page is displayed.

> **Note**
>
> The **Add New** button is disabled when Crosswork is in maintenance mode.

5. In the **Add New NSO Function Pack Bundle** page, enter this:
   - **Host Name/IP address**: Enter the IP address and subnet mask of the Cisco NSO server.
   - **Port**: For HTTPS, enter the port that corresponds with what is configured on the NSO VM in etc/ncs/ncs.conf to access NSO using HTTPS. NSO uses 8888 as the default port.
   - **User Name**: The username used to log in to the NSO server.
   - **Password**: The password credentials to authenticate into the NSO server.
   - **Server Path/Location**: The server path of the NSO server.
6. Click **Test SSH Connectivity** again to validate SSH-based connectivity. If the connection is successful, a confirmation message indicating that the NSO bundle upload is in-progress appears. Click **View Progress in Job History** to view the upload status.
7. Click **Add**.
8. Check the **Job History** to monitor the addition of the package to Crosswork. Once the package has been added, the next step is to deploy the function pack.

What to do next

After the function pack is added, deploy the function pack on NSO. Refer to Deploy NSO function pack bundles on page 289.

Deploy NSO function pack bundles

Deploy an NSO function pack bundle that is already uploaded to the Crosswork UI.

Deploy an uploaded NSO function pack bundle from the Crosswork UI. Review SSH connection and deployment target details before deploying, then monitor installation status in Job History.

Before you begin

- Ensure that the NSO function pack bundle is uploaded to the Crosswork UI. Refer to Manage NSO function pack bundles on page 287 for more information.
- If you plan to deploy the function pack bundle in an HA environment, you must have the primary and secondary server details readily available.
- If your primary and secondary NSO servers and Crosswork servers are in different subnets, you must configure either an IP static route (**Administration** > **Settings** > **Static Routes**) or an IP rule policy (run `ip rule add from all to 10.19.0.4 lookup cw_data` on the Crosswork server) to enable connectivity between the servers.

This topic explains the process to deploy the NSO function pack bundles.

The Cisco NSO sample function packs are provided as a starting point for VPN service provisioning functionality in Cisco Crosswork Network Controller. While the samples can be used “as is” in some limited network configurations, they are intended to demonstrate the extensible design of Cisco Crosswork Network Controller. Answers to common questions can be found on Cisco Devnet and Cisco Customer Experience representatives can provide answers to general questions about the samples. Support for customization of the samples for your specific use cases can be arranged through your Cisco account team.

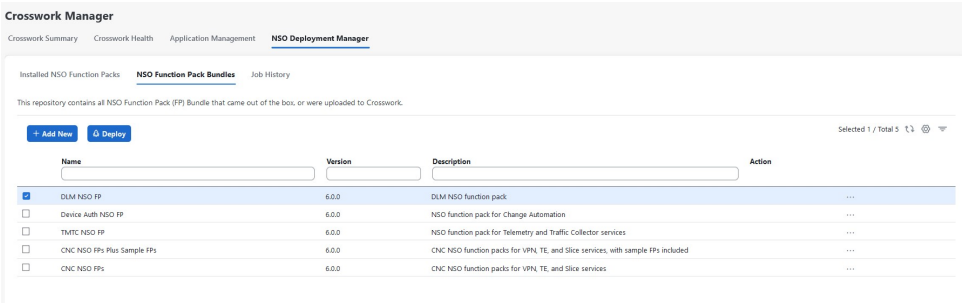
> **Note**
>
> NSO function packs cannot be deployed when Crosswork is in maintenance mode.

Procedure

1. From the main menu, choose **Administration > Crosswork Manager**, click the **NSO Deployment Manager** tab.
2. Click the **NSO Function Pack Bundles** tab.
3. Select the NSO function pack bundle and click **Deploy**.

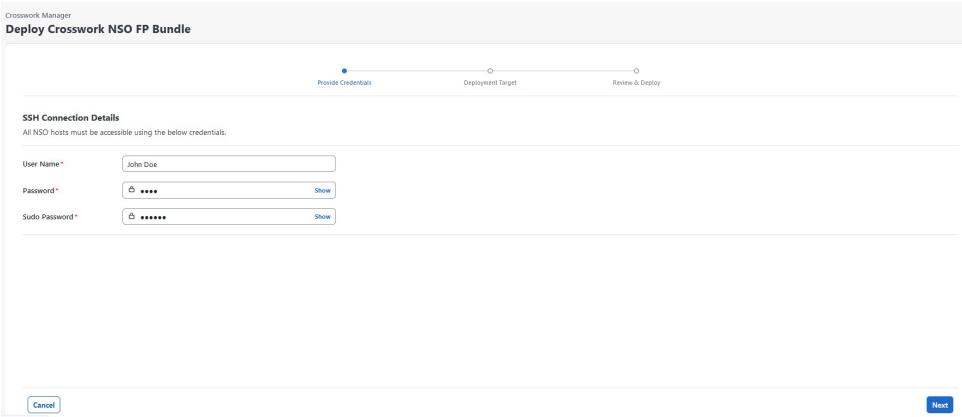
   > **Note**
   >
   > You can only select up to 3 Function Packs to be installed at a time. To install more, install the 3 function packs first and then repeat this process until you have installed all the Function Packs you will use.

Figure 39: NSO function pack bundles window



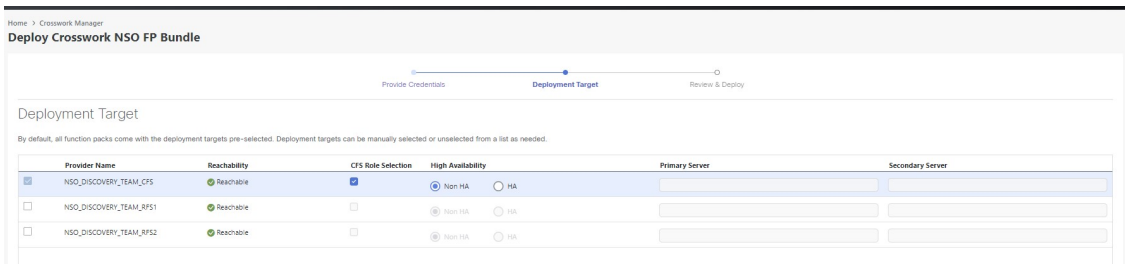
4. In the **Deploy Crosswork NSO FP Bundle** page, enter this SSH connection details:
   - **User Name**: The SSH username for server access.
   - **Password**: The SSH password for server access.
   - **Sudo Password**: The SSH sudo password.

Figure 40: SSH connection details page



5. Click **Next**.
6. In the **Deployment Target** section, review the target details:
   - **Provider Name**: Displays the name of the provider.
   - **Reachability**: Displays the reachability status of the provider.
   - **CFS Role Selection**: This column appears when a role is not assigned to a provider. Select the check box that corresponds to the provider row to assign the customer-facing service (CFS) role. The resource-facing service (RFS) role is automatically assigned to the other providers. For more information about CFS, RFS, and Cisco NSO Layered Service Architecture (LSA) deployment concepts, refer to the *Prepare Infrastructure for Device Management* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
   - **High Availability**: Depending on your deployment preferences for the function packs bundle on an NSO node, select either non-HA or HA. If you have selected HA, enter the server details in the **Primary Server** and **Secondary Server** fields.

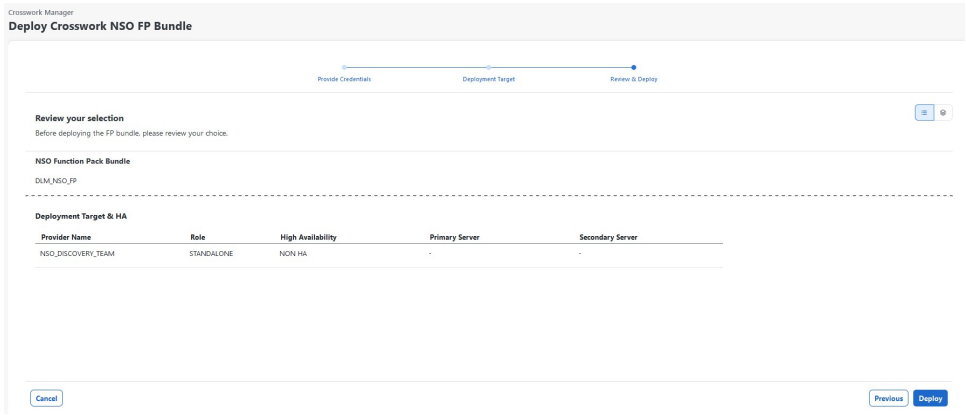
Figure 41: Deployment target page



7. Click **Next**.
8. In the **Review & Deploy** page, review the NSO bundle and deployment target details that you have configured. If you want to modify your selection, click **Previous** to view the earlier pages and modify it as required.

> **Note**
>
> If the provider is deployed on a standalone NSO node, the role is displayed as STANDALONE.

Figure 42: Review selection page



9. Deploy the bundle and monitor the installation.
   a) Click **Deploy**.
   b) Check the **Job History** to monitor the installation.

What to do next

Repeat the process for any additional Function Packs that you need to install.

NSO function pack installation troubleshooting

Lists common problems that can occur when you install or deploy Cisco NSO function packs and provides the actions specified for those problems.

This table lists common problems that can occur when you install or deploy a Cisco NSO function pack.

Function pack installation issues

Table 78: Troubleshooting the function pack installation issues

Issue	Action
The function pack deployment failed with this error: `Failed to open SSH connection to host coffee-nso1.cisco.com`	In an HA configuration, the NSO engine assumes that the NSO primary and secondary servers, and the Crosswork server reside in the same subnet. If the servers have different subnets, you must configure an IP route or an IP rule policy to ensure connectivity between the servers. When the routes are not configured, the engine cannot locate the subnet, and the function pack deployment fails. **Note** Static routes can only be configured when ZTP application is installed. Use one of this steps to resolve the issue: • To configure the static routes, from the main menu, select **Administration** > **Settings** > **Static Routes**. Click the + icon, enter the destination subnet IP address and mask (in slash notation), then click **Add**. • To configure the IP rule, log in to the Crosswork server and use this command: `ip rule add from all to 10.19.0.4 lookup cw_data`

View NSO function pack job history

View details and status for NSO function pack jobs in Crosswork Job History.

View NSO function pack job history in Crosswork, including job IDs, statuses, start and end times, task details, and descriptions.

The **Job History** tab shows detailed information of when jobs were started and ended, job ID, status, and other vital information.

Follow the steps below to view the details of the jobs.

Procedure

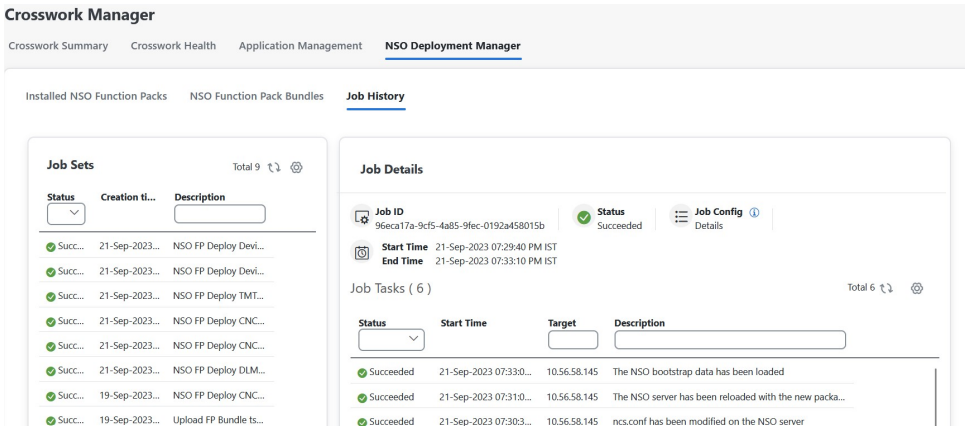
1. From the main menu, choose **Administration** > **Crosswork Manager**, click the **NSO Deployment Manager** tab.
2. Click the **Job History** tab.

In the **Job History** tab, the **Job Sets** pane displays the state of the job, job ID, and the job description. You can show or hide the columns based on the job creation time, status, and description.

3. In the **Job Sets** pane, select the job sets to view the associated job information in the **Job Details** pane. You can view the summary of the job tasks based on job task ID, task status, the task start and end time, and description.

   To view the job configuration information in JSON format, click the icon next to **Job Config**. A config window opens that lets you view the configuration in *View mode*.

Figure 43: Job history window



Install NSO function packs manually

Lists Cisco NSO function pack manuals for installation and configuration, including the transport, DLM, telemetry, change automation, and auto-discovery guides.

If you need to install or configure the NSO function packs manually, follow the instructions in these links:

Function pack guides

- *Cisco NSO Transport SDN Function Pack Bundle 7.2.x User Guide*
- *Cisco NSO Transport SDN Function Pack Bundle 7.2.x Installation Guide*
- *Cisco Network Services Orchestrator DLM Service Pack 7.2.x Installation Guide*
- *Cisco Crosswork NSO Telemetry Traffic Collector Function Pack 7.2.x Installation Guide*
- *Cisco Crosswork Change Automation NSO Function Pack 7.2.x Installation Guide*
- *Cisco NSO Auto Discovery Installation and User Guide 7.2.x*

Configure JSON web token authentication for NSO

Configure JSON Web Token (JWT) authentication so that Crosswork and NSO can use RESTCONF and JSON-RPC API workflows.

Use this task after the required Cisco NSO function packs are installed and before you validate Crosswork-to-NSO communication. The Crosswork certificate must be available on the NSO VM before you configure NSO to use the JWT authentication package.

Export the certificate from the browser and workstation that you use to access Crosswork. After you export the certificate, copy the certificate file to the NSO VM and complete the NSO-side configuration.

> **Warning**
>
> Do not perform this task until the required Cisco NSO function pack bundle is installed. The JWT authentication package is installed with the function pack bundle.

Procedure

1. In the browser that you use to access Crosswork, open the Crosswork website.
2. Click the padlock icon in the address bar, and then choose **Certificate is not Valid** > **View Certificate**.
3. In the **Certificate Viewer** window, click the **Details** tab.
4. Under **Certificate Hierarchy**, click **Crosswork**.
5. Click **Export**, choose a file name and location, and save the certificate as a Base64-encoded ASCII single certificate.

   Use the `.pem` extension. For example, save the file as `crosswork.pem`.

   > **Note**
   >
   > If you cannot save the file in `.pem` format, save it as a `.cer` file and use that file path in the bootstrap configuration.

6. Copy the certificate file to the NSO VM.

   > **Note**
   >
   > Make sure that the pem-key-path value and certificate file name are the same on the primary and secondary host.

7. On the NSO VM, open the `ncs.conf` file in the NSO directory, and add the JWT package authentication configuration under the `<aaa>` section.

   Example

   ```
   <aaa>
     <package-authentication>
       <enabled>true</enabled>
       <packages>
         <package>cisco-cfp-jwt-auth</package>
       </packages>
     </package-authentication>
   </aaa>
   ```

8. Restart NCS for the `ncs.conf` configuration to take effect.

Example

```
/etc/init.d/ncs restart
```

9. Configure the bootstrap authentication package by using `ncs_cli`.

   Log in to `ncs_cli`, enter configuration mode, and run these commands:

```
set jwt-auth cnc-host <Crosswork IP>
set jwt-auth port 30603
set jwt-auth pem-key-path /home/nso/crosswork.pem
commit dry-run
commit
```

12 Integrate SR-PCE

Topics:

- SR-PCE integration workflow
- SR-PCE certificate configurations
- Configure SR-PCE
- Add Cisco SR-PCE providers

Introduces the integration of Cisco SR-PCE with Crosswork Network Controller. It summarizes certificate setup, SR-PCE configuration, and provider onboarding.

This chapter covers the tasks and reference information for integrating SR-PCE with Crosswork Network Controller.

SR-PCE integration workflow

Summarizes the stages for integrating Cisco SR-PCE with Crosswork Network Controller.

This section explains the steps in integrating Cisco SR-PCE with Crosswork Network Controller.

Refer to the *Release Notes for Crosswork Network Controller, Release 7.2.1 SP* to know the SR-PCE version compatible with Crosswork Network Controller.

Table 79: SR-PCE integration stages

Stage	Action
1. Install the compatible version of Cisco SR-PCE	Select the type of SR-PCE and follow the relevant installation instructions in the Cisco IOS XRv 9000 Router Installation Guide.
2. Configure certificates	Refer to SR-PCE certificate configurations
3. Configure SR-PCE	Refer to Configure SR-PCE
4. Add SR-PCE provider and verify connectivity	Refer to Add Cisco SR-PCE providers

Use the related topics for certificate configuration, SR-PCE configuration, and provider setup.

SR-PCE certificate configurations

Provides certificate requirements and links to the default and custom certificate procedures for SR-PCE.

Before proceeding with the certificate configuration, determine your certificate requirements.

The default Cisco certificates are acceptable for Proof of concept or lab environments that require internal networks with basic security requirements.

If your organization requires custom certificates, you must complete the certificate configuration before attempting to onboard SR-PCE or devices. Integration will fail if certificates are not properly configured first.

Depending on whether you want to integrate custom or default certificates, refer to the relevant section.

- Configure SR-PCE with custom certificates or
- Configure SR-PCE with default certificates

Configure SR-PCE with default certificates

Create an SR-PCE credential profile, add the provider, and verify connectivity using default certificates.

Configure SR-PCE in your network using default Cisco certificates for authentication.

Before you begin

Use this procedure if you are using default Cisco certificates.

Procedure

1. Create SR-PCE credentials
   a) Navigate to **Administration > Credentials Management**.

b) Click **Add Credential Profile**.
c) Enter the required values:
- **Profile Name**: Enter a descriptive name.
- **Username**: SR-PCE admin username.
- **Password**: SR-PCE admin password.

d) Click Save.

2. Add SR-PCE provider and verify connectivity.
a) Navigate to **Administration** > **Manage Provider Access**, and then click +.
b) Enter the required values for the SR-PCE provider fields. For information about the fields, refer to the Add Cisco SR-PCE providers.
c) Click **Save** to add the SR-PCE provider.
d) Confirm that the SR-PCE provider shows a green Reachability status without any errors. You can also view the **Events** window (**Administration** > **Events**) to check whether the provider is configured correctly.
e) Repeat this process for each SR-PCE provider.

The SR-PCE provider(s) is added and validated using default certificates. Reachability status should indicate success, confirming proper configuration and connectivity.

Configure SR-PCE with custom certificates

Configure certificates and secure gRPC communication for SR-PCE, network devices, and Crosswork Network Controller.

Enable secure, encrypted communications for gRPC API in SR-PCE deployments using TLS, to protect data in transit and ensure compliance.

Configuring TLS for the SR-PCE gRPC API encrypts all API interactions. This security is essential for environments requiring data confidentiality and integrity, especially with Crosswork Network Controller device management. To enable secure gRPC communication, generate and install certificates on:

- All network devices (PCCs)
- SR-PCE
- Crosswork Network Controller

> **Note**
>
> Crosswork Network Controller supports only server authentication and not mutual authentication (client certificate validation).

Before you begin

Before configuring TLS for the SR-PCE gRPC API, make sure all prerequisites are met.

- Management IP addresses for all network devices (PCCs) and SR-PCE to add to the `subjectAltName` field in the certificate.
- Admin or privileged access to IOS XR devices
- SR-PCE access and readiness
  - SR-PCE is installed and operational, and

    - Administrative access to the SR-PCE CLI and file system is available.

- Certificates and key files
    - A root CA certificate is available (used for gNMI collection), and
    - Private key and certificate files are generated for SR-PCE.

Use this procedure to configure TLS on SR-PCE for gRPC API.

Procedure

---

1. Generate certificates using OpenSSL.
    a) On your Linux server, create a file named `extension.cnf`.
    b) Copy this content into the file and replace the parameters based on your environment.

```
[ req ]
distinguished_name = req_distinguished_name
req_extensions = req_ext
extensions = req_ext

[ req_distinguished_name ]
countryName = Country Name (2 letter code)
countryName_min = 2
countryName_max = 2
stateOrProvinceName = State or Province Name (full name)
localityName = Locality Name (eg, city)
organizationName = Organization Name (eg, company)
organizationalUnitName = Organizational Unit Name (eg, section)
commonName = Common Name (eg, fully qualified host name)
commonName_max = 64
emailAddress = Email Address
emailAddress_max = 64

[ req_attributes ]
challengePassword = A challenge password
challengePassword_min = 4
challengePassword_max = 20

[ req_ext ]
basicConstraints = critical,CA:true
subjectAltName = IP:<IP-address, IP-address...>

[ v3_ca ]
subjectKeyIdentifier = hash
authorityKeyIdentifier = keyid:always,issuer
basicConstraints = critical, CA:true
keyUsage = critical, digitalSignature, cRLSign, keyCertSign
```

    c) Replace the IP addresses in the `subjectAltName` line with your actual device and SR-PCE management IPs.
2. Generate the root CA key, device key, device Certificate Signing Request (CSR), sign the device certificate, and encrypt the device key.

> **Note**
>
> Record the password (`password1` in this example). You will need it later.

a) On your Linux server, run the command.

```
openssl genrsa -out rootCA.key 2048

openssl req -config extension.cnf -extensions v3_ca \
-subj /C=/ST=/L=/O=/CN=RootCA \
-x509 -days 365 -new -nodes -key rootCA.key -sha256 -out rootCA.pem

openssl genrsa -out device.key 2048

openssl req -subj /C=/ST=/L=/O=/CN=Device \
-new -config extension.cnf -key device.key -out device.csr

openssl x509 -extfile extension.cnf -extensions req_ext \
-req -days 365 -in device.csr \
-CA rootCA.pem -CAkey rootCA.key -CAcreateserial \
-out device.crt -sha256

openssl rsa -des3 -in device.key -out device.des3.key -passout
pass:password1
```

These are files are generated:

- `rootCA.key`
- `rootCA.pem`
- `device.key`
- `device.csr`
- `device.crt`
- `device.des3.key`

b) In your working directory, rename and prepare certificates for distribution.

```
mv device.crt ems.pem
mv device.key ems.key
mv rootCA.pem ca.cert
```

c) Using SCP, copy the three files (ems.pem, ems.key, ca.cert) to each device and SR-PCE **disk0:** location.

```
scp ems.pem admin@<device-ip>:/disk0:/
scp ems.key admin@<device-ip>:/disk0:/
scp ca.cert admin@<device-ip>:/disk0:/
```

Repeat this step for all devices and SR-PCE.

3. Install certificates on all devices and SR-PCE. For each device and SR-PCE:
   a) Log in via SSH or console.
   b) Back up any existing certificates.

c) Enable TLS in the gRPC configuration and disable mutual authentication to allow only server authentication.
d) Remove existing ems.pem, ems.key, and ca.cert files from /misc/config/grpc/ (if present).
e) Enter the run mode and navigate to the gRPC directory:

```
run
cd /misc/config/grpc/
```

Example

```
-rw-r--r--.  ems.pem
-rw-------.  ems.key
-rw-r--r--.  ca.cert
```

f) Copy the certificate files from disk0:

```
cp /disk0:/ems.pem .
cp /disk0:/ems.key .
cp /disk0:/ca.cert .
```

g) Verify file permissions and exit the run mode:

```
ls -lart
exit
```

h) Restart the EMS daemon:

```
process restart emsd
```

4. Enable secure gRPC on devices and SR-PCE. For each device and SR-PCE
a) Log in via SSH or console.
b) View the current gRPC configuration:

```
show running-config grpc
```

c) Enter global configuration mode and remove the `no-tls` directive.

Example

```
grpc
 port 57400
 no-tls          <-- Remove this line
 address-family ipv4
 service-layer
 !
```

Allow 1-2 minutes for the EMS daemon to restart and apply the changes.

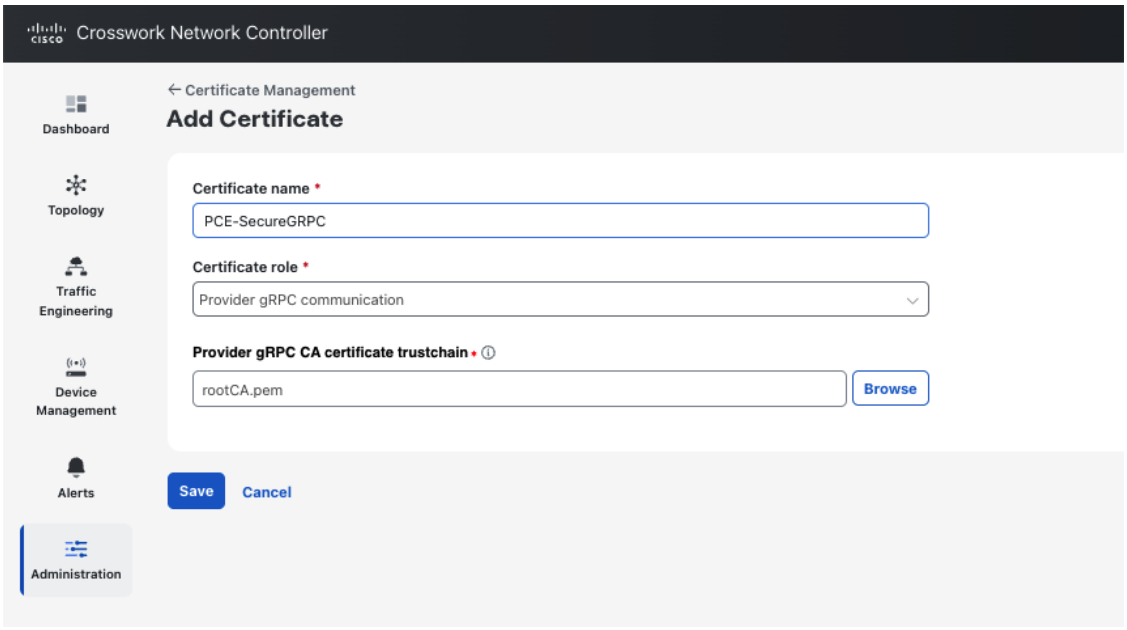
Secure gRPC is now enabled on the device or SR-PCE.

5. Create certificate profiles in Cisco Crosswork.
a) Log in to Crosswork Network Controller.
b) Navigate to **Administration > Certificate Management**.
c) Click **Create Certificate** and enter:

- Certificate name: Enter the certificate name as **PCE-SecureGRPC**.
- **Role**: Select **Provider gRPC communication**.

- **Provider gRPC CA certificate trustchain**: Upload rootCA.pem.

Figure 44: Add certificate configuration for PCE secure gRPC



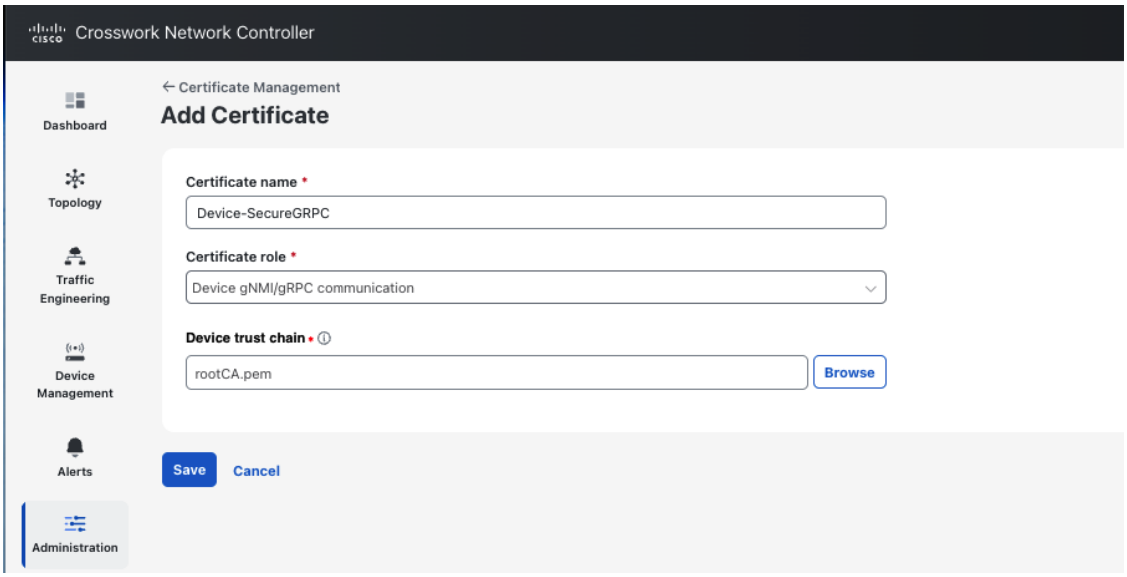
d) Click **Save**.

e) Click **Create Certificate** again and enter:

- **Certificate name**: Enter the certificate name as Device-SecureGRPC.
- **Role**: Select **Device gNMI/gRPC communication**.
- **Provider gRPC CA certificate trustchain**: Upload rootCA.pem.

f) Click **Save**.

Figure 45: Add certificate for secure device gRPC



6. Associate SR-PCE certificate and enable secure gRPC.

a) Log in to Crosswork Network Controller UI.

b) Navigate to **Administration > Manage Provider Access**.
c) Select the SR-PCE provider.
d) Set gRPC Mode to **GRPC_SECURE**.
e) Select the PCE-SecureGRPC certificate profile created in Step 5.
f) Click **Save**.

Crosswork will apply the certificate profile to the SR-PCE connection.

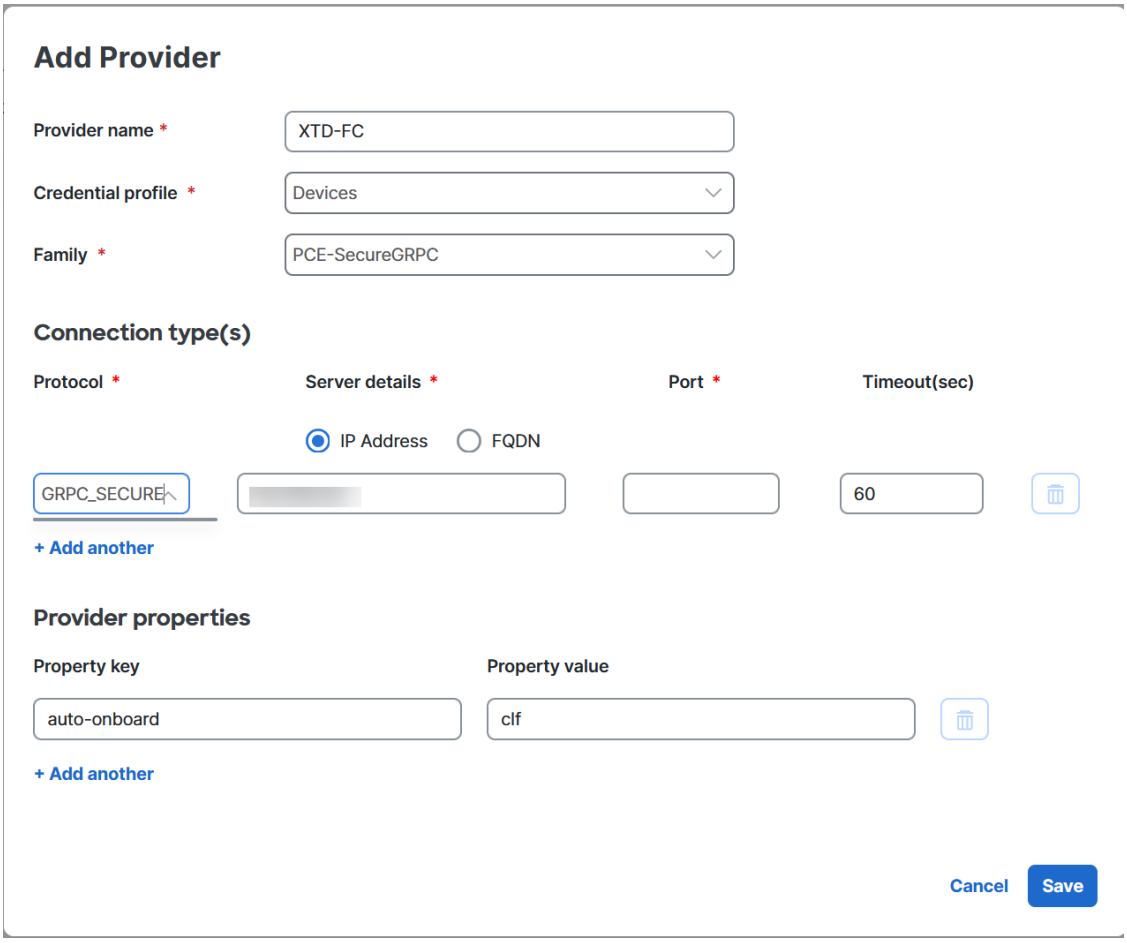
> **Note**
>
> If a gNMI certificate already exists and multiple trust chains are needed, update the existing .pem file to include all required CA certificates.

After successful addition, the gNMI certificate appears in the certificates listed in **Certificate Management > Certificates**. The gRPC configuration displays TLS enabled without mutual authentication.

7. Associate the certificate with SR-PCE.
   a) Navigate to **Administration > Manage Provider Access**.
   b) Select the SR-PCE provider.
   c) Select the **Family** as SR-PCE.
   d) Set the **gRPC Mode** to **GRPC_SECURE**.
   e) Select the certificate profile created earlier. For example, PCE-SecureGRPC.
   f) Click **Save**.

   Repeat these actions for all devices. Crosswork updates the connection settings for each device.

   Figure 46: Provider configuration



8. Verify secure gRPC configuration and connectivity.
   a) On each device and SR-PCE, verify the gRPC configuration:

   ```
   show running-config grpc
   ```

   Confirm that `no-tls` is absent and `port 57400` is present.

   b) In Crosswork Network Controller:
      - Navigate to Administration > Manage Provider Access and verify SR-PCE status shows "Connected" or "Active"
      - Navigate to Device Management > Network Devices and verify all devices show a "Connected" or "Active" status.

   c) (Optional) Review Crosswork application logs for successful gRPC TLS connections.

   Secure gRPC is now fully configured and operational.

After completing all the steps, you should observe:

- All devices and SR-PCE have certificates installed in `/misc/config/grpc/`.
- Device and SR-PCE gRPC configurations display `port 57400` without the `no-tls` directive.
- Cisco Crosswork shows all devices and SR-PCE as connected over secure gRPC (TLS).

For detailed provider management procedures including adding a provider, and advanced configurations, refer to the *Manage Provider Access* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

What to do next

After configuration, verify that TLS is functioning correctly.

1. Check the emsd process status:

```
show process emsd
```

2. Review gRPC service status:

```
show grpc
```

3. Check logs for TLS-related errors:

```
show logging | include TLS
show logging | include grpc
```

Example:

```
RP/0/RP0/CPU0:ASR9k-SRPCE1#show logging | include TLS
Mon Mar 16 17:01:41.239 UTC
RP/0/RP0/CPU0:Mar 16 16:57:50.669 UTC: emsd[1127]: %MGBL-EMS-6-EMSD_INFO :
gRPC is secure with TLS
```

```
RP/0/RP0/CPU0:ASR9k-SRPCE1#sh grpc
Mon Mar 16 16:57:50.593 UTC

Server name                              : DEFAULT
Address family                           : ipv4
Port                                     : 57400

Service ports
  gNMI                                   : none
  P4RT                                   : none
  gRIBI                                  : none

DSCP                                     : Default
TTL                                      : 64
VRF                                      : global-vrf
Server                                   : enabled
TLS                                      : enabled
TLS mutual                               : disabled
Trustpoint                               : none
Certificate Authentication               : disabled
Certificate common name                  : ems.cisco.com
TLS v1.0                                 : enabled
Maximum requests                         : 128
Maximum requests per user                : 10
Maximum streams                          : 32
Maximum streams per user                 : 32
Maximum concurrent streams               : 32
Memory limit (MB)                        : 1024
Keepalive time                           : 30
Keepalive timeout                        : 20
Keepalive enforcement minimum time       : 300
TLS Minimum Version                      : TLS 1.0
TLS Maximum Version                      : TLS 1.3

TLS cipher suites
  Default                                : none
```

```
  Default TLS1.3                          : aes_128_gcm_sha256
                                          : aes_256_gcm_sha384
                                          : chacha20_poly1305_sha256

  Enable                                  : none
  Disable                                 : none

  Operational enable                      : ecdhe-rsa-chacha20-poly1305
                                          : ecdhe-ecdsa-chacha20-poly1305
                                          : ecdhe-rsa-aes128-gcm-sha256
                                          : ecdhe-ecdsa-aes128-gcm-sha256
                                          : ecdhe-rsa-aes256-gcm-sha384
                                          : ecdhe-ecdsa-aes256-gcm-sha384
                                          : ecdhe-rsa-aes128-sha
                                          : ecdhe-ecdsa-aes128-sha
                                          : ecdhe-rsa-aes256-sha
                                          : ecdhe-ecdsa-aes256-sha
                                          : aes128-gcm-sha256
                                          : aes256-gcm-sha384
                                          : aes128-sha
                                          : aes256-sha
  Operational disable                     : none
Listen addresses                          : ANY
Latest Connection Id                      : 611159
Local Connection                          : Disabled
```

4. Test connectivity from a gRPC client using TLS.

For detailed certificate management procedures including renewal, revocation, and advanced configurations, refer to the *Manage Certificates* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Configure SR-PCE

Lists the configuration steps and commands for enabling gRPC on SR-PCE.

This section explains how to configure SR-PCE after you have installed it.

> **Note**
>
> The Cisco IOS XRv 9000 is the recommended platform to function as the SR-PCE. For additional details, refer to Configure Segment Routing Path Computation Element. Beginning with Crosswork Network Controller version 7.1, the gRPC API instrumentation for PCE is mandatory. For supported IOS XR platforms, contact Cisco Support.

Table 80: Configuring SR-PCE with gRPC enabled

Step	Command or Action	Description
1	**conf t** Example: `conf t`	Enters global configuration mode.

Step	Command or Action	Description

2	**lslib-server**
	Example:
	`lslib-server`
	`!`
		Enables the Link State Library (LSlib) server.

3	**grpc**
	Example:
	`grpc`
	`port 57400`
	`address-family ipv4`
	`tls profile <profile-name>`
	`service-layer`
	`!`
		Initiates the configuration mode for the gRPC server on the router. It is the top-level command for defining gRPC listener parameters. gRPC configuration enables a management and programmability interface for the router, which can be used to interact with various router components, including SR-PCE related data.
		• **port**–Specifies the TCP port number that the server will listen on for incoming client connections. Ideally gRPC server listening port is 57400, but it can be configured in the range: 57344-57999.
		• **address-family**– Specifies the address-family (such as IPv4 and IPv6) for the gRPC server to listen on.
		• **tls profile**–Enables TLS for secure communication. For secure gRPC deployment, TLS should be enabled.
		• **service-layer**–Defines that the gRPC server will expose operational data models and potentially configuration capabilities.

Step	Command or Action	Description

4	**linux networking**
	Example:
	`linux networking`
	` vrf default`
	`  address-family ipv4`
	`   default-route`
	`software-forwarding`
		Initiates the configuration mode for the Linux networking subsystem. It ensures that processes running in the Linux environment of the router (which might include the gRPC server) have their default route pointed to the IOS XR forwarding plane.
		• **vrf default**–Specifies that the subsequent routing configurations apply to the "default" Virtual Routing and Forwarding (VRF) instance.
		• **address-family**– Specifies the address-family (such as IPv4 and IPv6) for routing.
		• **default-route**–Sets the default route for traffic originating from the Linux kernel to be handled by the IOS XR routing engine. If `default-route mgmt` is not configured, the TPA's default route implicitly points to `software-forwarding`, which instructs the Linux kernel to use the `fwdintf` (forward interface) as the next-hop for its default route.

Step	Command or Action	Description

5	**pce**
	Example:
	`pce`
	` address ipv4`
	`192.168.0.1`
	` segment-routing`
	` state-sync ipv4`
	`192.168.0.3`
	` distribute link-state`
	` tcp-buffer size 1024000`
	` password encrypted pwd1`	

Step	Command or Action	Description
		Enables PCE and enters PCE configuration mode. • **address ipv4** *address*–Configures a PCE IPv4 address. • **segment-routing**–Enables Segment Routing capabilities for the PCE. • **state-sync ipv4** *address*–Configures the remote peer for state synchronization. • **distribute link-state** – Instructs the PCE to receive and build its topology database from the local Interior Gateway Protocol (IGP) (for example, OSPF or IS-IS). • **tcp-buffer size** *size*–Configures the transmit and receive TCP buffer size for each PCEP session, in bytes. The default buffer size is 256000. The valid range is from 204800 to 1024000. • **password** {**clear** | **encrypted**} *password*–Enables TCP MD5 authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured password will be rejected. Specify if the password is encrypted or clear text. Password can be configured globally under `pce` or per-peer under `peer <address>`. **Note** TCP-AO and TCP MD5 are never permitted to be used simultaneously. • **tcp-ao** *key-chain* [**include-tcp-options**] [**accept-ao-mismatch-connection**]–Enables TCP Authentication Option

Step	Command or Action	Description
		(TCP-AO) authentication for all PCEP peers. Any TCP segment coming from the PCC that does not contain a MAC matching the configured key chain will be rejected. TCP-AO can be configured globally under `pce` or per-peer under `peer <address>`. • **include-tcp-options**—Includes other TCP options in the header for MAC calculation. • **accept-ao-mismatch-connection**—Accepts connection even if there is a mismatch of AO options between peers. **Note** TCP-AO and TCP MD5 are never permitted to be used simultaneously.
6	For PCEP sessions: **peer** *SR-PCE loopback IP address* Example: `peer ipv4 209.165.255.2 password encrypted pwd1`	Specifies peer settings. • **peer**— Specifies the IP address of the SR-PCE peer (typically loopback). • **!**—Exits peer mode.

Step	Command or Action	Description
7	**segment-routing** {**strict-sid-only** | **te-latency**} Example: `segment-routing strict-sid-only`	Configures the segment routing algorithm to use strict SID or TE latency. **Note** This setting is global and applies to all LSPs that request a path from this controller.
8	**timers** Example: `timers keepalive 60 minimum-peer-keepalive 30 reoptimization 600 exit`	Enters timer configuration mode. • **keepalive** *time*–Configures the timer value for locally generated keep-alive messages. The default time is 30 seconds. • **minimum-peer-keepalive** *time*–Configures the minimum acceptable keep-alive timer that the remote peer may propose in the PCEP OPEN message during session establishment. The default time is 20 seconds. • **reoptimization** *time*–Configures the re-optimization timer. The default timer is 1800 seconds. • **exit**–Exits timer configuration mode and returns to PCE configuration mode.
9	**exit**	Exits PCE mode.
9	**commit**	Applies changes and exits the configuration.

What to do next:

Return to the installation workflow:

- VMware: VMware installation workflow
- AWS EC2: AWS EC2 installation workflow

Sample SR-PCE configurations

Provides sample SR-PCE, telemetry, traffic collector, and routing configurations.

These configurations are *examples* to guide you in a multiple SR-PCE setup for HA. Modify the examples for your deployment.

ISIS single topology configuration for dual-stack networks

Cisco Crosswork Network Controller supports ISIS Single Topology in addition to Multi-Topology. To use this, your XTC devices must be configured for ISIS Single Topology. For Single Topology configurations, only global IPv6 addressing is supported; support for link-local IPv6 addressing is not included.

Device-side configuration example

```
RP/0/RP0/CPU0:iosxrv-2(config)#router isis [NAME]
RP/0/RP0/CPU0:iosxrv-2(config-isis)#address-family ipv6 unicast
RP/0/RP0/CPU0:iosxrv-2(config-isis-af)#single-topology
```

Configuration requirements for deploying and reporting SR MSL policies to PCE

Enable gRPC on devices and for SR-TE policies

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config grpc
grpc
 segment-routing
  traffic-eng
   policy-service
  !
 !
 port 57400
 no-tls
```

Advertise all SR policies to BGP-LS peers

This configuration enables your router to report all configured SR MSL policies–both active and inactive–into the link-state database. As a result, these policies can be advertised via BGP-LS to controllers or peers, providing full visibility and supporting network orchestration.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config segment-routing traffic-eng
distribute link-state
segment-routing
 traffic-eng
  distribute link-state
   report-candidate-path-inactive
  !
 !
!
```

Prevent reporting MSL policies in PCEP

This configuration prevents SR MSL policies from being reported via PCEP. Since PCEP does not fully support MSL policies (it only advertises a single segment list, which can cause operational issues), it is recommended to remove the report-all command from the PCC configuration on the headend router.

```
RP/0/RP0/CPU0:L4-NCS560#sh running-config segment-routing traffic-eng pcc
segment-routing
 traffic-eng
  pcc
   source-address ipv4 192.100.0.4
   pce address ipv4 100.100.0.1
    precedence 25
   !
   pce address ipv4 100.100.0.2
    precedence 50
```

```
   !
   !  Remove this line to prevent reporting MSL policies to PCE
   !  report-all
   redundancy pcc-centric
   profile 1981
    autoroute
     include ipv4 all
     force-sr-include
    !
   !
  !
 !
!
```

Advertise SR MSL policies in link-state to PCE neighbor via BGP-LS

This configuration enables your router to advertise SR MSL policies in the link-state address family to a PCE neighbor over BGP. By establishing a BGP session with the PCE and including the `address-family link-state link-state` configuration, the router ensures that SR MSL policies are advertised and can be learned by the PCE.

> **Note**
>
> The link-state address family must be configured on both the headend and the PCE for successful exchange.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config router bgp
router bgp 60
 neighbor <NEIGHBOR_IP>   ! PCE neighbor
  remote-as 60
  update-source Loopback0
  address-family ipv4 unicast
   next-hop-self
  !
  address-family ipv6 unicast
  !
  address-family link-state link-state. ! Enable BGP-LS for SR MSL policy
advertisement
  !
 !
```

SRv6 data collection and traffic steering for DDM (Deterministic Demand Matrix) integration on Cisco IOS XR

Enable SRv6 locator accounting

This configuration enables the router to perform detailed accounting for IPv6 traffic specifically related to SRv6 locators. By tracking traffic on a per-prefix and per-nexthop basis, operators gain granular visibility into the usage and flow of SRv6-enabled services.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config accounting
accounting
 prefixes
  ipv6
   mode per-prefix per-nexthop srv6-locators
  !
 !
!
```

Enable SRv6 accounting data to telemetry

This configuration sets up model-driven telemetry on the router to stream SRv6 accounting data to external collectors. By defining specific sensor paths, the router can push operational data related to SRv6 locator accounting, enabling real-time monitoring, analysis, and orchestration of SRv6 network performance and traffic patterns.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config telemetry model-driven
telemetry model-driven
 sensor-group cisco_models
  sensor-path
Cisco-IOS-XR-infra-xtc-agent-oper:xtc/forwarding/policy-forwardings/policy-forwarding

  sensor-path
Cisco-IOS-XR-fib-common-oper:cef-accounting/vrfs/vrf[vrf-name='default']/afis/afi[afi-type=ipv6]/pfx/srv6locs/srv6loc

 !
!
```

Enable customer/VRF traffic steering to SRv6 locators via BGP

This configuration enables an edge router to steer customer or VRF (Virtual Routing and Forwarding) IPv4 and IPv6 traffic into specific SRv6 locators using BGP.

```
RP/0/RP0/CPU0:L1-NCS5501#sh running-config router bgp
router bgp 60
 bgp router-id <ROUTER_ID_IP>
 segment-routing srv6
  locator L1algo0
 !
 address-family ipv4 unicast
  network <ROUTER_ID_IP>/32
 !
 address-family vpnv4 unicast
  vrf all                ! If there are multiple VRF where traffic is
ingressing, add srv6 locator in vrf all.
   segment-routing srv6
    locator L1algo0
    alloc mode per-vrf
   !
  !
 !
 vrf ntt
  rd 200:200
  address-family ipv4 unicast
   segment-routing srv6   ! If there is only one VRF where traffic is
ingressing, add srv6 locator in this vrf alone, if there is no VRF, then add
the locator in neighbor address family
    locator L1algo0
    alloc mode per-vrf
   !
   redistribute connected
  !
  neighbor <NEIGHBOR_IP>
   remote-as 61
   update-source GigabitEthernet0/0/0/0
   address-family ipv4 unicast
    route-policy PASS_ALL in
    route-policy PASS_ALL out
   !
  !
 !
```

Verify SRv6 traffic steering via CEF accounting

This command is used to verify that IPv6 traffic is being steered into SRv6 locators, rather than MPLS labels, by inspecting the CEF accounting statistics. It provides granular visibility, showing packet and byte counts for specific IPv6 prefixes that are associated with SRv6 locators.

```
sh cef ipv6 accounting
fccc:cc3e:3::/48
Accounting: 0/0 packets/bytes output (per-prefix-per-path mode)
 via fe80::2/128, Bundle-Ether1201
  path-idx 0
  next hop fe80::2/128
  Accounting: 200000/58400000 packets/bytes output  <<< Traffic packets for
prefix fccc:cc3e:3::
```

Other sample SR-PCE configurations

Redundant SR-PCE configuration (on PCE with Cisco IOS-XR 7.x.x)

```
pce
 address ipv4 100.100.0.7
 state-sync ipv4 100.100.0.1
 api
  sibling ipv4 100.100.0.1
```

PCE configuration for enabling gRPC API on XR 25.2.1.x (IPv4 deployment)

```
conf t
  lslib-server
  !
  grpc
    port 57400
    no-tls
    address-family ipv4
    service-layer
    !
!
pce
  distribute link-state
  !
!
linux networking
  vrf default
    address-family ipv4
      default-route software-forwarding
    !
    address-family ipv6
      default-route software-forwarding
    !
  !
!
commit
```

Note

For secure gRPC deployment, remove `no-tls`.

Configure distribute link-state on all PCEs to inject SR policies into BGP-LS.

Enable gRPC API on XR 25.2.1.x (IPv6 deployment)

```
conf t
  lslib-server
  !
  grpc
    port 57400
    no-tls
    address-family ipv6
    service-layer
    !
!
pce
  distribute link-state
  !
!
linux networking
  vrf default
    address-family ipv4
      default-route software-forwarding
    !
    address-family ipv6
      default-route software-forwarding
    !
  !
!
commit
```

Note

For secure gRPC deployment, remove `no-tls`.

Configure distribute link-state on all PCEs to inject SR policies into BGP-LS.

Verify whether the topology is published in gRPC

```
sh lslib server topology-db
```

Verify the SR-MPLS LSP published in gRPC

```
show lslib server topology-db detail protocol sr
```

Redundant SR-PCE configuration (PCC)

```
segment-routing
 traffic-eng
```

```
  pcc
   source-address ipv4 100.0.0.1
   pce address ipv4 100.0.0.2
    precedence 200
   !
   pce address ipv4 100.0.0.3
    precedence 100
   !
   report-all
   redundancy pcc-centric
```

Redundant SR-PCE configuration (on PCC) for RSVP-TE

> **Note**
>
> `Loopback0` represents the TE router ID.

```
ipv4 unnumbered mpls traffic-eng Loopback0
!
mpls traffic-eng
 pce
  peer source ipv4 200.100.200.1
  peer ipv4 209.165.0.6
   precedence 200
  !
  peer ipv4 100.100.0.0
   precedence 100
  !
  stateful-client
   instantiation
   report
   redundancy pcc-centric
   autoroute-announce
  !
 !
 auto-tunnel pcc
  tunnel-id min 1000 max 1999
 !
!
```

Sample telemetry configurations

SR-TM configuation

```
telemetry model-driven
 destination-group crosswork
  address-family ipv4 5.5.5.5 port 9000
   encoding self-describing-gpb
   protocol tcp
  !
 !
 sensor-group SRTM
  sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/afs/af/counters/tunnels
  sensor-path
```

```
Cisco-IOS-XR-infra-tc-oper:traffic-collector/vrf-table/default-vrf/afs/af/counters/prefixes

 !
 subscription OE
  sensor-group-id SRTM sample-interval 60000
  destination-id crosswork
  source-interface Loopback0
!
traffic-collector
 interface GigabitEthernet0/0/0/3
 !
 statistics
  history-size 10
```

Note

The destination address uses the southbound data interface (eth1) address of the Crosswork Data Gateway VM.

It is required to push sensor path on telemetry configuration via NSO to get prefix and tunnel counters. It is assumed that the Traffic Collector has been configured with all the traffic ingress interface. This configuration is needed for demands in the Bandwidth on Demand feature pack to work.

Telemetry sensor path

```
sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/afs/af/counters/tunnels/tunnel
sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/vrf-table/default-vrf/afs/af/counters/prefixes/prefix
```

Telemetry configuration pushed by Crosswork Network Controller to all the headend routers via NSO

```
telemetry model-driven
  destination-group CW_43dc8a5ea99529715899b4f5218408a785e40fce
    vrf default
    address-family ipv4 5.5.5.5 port 30500
      encoding self-describing-gpb
      protocol top
    !
  !
destination-group CW_4b3c69a200668b0a8dc155caff295645c684a8f8
  vrf default
  address-family ipv4 5.5.5.5 port 30500
    encoding self-describing-gpb
    protocol top
  !
!
sensor-group CW_43dc8a5ea99529715899b4f5218408a785e40fce
  sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/afs/af/counters/tunnels/tunnel
!
sensor-group CW_4b3c69a200668b0a8dc155caff295645c684a8f8
  sensor-path
Cisco-IOS-XR-infra-tc-oper:traffic-collector/vrf-table/default-vrf/afs/af/counters/prefixes/prefix
!
```

```
subscription CW_43dc8a5ea99529715899b4f5218408a785e40fce
  sensor-group-id CW_43dc8a5ea99529715899b4f5218408a785e40fce sample-interval
 300000
  destination-id CW_43dc8a5ea99529715899b4f5218408a785e40fce
!
subscription CW_4b3c69a200668b0a8dc155caff295645c684a8f8
  sensor-group-id CW_4b3c69a200668b%a8dc155caff295645c684a8f8 sample-interval
 300000
  destination-id CW_463c69a200668b0a8dc155caff295645c684a8f8
  !
!
```

Traffic Collector configurations

Traffic Collector configurations (all Ingress traffic interface to be added below in the Traffic Collector)

```
RP/0/RSP0/CPU0:PE1-ASR9k#sh running-config traffic-collector
Fri May 22 01:14:35.845 PDT
traffic-collector
  interface GigabitEthernet0/0/0/0
  !
  statistics
    history-size 1
    collection-interval 1
    history-timeout 1
    history-minute-timeout
  !
!
```

Add BGP neighbor next-hop-self for all the prefix (to show TM rate counters)

```
bgp router-id 5.5.5.5
address-family ipv4 unicast
  network 5.5.5.5/32
  redistribute static
!
address-family link-state link-state
!
neighbor 1.1.1.1
  remote-as 65000
  update-source Loopback0
  address-family ipv4 unicast
   next-hop-self
  !
!
```

Traffic collector tunnel and prefix counters

```
RP/0/RSP0/CPU0:PE1-ASR9k#show traffic-collector ipv4 counters prefix
Fri May 22 01:13:51.458 PDT
Prefix              Label         Base rate         TM rate        State
                                  (Bytes/sec)       (Bytes/sec)
------------------  ------------  ---------------  --------------
------------------
1.1.1.1/32          650001         3                0              Active
2.2.2.2/32          650002         3                0              Active
3.3.3.3/32          650003         6                0              Active
4.4.4.4/32          650004         1                0              Active
6.6.6.6/32          650200         6326338          6326234        Active
7.7.7.7/32          650007         62763285         62764006       Active
8.8.8.8/32          650008         31129168         31130488       Active
```

```
9.9.9.9/32          650009         1                0              Active
10.10.10.10/32      650010         1                0              Active
RP/0/RSP0/CPU0:PE1-ASR9k#stt
RP/0/RSP0/CPU0:PE1-ASR9k#show traffic-collector ipv4 counters tunnel
Fri May 22 01:13:52.169 PDT
RP/0/RSP0/CPU0:PE1-ASR9k#]
```

Add Cisco SR-PCE providers

Add Cisco SR-PCE providers to discover and manage SR-PCE devices in Crosswork Network Controller.

Add Cisco SR-PCE providers to supply device discovery, management, configuration maintenance, and route calculation services to Crosswork Network Controller.

Cisco Segment Routing Path Computation Elements (Cisco SR-PCE) providers:

- supply device discovery, management, configuration maintenance, and route calculation services to Crosswork Network Controller components
- enable system access as part of SDN controllers in the management domain, and
- support multi-AS topology and path calculations.

Multi-AS topology and path calculations are supported if the complete topology is accessible to both the Crosswork Network Controller and each PCE. A single PCE cannot view a specific AS topology while another PCE views a different topology. Each PCE must have access to the entire topology view. To learn and discover SR policies, Layer 3 links, and devices, at least one SR-PCE provider is required. Additionally, a second SR-PCE can be configured as a backup.

Follow the steps below to add one or more instances of Cisco SR-PCE as providers.

Before you begin

Ensure these requirements are met prior to adding an SR-PCE provider:

- Configure a device to act as the SR-PCE. Refer to SR configuration documentation for your specific device platform to enable SR (for IS-IS or OSPF protocols) and configure an SR-PCE (for example: *Segment Routing Configuration Guide for Cisco NCS 540 Series Routers*).
- Create a credential profile for the Cisco SR-PCE provider with these connection types. Refer to *Create credential profiles* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* for more details:
  - gRPC–Required to discover topology, SR-MPLS, and SRv6 policies. Refer to Sample PCE configuration for enabling gRPC API on XR for configuration examples.
  - Basic HTTP text-authentication– Required to process for RSVP, TreeSID and PCEP sessions. MD5 authentication is currently not supported.

    If the Cisco SR-PCE server you are adding does not require authentication, you must still supply a credential profile for the provider. The credential profile can be any profile that does not use the HTTP protocol.
- If you plan to set up gRPC with Transport Layer Security (TLS), a certificate must be generated and added with the **Secure gRPC communication** role. The certificate is used to secure communication through TLS between gRPC clients and the EMS server. The client should use ems.pem and ca.cert to initiate the TLS authentication. To update the certificate, ensure to copy the new certificate that has been generated earlier to the location and restart the server. Refer to *Manage Certificates* chapter in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* for more details.
- For high availability, ensure that you set up two separate Cisco SR-PCE providers with unique names and IP addresses, but with matching configurations.

- Know the name you want to assign to the Cisco SR-PCE provider. This is usually the DNS hostname of the Cisco SR-PCE server.
- Know the Cisco SR-PCE server IP address.
- Know the interface you want to use to communicate between Cisco SR-PCE and the Crosswork application server.
- Know the SSH credentials for the PCE device to enable gRPC communication. Note that PCE API credentials are used exclusively for HTTP-based communication.

For device auto-onboarding:

- Determine whether you want to auto-onboard the devices that Cisco SR-PCE discovers and, if so, whether you want the new devices to have their management status set to **`off`**, **`managed`** or **`unmanaged`** when added.
- If you plan to auto-onboard devices that the Cisco SR-PCE provider discovers, and set them to a managed state when they are added to the database:
  - Assign an existing credential profile for communication with the new managed devices.
  - The credential profile must be configured with an SNMP protocol.

Procedure

1. From the main menu, choose **Administration** > **Manage Provider Access** >  .
2. Enter these values for the SR-PCE provider fields:
   a) Required fields:
      - **Provider Name**: Name of the SR-PCE provider.
      - **Credential Profile**: Select the previously created Cisco SR-PCE credential profile.
      - **Family**: Select **`SR_PCE`**. All other options should be ignored.
      - **Protocol**:

        Note

        HTTP is mandatory when RSVP TE is configured in the network, and gRPC is mandatory for SR TE and L3 links discovery.

        - Select **`HTTP`** and enter required fields. HTTP is required to process RSVP, TreeSID and PCEP sessions. The default port is 8080.
        - Select **`GRPC`** or **`GRPC_SECURE`** (gRPC with Transport Layer Security (TLS)) and enter required fields. These settings are required to process topology, SR-MPLS, and SRv6 policies. Only one of these options can be used. If GRPC_SECURE is selected, you must provide the trusted certificate in the **Certificate profile** field

      - **Provider Properties**: Enter property keys and values:

Table 81: SR-PCE provider property values

Property Key	Value

`auto-onboard`	`off`
	
	Note
	Use this option if you plan to manually (via UI or CSV import) enter all of your network devices.
	
	When devices are discovered, the device data is recorded in the Cisco SR-PCE database, but is not registered in Crosswork Network Controller Inventory Management database.
	
	`unmanaged`
	If this option is enabled, all devices that Crosswork Network Controller discovers will be registered in the Crosswork Network Controller Inventory Management database, with their configured state set to `unmanaged`. SNMP polling will be disabled for these devices, and no management IP information will be included. To get these devices into the `managed` state later, you will need to either edit them via the UI or export them to a CSV make modifications and then import the updated CSV. You can also assign credential profiles by adding them to the device CSV file before import (the credential profiles must already exist).

Property Key	Value
	`managed` If this option is enabled, all devices that Cisco SR-PCE discovers will be registered in the Crosswork Network Controller Inventory Management database, with their configured state set to **`managed`**. Typically suitable for an environment that has same device profiles, devices are managed by their TE router-ID, and all devices can be discovered by the Cisco SR-PCE. SNMP polling will be enabled for these devices, and Cisco SR-PCE will also report the management IP address (TE Router ID for IPv4, or IPv6 Router ID for IPv6 deployment). The devices will be added with the credential profile associated with the device-profile key in the SR-PCE provider configuration. If you enable this option for IPv6 deployment, devices will still register as **unmanaged** in the inventory. When you delete an onboarded device that was added via SR-PCE discovery with auto-onboard set to **managed**, the topology service adds it again as **unmanaged**. This ensures that devices that have been removed are not automatically managed again unless they acquire a new TE-ID. To manage a rediscovered device, update its status manually.

Property Key	Value
`device-profile`	The name of a credential profile that contains SNMP credentials for all the new devices. **Note** This field is necessary only if `auto-onboard` is set to `managed` or `unmanaged`. If the **auto-onboard** is set to **managed** and there is no valid **device-profile** set, the device will instead be onboarded as **unmanaged**.
`outgoing-interface`	`eth1` **Note** You have to set this only if you want to enable Crosswork application access to SR-PCE via the data network interface when using the two NIC configuration.
`preferred-stack`	`ipv4` or `ipv6` or `NOT SET`. This property is applicable only for dual stack deployments.
`pce`	`off` or `on`. This is an optional property. If not specified, the default value is `on`. If value is specified as `off`, it means that LSPs and policies are not accessible for the SR-PCE provider.

> **Note**
>
> Topology can be visualized even with **auto-onboard** as **off** and a **device-profile** is not specified.
>
> If **managed** or **unmanaged** options are set and you want to delete a device later, you must either:.
>
> - Reconfigure and remove the devices from the network before deleting the device from Crosswork Network Controller. This avoids Crosswork Network Controller from rediscovering and adding the device back.
> - Set auto-onboard to **off**, and then delete the device from Crosswork Network Controller. However, doing so will not allow Crosswork Network Controller to detect or auto-onboard any new devices in the network.
>
> It is not recommended to modify **auto-onboard** options once set. If you need to modify them:
>
> 1. Delete the provider and wait until deletion confirmation is displayed in the Events window.
> 2. Re-add the provider with the updated **auto-onboard** option.
> 3. Confirm the provider has been added with the correct **auto-onboard** option in the Events window.

b) Optional values:

- **Timeout**: The amount of time (in seconds) to wait before timing out the connection to the SR-PCE server. The default is 30 seconds.

3. Click **Save** to add the SR-PCE provider.
4. Confirm that the SR-PCE provider shows a green Reachability status without any errors. You can also view the Events window (**Administration** > **Events**) to check whether the provider is configured correctly.
5. Repeat this process for each SR-PCE provider.

What to do next

- If `auto-onboard` is set to `off`, navigate to **Device Management** > **Network Devices** to onboard devices.
- If you opted to automatically onboard devices, navigate to **Device Management** > **Network Devices** to view the device list. To add more node information such as geographical location details, export the device list (.csv), update it, and import it back. If geographical location data is missing, you will only be able to see device topology using the logical map.

Return to the installation workflow:

- VMware: VMware installation workflow
- AWS EC2: AWS EC2 installation workflow

Part VII

Install Cisco Crosswork Network Controller on a Single VM

Topics:

- Install Crosswork Network Controller on a Single VM

13 Install Crosswork Network Controller on a Single VM

Topics:

- Introduction
- Installation requirements
- Installation parameters
- Install Crosswork Network Controller using the vCenter vSphere UI
- Install Crosswork Network Controller via the OVF Tool
- Install Crosswork Network Controller using the Docker installer tool
- Install Crosswork Network Controller VM using CLI
- Post installation tasks (for Premier tier installation)

Introduces single-VM installation of Crosswork Network Controller. Summarizes the requirements, parameters, and installation methods covered in this chapter.

This chapter describes how to install Crosswork Network Controller on a single VM.

The chapter covers installation requirements and parameters, vSphere UI and CLI installation methods, KVM configuration, and post-installation NSO function pack deployment.

Introduction

Introductory overview of single-VM deployment and the available product tiers.

Crosswork Network Controller single-VM deployment is an installation option that runs on a single VM or node.

This chapter explains the requirements and processes to install Crosswork Network Controller on a single VM or node. Cisco Crosswork Network Controller enables you to proactively manage your end-to-end networks, by providing automation solutions to ensure faster innovation, optimal user experience, and operational excellence.

Crosswork Network Controller, when deployed on a single VM, is available in these tiers:

- **Essentials**: includes Element Management Functions and embedded Collectors.
- **Advantage**: includes the **Essentials** package, Optimization Engine, Active Topology, Service Health, and embedded NSO.
- **Premier**: includes the **Advantage** package, and Crosswork Workflow Manager Solutions.

> **Note**
>
> Crosswork Network Controller Add-on package (Change Automation and Health Insights) is not supported for single VM deployments.

Installation requirements

Lists host, hypervisor, VM, network, and port requirements for single-VM deployment.

these sections list the host, VMware, KVM, VM, network, and port requirements for single-VM deployment.

Refer to Physical server requirements on page 29 for more information.

- Hypervisor and vCenter supported:
  - VMware vCenter Server 8.0 (U2c or later) and ESXi 8.0 (U2b or later)
  - VMware vCenter Server 7.0 (U3p or later) and ESXi 7.0 (U3p or later)
- Crosswork Network Controller VM must be hosted on hardware with Hyper Threading disabled.
- Ensure that profile-driven storage is enabled by the vCenter admin user. Query permissions for the vCenter user at the root level (for all resources) of the vCenter.
- We also recommend you to enable vCenter storage control.
- The networks required for the Crosswork Management and Data networks need to be built and configured in the data centers, and must allow low latency L2 communication (latency with RTT <= 10 ms).
- Ensure the user account you use for accessing vCenter has these privileges:
  - VM (Provisioning): Clone VM on the VM you are cloning.
  - VM (Provisioning): Customize on the VM or VM folder if you are customizing the guest operating system.
  - VM (Inventory): Create from the existing VM on the data center or VM folder.
  - VM (Configuration): Add a new disk on the data center or VM folder.

- Resource: Assign a VM to a resource pool on the destination host or resource pool.
- Datastore: Allocate space on the destination datastore or datastore folder.
- Network: Assign the network to which the VM will be assigned.
- Profile-driven storage (Query): This permission setting needs to be allowed at the root of the data center tree level.

these requirements are mandatory if you are planning to install Crosswork Network Controller on RHEL KVM.

Table 82: Host bare metal requirements (per bare metal server)

Component	Minimum requirement per node
Processor	Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz or latest
NIC	2 x 10 Gbps NICs
OS	Red Hat Enterprise Linux 8.10 Red Hat Enterprise Linux 9.4

Note

- Ensure the networks required for the Crosswork Management and Data networks are built and configured in the data centers. These networks must allow low-latency L2 communication with a round-trip time (RTT) of 10 ms or less.
- The same network name must be used and configured on the RHEL bare metal host machine that is hosting the Crosswork Network Controller VM.
- Crosswork Network Controller VM must run on hardware with Hyper Threading disabled to ensure consistent, real-time performance for CPU-intensive workloads, as Hyper Threading may cause resource contention and unpredictable performance.
- You cannot reuse the bare metal IP addresses for Crosswork VM.

Refer to Resource footprint for single VM deployments on page 28 for information on the resource requirements for single VM deployments.

Table 83: Network requirements

Requirement	Description
Network Connections	For production deployments, we recommend that you use dual interfaces, one for the Management network and one for the Data network. For optimal performance, the Management and Data networks should use links configured at a minimum of 10 Gbps with a latency of less than 10 milliseconds.

Requirement	Description
IP Addresses	**4 IPv4 and/or IPv6 addresses**: A management and data IP address for the Hybrid node being deployed, and two additional IP addresses to be used as the Virtual IP (VIP) address (one for the Management network and one for the Device network). **Note** • The IP addresses must be able to reach the gateway address for the network, or the installation will fail. • When deploying with IPv6 or dual stack, the installation needs to run on an IPv6 enabled container/VM. • At this time, your IP allocation is permanent and cannot be changed without re-deployment. For more information, contact the Cisco Customer Experience team.
Interfaces	Crosswork is deployed on a single VM with **2 interfaces**. • **No. of NICs**: 2 • **vNIC0**: Management Traffic (for accessing the interactive console and passing the Control/Data information between servers). • **vNIC1**: Device Access Traffic (for device access and data collection). **Note** Due to security policies, traffic from subnets of a vNIC received on other vNICs is dropped. For example, in a setup with two vNICs, all device traffic (incoming and outgoing) must be routed through the default vNIC1.
NTP Server	The IPv4 and/or IPv6 addresses or host names of the NTP server you plan to use. If you want to enter multiple NTP servers, separate them with spaces. These should be the same NTP servers you use to synchronize the Crosswork application VM clock, devices, clients, and servers across your network. Ensure that the NTP servers are reachable on the network before attempting installation. The installation will fail if the servers cannot be reached.

Requirement	Description
DNS Servers	The IPv4 and/or IPv6 addresses of the DNS servers you plan to use. These should be the same DNS servers you use to resolve host names across your network. Ensure that the DNS servers are reachable on the network before attempting installation. The installation will fail if the servers cannot be reached.
DNS Search Domain	The search domain you want to use with the DNS servers, for example, cisco.com. You can have only one search domain.
Backup Server	Cisco Crosswork will back up the configuration of the system to an external server using SCP. The SCP server storage requirements will vary slightly but you must have at least 25 GB of available storage.
FQDN (Optional)	The installation process supports using either a VIP (Virtual IP address) or an FQDN (Fully Qualified Domain Name) to access the VM. If you choose to use the FQDN, you will need one for the Management and one for the Data network. **Note** If you choose to supply the FQDNs during the initial installation, the DNS server must be populated with them before the VM is powered on; otherwise, the installation script will fail to complete the environment setup.

For information on the ports used in single VM deployments, refer to Single VM communication on page 64.

Installation parameters

Lists the parameters required to install the Crosswork Network Controller VM.

The installation parameters in this topic define the values to provide when you install the Crosswork VM.

This section explains the important parameters that must be specified while installing the Crosswork VM. Kindly ensure that you have relevant information to provide for each of the parameters mentioned in the table.

Note

Please use the latest template file that comes with the Crosswork build file.

Note

In case of dual stack deployment, you need to configure the IPv4 and IPv6 values for the Management, Data, and DNS parameters.

- `ManagementIPv4Address, ManagementIPv6Address`
- `ManagementIPv4Netmask, ManagementIPv6Netmask`
- `ManagementIPv4Gateway, ManagementIPv6Gateway`
- `ManagementVIPv4, ManagementVIPv6`
- `DataIPv4Address, DataIPv6Address`
- `DataIPv4Netmask, DataIPv6Netmask`
- `DataIPv4Gateway, DataIPv6Gateway`
- `DataVIPv4, DataVIPv6`
- `DNSv4, DNSv6`

Table 84: General parameters

Parameter Name	Description
`ClusterIPStack`	The IP stack protocol: IPv4 or IPv6
`ManagementIPAddress`	The Management IP address of the VM (IPv4 or IPv6).
`ManagementIPNetmask`	The Management IP subnet in dotted decimal format (IPv4 or IPv6).
`ManagementIPGateway`	The Gateway IP on the Management Network (IPv4 or IPv6). The address must be reachable, otherwise the installation will fail.
`ManagementVIP`	The Management Virtual IP for the Crosswork VM.
`DataIPAddress`	The Data IP address of the VM (IPv4 or IPv6).
`DataIPNetmask`	The Data IP subnet in dotted decimal format (IPv4 or IPv6).
`DataIPGateway`	The Gateway IP on the Data Network (IPv4 or IPv6). The address must be reachable, otherwise the installation will fail.

Parameter Name	Description
`DataVIP`	The Data Virtual IP for the Crosswork VM.
`DNS`	The IP address of the DNS server (IPv4 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 VM.
`CWPassword`	Password to log into Cisco Crosswork. When setting up a VM, ensure the password is strong and meets these criteria: • It must be at least 8 characters long and include uppercase and lowercase letters, numbers, and at least one special character. • these special characters are not allowed: backslash (\), single quote ('), or double quote ("). • 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.
`VMSize`	Size of the VM. Crosswork Network Controller supports both the `"XLarge"` and `"Large"` profiles. For more information, refer to Resource footprint for single VM deployments on page 28.
`VMName`	Name of the VM.
`NodeType`	Indicates the type of VM. Choose **Hybrid**.
`IsSeed`	Set to `"True"`.
`InitNodeCount`	Set value to **1**.
`InitMasterCount`	Set value to **1**.

Parameter Name	Description
`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. **Note** The final backup partition size will be calculated dynamically. This parameter defines the minimum.
`ThinProvisioned`	Set to **false** for production deployments.
`SchemaVersion`	The configuration Manifest schema version. This indicates the version of the installer to use with this template. Schema version should map to the version packaged with the sample template in the installer tool on cisco.com. 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.
`LogFsSize`	Log partition size (in Giga Bytes). Minimum value is 20 GB and Maximum value is 1000 GB. If left blank, the default value (20 GB) is selected.
`EnableSkipAutoInstallFeature`	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"`. For single VM deployment, you must set the value as `"True"`. **Note** • If left blank, the default value (`"False"`) is automatically selected. • This parameter accepts a string value, so be sure to enclose the value in double quotes.
`EnforcePodReservations`	Enforces minimum resource reservations for the pod. If left blank, the default value (`"True"`) is selected. This parameter accepts a string value, so be sure to enclose the value in double quotes.

Parameter Name	Description
`K8sServiceNetwork`	The network address for the kubernetes service network. By default, the CIDR range is fixed to '/16'.
`K8sPodNetwork`	The network address for the kubernetes pod network. By default, the CIDR range is fixed to '/16'.
`IgnoreDiagnosticsCheckFailure`	Used to set the system response in case of a diagnostic check failure. 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. You are recommended to select the default value. This parameter accepts a string value, so be sure to enclose the value in double quotes. **Note** • The log files (`diagnostic_stdout.log` and `diagnostic_stderr.log`) can be found at **/var/log**. The result from each diagnostic execution is kept in a file at **/home/cw-admin/diagnosis_report.txt**. • Use **diagnostic all** command to invoke the diagnostic manually on day N. • Use **diagnostic history** command to view previous test report.
`ManagementVIPName`	Name of the Management Virtual IP for the Crosswork VM. This is an optional parameter used to reach Crosswork Management VIP via DNS name. If this parameter is used, the corresponding DNS record must exist in the DNS server.
`DataVIPName`	Name of the Data Virtual IP for the Crosswork VM. This is an optional parameter used to reach Crosswork Data VIP via DNS name. If this parameter is used, the corresponding DNS record must exist in the DNS server.

Parameter Name	Description
`EnableHardReservations`	Determines the enforcement of VM CPU and Memory profile reservations. This is an optional parameter and the default value is `"True"`, if not explicitly specified. This parameter accepts a string value, so be sure to enclose the value in double quotes. 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. 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.
`ManagerDataFsSize`	This parameter is applicable only when installing with the Docker installer tool. Refers to the data disk size for the Crosswork node (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.
`RamDiskSize`	Size of the Ram disk. **Warning** Use the `RamDiskSize` option only for temporary lab testing or product demonstrations. This option is intended for environments where disk performance does not meet the application requirements and where rebuilding the VMs after a restart is an acceptable trade-off. Do not use this option for deployments that require resiliency or reliability. The minimum value for `RamDiskSize` is `2`. When you set `RamDiskSize` to a non-zero value, leave `HSDatastore` unset as `""`. If the VM is powered down or restarted after installation with `RamDiskSize` configured, the system may be unable to launch applications, and Kubernetes may report errors.
`Timezone`	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. This is an optional parameter.

Table 85: VMware template parameters

Parameter Name	Description
`VCenterAddress`	The vCenter IP or host name.
`VCenterUser`	The username needed to log into vCenter.
`VCenterPassword`	The password needed to log into vCenter.
`DCname`	The name of the Data Center resource to use. Example: `DCname = "WW-DCN-Solutions"`
`MgmtNetworkName`	The name of the vCenter network to attach to the VM's Management interface. This network must already exist in VMware or the installation will fail.
`DataNetworkName`	The name of the vCenter network to attach to the VM's Data interface. This network must already exist in VMware or the installation will fail.
`Host`	The ESXi host, or ONLY the vcenter VM/resource group name where the VM is to be deployed. The primary option is to use the host IP or name (all the hosts should be under the data center). If the hosts are under a VM in the data center, only provide the VM name (all hosts within the VM will be picked up). The subsequent option is to use a resource group. In this case, a full path should be provided. Example: `Host = "Main infrastructure/Resources/00_trial"`
`Datastore`	The datastore name available to be used by this host or resource group. The primary option is to use host IP or name. The subsequent option is to use a resource group. Example: `Datastore = "SDRS-DCNSOL-prodexsi/bru-netapp-01_FC_Prodesx_ds_15"`
`HSDatastore`	The high speed datastore available for this host or resource group. When not using a highspeed data store, set to same value as Datastore.
`Cw_VM_Image`	The name of Crosswork VM image in vCenter. This value is set as an option when running the installer tool and does not need to be set in the template file.
`HostedCwVMs`	The ID of the VM to be hosted by the ESXi host or resource.

Install Crosswork Network Controller using the vCenter vSphere UI

Describes how to deploy Crosswork Network Controller on a single VM through the vCenter vSphere UI.

Use this task to install Crosswork Network Controller using the vCenter vSphere UI.

Before you begin

Import the required CA certificate chain into the VMware trust store before you deploy the signed OVA image from the VMware vSphere UI. For more information, refer to Import the CA certificate chain on VMware on page 102.

This topic explains how to deploy Crosswork Network Controller on a single VM using the vCenter user interface.

This is the recommended method for installing Crosswork Network Controller on a single VM.

Procedure

1. Prepare the OVA and vSphere host.
   a) Download the latest available Cisco Crosswork platform image file (*.ova) to your system.
   b) With VMware ESXi running, log into the VMware vSphere Web Client. On the left navigation pane, choose the ESXi host where you want to deploy the VM.
   c) In the vSphere UI, go to **Host** > **Configure** > **Networking** > **Virtual Switches** and select the virtual switch for the Management Network that will be used to access the UI of the VM. In the virtual switch, select **Edit** > **Security**, and configure the following DVS port group properties:
      - Set **Promiscuous mode** as *Reject*
      - Set **MAC address changes** as *Reject*

      Confirm the settings and repeat the process for the virtual switch that will be used for the Data Network.
   d) Review and confirm that your network settings meet the requirements.

      Ensure that the networks that you plan to use for Management Network and Data network are connected to the host. Contact your Cisco Experience team for assistance.

2. Deploy the OVA.
   a) Choose **Actions** > **Deploy OVF Template**.

      Caution

      The default VMware vCenter deployment timeout is 15 minutes. If vCenter times out during deployment, the resulting VM will not be bootable. To prevent this, we recommend that you document the choices (such as IP address, gateway, DNS server, etc.) so that you can enter the information quickly and avoid any issues with the VMware configuration.

   b) The VMware **Deploy OVF Template** window appears, with the first step, **1 - Select an OVF template**, highlighted. Click **Choose Files** to navigate to the location where you downloaded the OVA image file and select it. Once selected, the file name is displayed in the window.
   c) Click **Next**. The **Deploy OVF Template** window is refreshed, with **2 - Select a name and folder** now highlighted. Enter a name and select the respective data center for the Cisco Crosswork VM you are creating.

      We recommend that you include the Cisco Crosswork version and build number in the name, for example: Crosswork-7-2-1-SP.

d) Click **Next**. The **Deploy OVF Template** window is refreshed, with **3 – Select a compute resource** highlighted. Select the host for your Cisco Crosswork VM.

e) Click **Next**. The VMware vCenter Server validates the OVA. Network speed will determine how long validation takes. After the validation is complete, the **Deploy OVF Template** window is refreshed, with **4 – Review details** highlighted.

f) Review the OVF template that you are deploying. This information is gathered from the OVF, and cannot be modified.

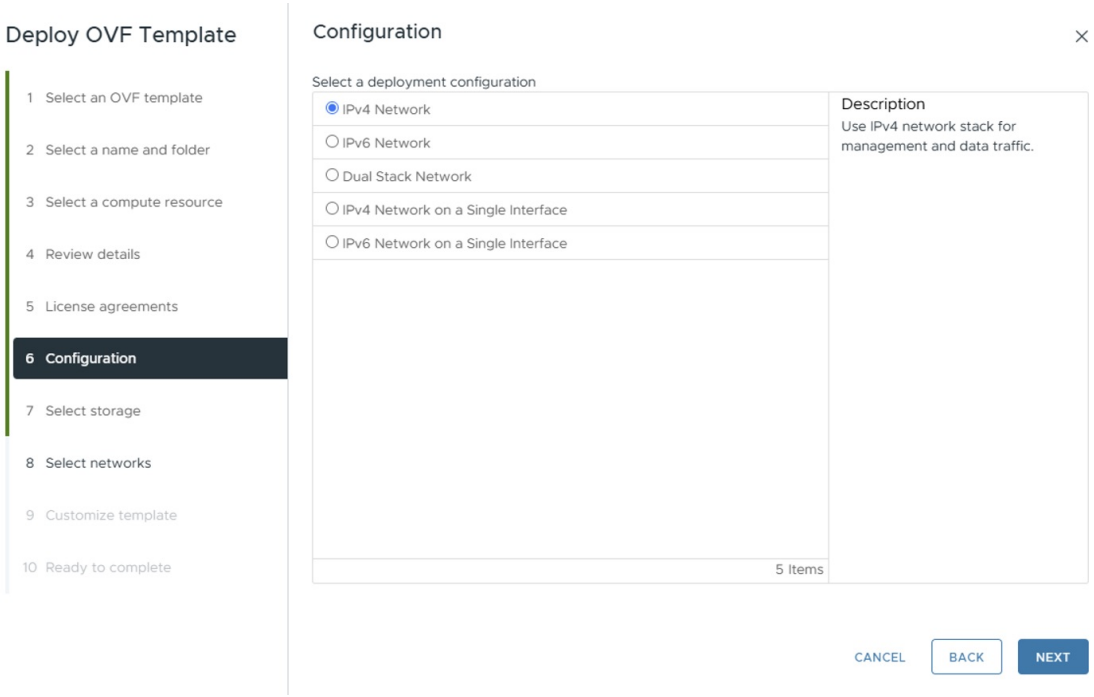
> **Note**
>
> You may refer to alerts regarding the OVF package containing advanced configuration options and/or about trusted certificates. These are common and you can safely select the "Ignore" option.

g) Click **Next**. The **Deploy OVF Template** window is refreshed, with **5 – License agreements** highlighted. Review the End User License Agreement and if you agree, click the **I accept all license agreements** checkbox. Otherwise, contact your Cisco Experience team for assistance.

3. Configure the deployment.

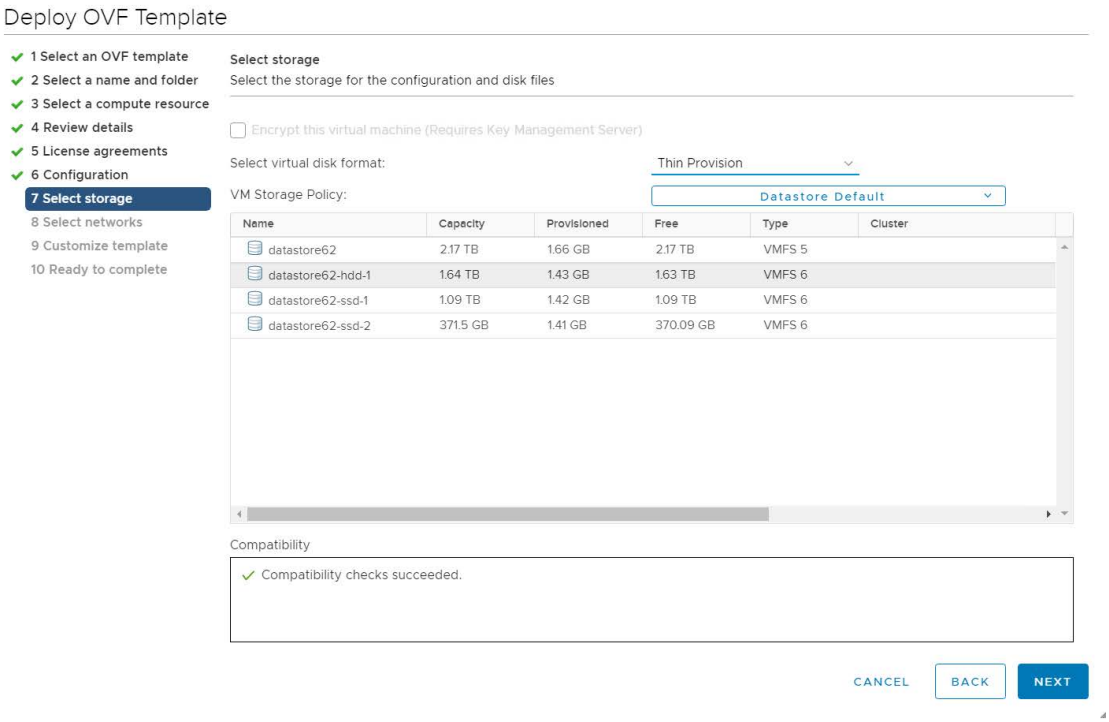
a) Click **Next** The **Deploy OVF Template** window is refreshed, with **6 – Configuration** highlighted. Choose the desired deployment configuration.

Figure 47: Select a deployment configuration



b) Click **Next**. The **Deploy OVF Template** window is refreshed, with **7 – Select Storage** highlighted. Choose the relevant option from the **Select virtual disk format** drop-down list. From the table, choose the datastore you want to use, and review its properties to ensure there is enough available storage.

Figure 48: Select storage

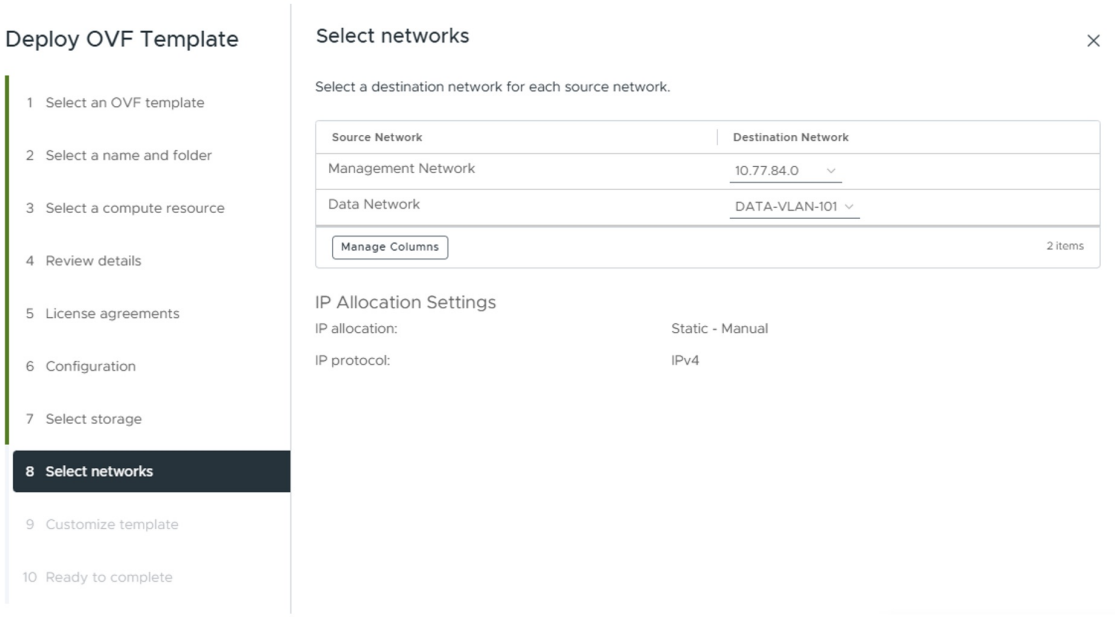


Note

For production deployment, choose the **Thick Provision Eager Zeroed** option because this will preallocate disk space and provide the best performance. For lab purposes, we recommend the **Thin Provision** option because it saves disk space.

c) Click **Next**. The **Deploy OVF Template** window is refreshed, with **8 - Select networks** highlighted. From the **Destination Network** drop-down list, select the proper networks for the Management Network and the Data Network.

Figure 49: Select networks



4. Click **Next**. The **Deploy OVF Template** window is refreshed, with **9 – Customize template** highlighted.
   a) Expand the **Management Network** settings. Provide information for the IPv4 and/or IPv6 deployment (as per your selection) such as IP address, IP netmask, IP gateway, virtual IP address, and virtual IP DNS name.
   b) Expand the **Data Network** settings. Provide information for the IPv4 and/or IPv6 deployment (as per your selection) such as IP address, IP netmask, IP gateway, virtual IP address, and virtual IP DNS name.
   c) Expand the **Deployment Credentials** settings. Enter relevant values for the VM Username and Password.

> **Note**
>
> Avoid using passwords that resemble dictionary words (for example, 'Pa55w0rd!') or easily guessable patterns. While such passwords might meet the initial criteria, they are considered weak and could cause the VM setup to fail without a clear explanation. To ensure a successful installation, use a complex password with a minimum of 8 characters that combines uppercase and lowercase letters, numbers, and special characters in a non-predictable sequence.

   d) Expand the **DNS and NTP Servers** settings. According to your deployment configuration (IPv4 and/or IPv6), the fields that are displayed are different. Provide information in the following three fields:
      - **DNS IP Address**: The IP addresses of the DNS servers you want the Cisco Crosswork server to use. Separate multiple IP addresses with spaces.
      - **NTP Servers**: The IP addresses or host names of the NTP servers you want to use. Separate multiple IPs or host names with spaces.
      - **DNS Search Domain**: The name of the DNS search domain.
      - **Timezone**: Enter the timezone details. Default value is UTC.

> **Note**
>
> The DNS and NTP servers must be reachable using the network interfaces you have mapped on the host. Otherwise, the configuration of the VM will fail.

e) Expand the **Disk Configuration** settings. Provide relevant values for these fields:

- **Logfs Disk Size**
- **Datafs Disk Size**
- **Corefs Partition Size**
- **High Speed Disk Size**
- **Minium backup partition size**

The default disk configuration settings should work for most environments. Change the settings only if you are instructed to by the Cisco Customer Experience team.

f) Expand **Crosswork Configuration** and enter your legal disclaimer text (users will refer to this text if they log into the CLI).

g) Expand **Crosswork Cluster Configuration**. Provide relevant values for these fields:

Table 86: Crosswork cluster configuration parameters

Parameter	Value or instruction
VM Type	Choose **Hybrid**.
Cluster Seed node	Choose True .
Crosswork Management Cluster Virtual IP	Enter virtual IP of the management network.
Crosswork Management Cluster Virtual IP Name	Enter DNS hostname of virtual IP interface of the management network.
Crosswork Data Cluster Virtual IP	Enter virtual IP of the data network.
Crosswork Data Cluster Virtual IP Name	Enter DNS hostname of virtual IP interface of the data network.
Initial hybrid node count	Set to 1.
Initial total node count	Set to 1.
Location of VM	Enter the location of VM.
Disclaimer	Enter your legal disclaimer text (users will refer to this text if they log into the CLI).
Installation type	Not applicable to single VM deployment. Do not select any checkbox.
Enable Skip Auto Install Feature	Set to **True**.
Auto Action Manifest Definition	Use the default value (Empty).

Parameter	Value or instruction
Product specific definition	Enter the product specific definition. While deploying an arbiter VM, set the value as `<![CDATA[{"product_image_id": "CNC", "attributes": {"is_arbiter": "true"}}]]>`.
Ignore Diagnostic Failure?	Use the default value (**False**).

h) Expand **Crosswork Cluster Configuration**. Provide relevant values for the following fields:

5. Click **Next**. The **Deploy OVF Template** window is refreshed, with **10 – Ready to Complete** highlighted.
6. Review your settings and then click **Finish** if you are ready to begin deployment. Wait for the deployment to finish before continuing. To check the deployment status:
   a) Open a VMware vCenter client.
   b) In the **Recent Tasks** tab of the host VM, view the status of the **Deploy OVF template** and **Import OVF package** jobs.
7. Once the deployment is completed, right-click on the VM and select **Edit Settings**. The **Edit Settings** dialog box is displayed. Under the **Virtual Hardware** tab, update these attributes:

   Table 87: VM attributes

VM profile	CPU	Memory
`XLarge`	24	128 GB
`Large`	12	96 GB
`Small`	8	48 GB

> **Note**
>
> Used only for deploying an arbiter VM.

   For more information, refer to Resource footprint for single VM deployments on page 28.

   Click **OK** to save the changes.

8. Power on the Crosswork VM. To power on, expand the host's entry, click the Cisco Crosswork VM, and then choose **Actions** > **Power** > **Power On**.

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

The OVF deployment is complete.

What to do next

After the VM boots, log in to the VM and complete the terminal-based UI review and confirmation flow. Review the detected system profile and configuration values, correct supported values if needed, validate the configuration, and confirm. Peer Node Management is hidden for single-VM deployments.

Install Crosswork Network Controller via the OVF Tool

Describes how to deploy Crosswork Network Controller on a single VM with the OVF Tool.

Use this task to install Crosswork Network Controller via the OVF Tool.

Before you begin

- In your vCenter data center, go to **Host** > **Configure** > **Networking** > **Virtual Switches** and select the virtual switch. In the virtual switch, select **Edit** > **Security**, and ensure that these DVS port group properties are as shown:
  - Set **Promiscuous mode** as Reject
  - Set **MAC address changes** as Reject

  Confirm the settings and repeat the process for each virtual switch used by Crosswork.
- Ensure you are using the OVF tool version 4.4 or higher.

This topic explains how to deploy Crosswork Network Controller on a single VM using the OVF tool. You must modify the list of mandatory and optional parameters in the script as per your requirements and run the OVF tool.

Follow these steps to log in to the Cisco Crosswork VM from SSH:

Procedure

1. On the machine where you have the OVF tool installed, use these command to confirm that you have OVF tool version 4.4:

   ```
   ovftool --version
   ```

2. Create the script file (refer to example in this step) and provide relevant information as per your target environment (such as IP addresses, gateway, netmask, password, and VCENTER_PATH, etc.).

Note

The file names mentioned in this topic are sample names and may differ from the actual file names on cisco.com.

Note

This is a sample script for deploying an `XLarge` VM profile. If you need to deploy a `Large` to `Small` VM profile, please replace the `XLarge` values with corresponding values for your desired profile.

- **`XLarge`**

```
--numberOfCpus:"*"=24 --viCpuResource=:50000: \
--memorySize:"*"=131072 --viMemoryResource=:131072: \
```

- **`Large`**

```
--numberOfCpus:"*"=12  --memorySize:"*"=98304 \
--viCpuResource=-1:18000:-1 --viMemoryResource=-1:98304:-1 \
```

- **`Small`**

  Note

  Only used to deploy an arbiter VM.

```
--numberOfCpus:"*"=8  --memorySize:"*"=49152 \

ProductDefinition="<![CDATA[{"product_image_id": "CNC",
"attributes": {"is_arbiter": "true"}}]]>"]
EnableSkipAutoInstallFeature="True"
ddatafs=100
```

```
cat svm_install.sh
#!/usr/bin/env bash
Host="X.X.X.X"
DM="thick"
DS="DS36"
Deployment="cw_ipv4"
DNSv4="10.10.0.99"
NTP="ntp.cisco.com"
Timezone="US/Pacific"
EnforcePodReservations="True"
EnableSkipAutoInstallFeature="True"
```

```
Domain="cisco.com"
Disclaimer="ACCESS IS MONITORED"
VM_NAME="svmEMS"
DataNetwork="DataNet"
ManagementNetwork="MgmtNet"
DataIPv4Address="x.x.x.x"
DataIPv4Gateway="x.x.x.x"
DataIPv4Netmask="x.x.x.x"
ManagementIPv4Address="x.x.x.x"
ManagementIPv4Gateway="x.x.x.x"
ManagementIPv4Netmask="x.x.x.x"
K8sServiceNetworkV4="10.75.0.0"
K8sPodNetworkV4="10.225.0.0"
Password="CLI Password"
Username="cw-admin"
ManagementVIP="x.x.x.x"
DataVIP="x.x.x.x"
VMType="Hybrid"
IsSeed="True"
InitNodeCount="1"
InitMasterCount="1"

SVM_OVA_PATH=$1

VCENTER_LOGIN="Administrator%40vsphere%2Elocal:Password%981123%21@x.x.x.x"
VCENTER_PATH="DC1/host"

ovftool --version
ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv -ds=$DS \
--numberOfCpus:"*"=24 --viCpuResource=:50000: \
--memorySize:"*"=131072 --viMemoryResource=:131072: \
--diskMode=$DM --overwrite --powerOffTarget --powerOn --noSSLVerify \
--allowExtraConfig \
--deploymentOption=$Deployment \
--prop:"DNSv4=${DNSv4}" \
--prop:"NTP=${NTP}" \
--prop:"Timezone=${Timezone}" \
--prop:"EnforcePodReservations=${EnforcePodReservations}" \
--prop:"EnableSkipAutoInstallFeature=${EnableSkipAutoInstallFeature}" \
--prop:"Domain=${Domain}" \
--prop:"Disclaimer=${Disclaimer}" \
--name=$VM_NAME \
--net:"Data Network=${DataNetwork}" \
--net:"Management Network=${ManagementNetwork}" \
--prop:"DataIPv4Address=${DataIPv4Address}" \
--prop:"DataIPv4Gateway=${DataIPv4Gateway}" \
--prop:"DataIPv4Netmask=${DataIPv4Netmask}" \
--prop:"ManagementIPv4Address=${ManagementIPv4Address}" \
--prop:"ManagementIPv4Gateway=${ManagementIPv4Gateway}" \
--prop:"ManagementIPv4Netmask=${ManagementIPv4Netmask}" \
--prop:"K8sServiceNetworkV4=${K8sServiceNetworkV4}" \
--prop:"K8sPodNetworkV4=${K8sPodNetworkV4}" \
--prop:"CWPassword=${Password}" \
--prop:"CWUsername=${Username}" \
--prop:"ManagementVIP=${ManagementVIP}" \
--prop:"DataVIP=${DataVIP}" \
--prop:"VMType=${VMType}" \
--prop:"IsSeed=${IsSeed}" \
--prop:"InitNodeCount=${InitNodeCount}" \
--prop:"InitMasterCount=${InitMasterCount}" \
$SVM_OVA_PATH \
vi://$VCENTER_LOGIN/$VCENTER_PATH/$Host
```

3. Download the OVA and install scripts from cisco.com. For the purpose of these instructions, we use the file name as **cnc-advantage-single-node-deployment-7.2.1.ova**.

   Use these command to extract the files from the tar bundle:

   ```
   tar -xvzf cnc-advantage-single-node-deployment-7.2.1.ova
   ```

   The OVA is extracted:

   ```
   svm]# ls -al
   -rw-r--r--   1 root root 15416145920 Mar 28 11:12
   cw-na-cnc-advantage-svm-7.2.1.ova
   -rwxr-xr-x   1 root root        2324 Apr  2 14:06 svm_install.sh
   ```

4. Use these command to make the scripts executable:

   ```
   chmod +x {filename}
   ```

   For example:

   ```
   chmod +x svm_install.sh
   ```

5. Execute the script with the OVA file name as parameter:

   ```
   svm]# ./svm_install.sh cw-na-cnc-advantage-svm-7.2.1.ova
   VMware ovftool 4.4.0 (build-16360108)
   Opening OVA source: cw-na-cnc-advantage-svm-7.2.1.ova
   <Removed some output >
   Completed successfully
   ```

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

Install Crosswork Network Controller using the Docker installer tool

Describes the single-VM installation procedure using the Docker installer tool.

Use this task to install Crosswork Network Controller using the Docker installer tool.

Before you begin

- Make sure that your environment meets all the vCenter requirements specified in Installation requirements on page 332.
- The edited template in the `/data` directory contains sensitive information (VM passwords and the vCenter password). The operator needs to manage access to this content. Store the templates used for your install in a secure environment or edit them to remove the passwords.
- The `install.log`, `install_tf.log`, and `.tfstate` files will be created during the install and stored in the `/data` directory. If you encounter any trouble with the installation, provide these files to the Cisco Customer Experience team when opening a case.
- The install script is safe to run multiple times. Upon error, input parameters can be corrected and re-run. You must remove the `install.log`, `install_tf.log`, and `tfstate` files before each re-run. Running the installer tool multiple times may result in the deletion and re-creation of VMs.

- In case you are using the same installer tool for multiple Crosswork installations, it is important to run the tool from different local directories, allowing for the deployment state files to be independent. The simplest way for doing so is to create a local directory for each deployment on the host machine and map each one to the container accordingly.
- Docker version 19 or higher is required while using the installer tool. For more information on Docker, refer to https://docs.docker.com/get-docker/.
- In order to change install parameters or to correct parameters following installation errors, it is important to distinguish whether the installation has managed to deploy the VM or not. Deployed VM is evidenced by the output of the installer similar to:

  ```
  vsphere_virtual_machine.crosswork-IPv4-vm["1"]: Creation complete after 2m50s
  [id=4214a520-c53f-f29c-80b3-25916e6c297f]
  ```
- If you do not have Python installed, go to python.org and download the version of Python that is appropriate for your work station.

Known limitations:

- The vCenter host VMs defined must use the same network name (vSwitch) across all hosts in the data center.
- The vCenter storage folders or datastores organized under a virtual folder structure, are not currently supported. Ensure that the datastores referenced are not grouped under a folder.

This section explains the procedure to install Crosswork Network Controller on a single VM using the docker installer tool. This method is the least recommended compared to using the vCenter UI or the OVF tool for installation.

Procedure

1. Prepare the installer and configuration.
   a) In your Docker-capable machine, create a directory where you will store everything you will use during this installation.

   > **Note**
   >
   > If you are using a Mac, ensure that the directory name is in lower case.

   b) Download the installer bundle (.tar.gz file) and the OVA file from cisco.com to the directory you created previously. For the purpose of these instructions, we will use the file name as **cnc-premier-single-node-docker-deployment-7.2.1.tar.gz** and **cnc-premier-single-node-deployment-7.2.1.ova**.

   > **Note**
   >
   > The file names mentioned in this topic are sample names and may differ from the actual file names in cisco.com.

   c) Use these command to extract the installer bundle:

   ```
   tar -xvf cnc-premier-single-node-docker-deployment-7.2.1.tar.gz
   ```

   The contents of the installer bundle are unzipped to a new directory (e.g. `cnc-premier-single-node-docker-deployment-7.2.1`). The extracted files will contain the installer

image (**cnc-premier-single-node-docker-deployment-7.2.1-123.tar.gz**) and files necessary to validate the image.

d) Review the contents of the README file to understand everything that is in the package and how it will be validated in these steps.

e) Use these command to verify the signature of the installer image:

> **Note**
>
> Use `python --version` to find out the version of python on your machine.

If you are using Python 2.x, use these command to validate the file:

```
python cisco_x509_verify_release.py -e <.cer file> -i <.tar.gz file> -s
<.tar.gz.signature file>
-v dgst -sha512
```

If you are using Python 3.x, use these command to validate the file:

```
python3 cisco_x509_verify_release.py3 -e <.cer file> -i <.tar.gz file> -s
 <.tar.gz.signature file>
-v dgst -sha512
```

f) Use these command to load the installer image file into your Docker environment.

```
docker load -i <.tar.gz file>
```

For example:

```
docker load -i cnc-premier-single-node-docker-deployment-7.2.1-123.tar.gz
```

g) Run the Docker image list or Docker images command to get the "image ID" (which is needed in the next step).

For example:

```
docker images
```

The result will be similar to these: (section we will need is underlined for clarity)

```
My Machine% docker images
REPOSITORY                          TAG
             IMAGE ID             CREATED        SIZE
dockerhub.cisco.com/cw-installer  cnc-premier-svm-7.2.1
          a4570324fad30    7 days ago     276MB
```

> **Note**
>
> Pay attention to the "CREATED" time stamp in the table presented when you run `docker images`, as you might have other images present from the installation of prior releases. If you wish to remove these, the `docker image rm {image id}` command can be used.

2. Run the installer and complete deployment.
   a) Launch the Docker container using these command:

```
docker run --rm -it -v `pwd`:/data {image id of the installer container}
```

   To run the image loaded in our example, use these command:

```
docker run --rm -it -v `pwd`:/data a4570324fad30
```

Note

- You do not have to enter that full value. In this case, "docker run --rm -it -v `pwd`:/data a45" was adequate. Docker requires enough of the image ID to uniquely identify the image you want to use for the installation.
- In the above command, we are using the backtick (`). Do not use the single quote or apostrophe (') as the meaning to the shell is very different. By using the backtick (recommended), the template file, and OVA will be stored in the directory where you are on your local disk when you run the commands, instead of inside the container.
- When deploying a IPv6 setup, the installer needs to run on an IPv6 enabled container/VM. This requires additionally configuring the Docker daemon before running the installer, using these method:
  - **Linux hosts (ONLY)**: Run the Docker container in host networking mode by adding the "-network host" flag to the Docker run command line.

    ```
    docker run --network host <remainder of docker run
     options>
    ```

- Centos/RHEL hosts, by default, enforce a strict SELinux policy which does not allow the installer container to read from or write to the mounted data volume. On such hosts, run the Docker volume command with the Z option as shown below:

  ```
  docker run --rm -it -v `pwd`:/data:Z <remainder of docker
   options>
  ```

Note

The Docker command provided will use the current directory to read the template and the ova files, and to write the log files used during the install. If you encounter either of these errors you should move the files to a directory where the path is in lowercase (all lowercase, no spaces or other special characters). Then navigate to that directory and rerun the installer.

Error 1:

```
% docker run --rm -it -v `pwd`:/data a45
docker: invalid reference format: repository name must be lowercase.
See 'docker run --help'
```

Error 2:

```
docker: Error response from daemon: Mounts denied: approving
/Users/Desktop: file does not exist
ERRO[0000] error waiting for container: context canceled
```

b) Navigate to the directory with the VMware template.

```
cd /opt/installer
```

c) Copy the template file found under `/opt/installer/svm_deployment_template_tfvars` to the `/data` folder using a different name.

   For example: `cp svm_deployment_template_tfvars /data/deployment.tfvars`

   For the rest of this procedure, we will use `deployment.tfvars` in all the examples.

d) Edit the template file located in the `/data` directory in a text editor, to match your planned deployment (for reference, refer to Crosswork Network Controller VM templates for single VM deployments on page 546).

e) From the `/opt/installer` directory, run the installer.

```
./cw-installer.sh install -m /data/<template file name> -o /data/<.ova
file>
```

   For example:

```
./cw-installer.sh install -m /data/deployment.tfvars -o
/data/cnc-premier-single-node-deployment-7.2.1.ova
```

f) Read, and then enter "yes" if you accept the End User License Agreement (EULA). Otherwise, exit the installer and contact your Cisco representative.

g) Enter "yes" when prompted to confirm the operation.

Note

It is not uncommon to refer to some warnings like these during the install:

```
Warning: Line 119: No space left for device '8' on parent controller
'3'.
Warning: Line 114: Unable to parse 'enableMPTSupport' for attribute 'key'
 on element 'Config'.
```

If the install process proceeds to a successful conclusion (refer to sample output below), these warnings can be ignored.

Sample output:

```
cw_vms = <sensitive>
INFO: Copying day 0 state inventory to CW
INFO: Waiting for deployment status server to startup on 10.90.147.66.
Elapsed time 0s, retrying in 30s
Crosswork deployment status available at
http://{VIP}:30602/d/NK1bwVxGk/crosswork-deployment-readiness?orgId=1&refresh=10s&theme=dark

Once deployment is complete login to Crosswork via:
https://{VIP}:30603/#/logincontroller
INFO: Cw Installer operation complete.
```

Note

If the installation fails, open a case with Cisco and provide the .log files that were created in the /data directory (and the local directory where you launched the installer Docker container), to Cisco for review. The two most common reasons for the install to fail are: (a) password that is not adequately complex, and (b) errors in the template file. If the installer fails for any errors in the template (for example, mistyped IP address), correct the error and rerun the install script.

Install Crosswork Network Controller VM using CLI

Summarizes the KVM installation workflow and its known limitations.

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

Table 88: Installation workflow

Step	Action
1. Ensure you have performed the preliminary checks.	Refer to Preliminary checks on page 358 for details.
2. Set up and validate the KVM environment.	Refer to Set up and validate KVM on RHEL on page 179.

Step	Action
3. Configure network bridges and SRIOV	Refer to Configure network bridges or SRIOV on page 359.
4. Install Crosswork Network Controller VM on KVM.	Refer to Install Crosswork Network Controller VM on KVM using CLI on page 363.

- 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

Lists virtualization and KVM module checks for the KVM environment.

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

Describes how to install KVM on RHEL and validate the configuration.

This topic describes how to install KVM on RHEL.

Before you begin

Confirm you have administrator (sudo) privileges.

Procedure

1. Ensure you have the latest packages (RHEL version 8.10 or 9.4) needed for KVM installation.
2. Install the required virtualization tools.
   a) 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.

   b) Install libvirt virtualization daemon, which is necessary for managing VMs.

   ```
   sudo dnf install -y libvirt
   ```

   c) Install virt-manager, a graphical interface for managing VMs.

   ```
   sudo dnf install virt-manager -y
   ```

   d) Install additional virtualization tools for managing VMs.

   ```
   sudo dnf install -y virt-top libguestfs-tools
   ```

3. Start and enable the libvirtd virtualization daemon.
   a) Start the libvirtd daemon.

```
sudo systemctl start libvirtd
```

   b) Enable the libvirtd daemon.

```
sudo systemctl enable libvirtd
```

   c) Verify that the daemon is running.

```
sudo systemctl status libvirtd
```

4. Add users to the required groups, such as libvirt and qemu.

   In these commands, replace *your_username* with the actual username.

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

5. 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
```

6. 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 these commands to enable it.

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

KVM is successfully installed on RHEL and validated.

Configure network bridges or SRIOV

Describes the network configuration choices for KVM deployment.

Crosswork requires a 10G interface for all the data layer communications to operate at a scale. You can choose any networking configuration that provides 10G throughput.

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

For detailed instructions, refer to these topics:

- Configure network bridges on page 360

- Configure SRIOV on page 362

Configure network bridges

Describes how to configure the management and data network bridges.

Use this task to configure network bridges.

A network bridge, such as Linux bridge and Open vSwitch (OVS), acts like a virtual network switch, allowing multiple network interfaces to communicate as if they are on the same physical network. For detailed information, refer to RHEL 8.x or RHEL 9.x documentation.

Complete these steps to configure network bridges.

Procedure

1. Create network bridge connection for management interface:
   a) 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
```

   b) 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> master intMgmt
```

Example

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

   c) Assign IP address, netmask, gateway, DNS, and domain to `intMgmt`.

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

Example

```
nmcli con modify intMgmt ipv4.addresses <hostmgmtIp/mask>
```

```
nmcli con modify intMgmt ipv4.gateway <mgmtgw>
```

```
nmcli con modify intMgmt ipv4.dns <dnsIp>
```

```
nmcli con modify intMgmt ipv4.dns-search 'cisco.com'
```

```
nmcli con modify intMgmt ipv4.method manual
```

```
nmcli con modify intMgmt ipv4.route-metric 50
```

d) Bring up the `intMgmt` network connection.

```
nmcli connection up intMgmt
```

Example

```
nmcli con up intMgmt
```

2. Repeat the above steps and create network bridge connection for data interface:
   a) Create a 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
```

   b) 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> master intData
```

Example

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

   c) Assign IP address and other details to `intData`.

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

Example

```
nmcli con modify intData ipv4.addresses <hostdataIp/mask>
```

```
nmcli con modify intData ipv4.method manual
```

```
nmcli con modify intData ipv4.gateway <datagw>
```

```
nmcli con modify intData ipv4.route-metric 90
```

   d) Bring up the `intData` network connection.

```
nmcli connection up intData
```

Example

```
nmcli con up intData
```

Both network bridges, intMgmt and intData, are configured and active, enabling communication across associated network interfaces as if connected to the same physical network.

Configure SRIOV

Describes how to configure SR-IOV for the single-VM deployment.

Use this task to configure SRIOV.

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

Complete these steps to configure SRIOV.

Procedure

---

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

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

2. Set the number of VFs for the network interfaces according to your requirement. In a single VM installation, you need a minimum of two network interfaces: one for management and one for data. By default, two VFs are configured for each interface. You can 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
   ```

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

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

4. If any of the interfaces are configured for VLAN, assign VLAN IDs to the interfaces.

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

5. Save the changes and reboot the system.
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 184).

Install Crosswork Network Controller VM on KVM using CLI

Describes how to install Crosswork Network Controller on KVM using CLI.

Use this task to install Crosswork Network Controller VM on KVM using CLI.

Before you begin

Ensure that

- the KVM host meets all the requirements specified in KVM host bare metal requirements on page 170.
- the checks specified in Preliminary checks on page 358 is completed.
- your KVM environment is set up (refer to Set up and validate KVM on RHEL on page 179), and
- the network bridge or SRIOV is configured (refer to Configure network bridges or SRIOV on page 359).

Follow these steps to install Crosswork on KVM.

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

Procedure

1. As a first step, prepare the config ISO file (`ovf-env.xml`) for the Crosswork Network Controller VM. For sample templates, refer to Crosswork Network Controller deployment templates for KVM on page 535.

> **Note**
>
> 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, refer to Installation parameters on page 335.

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

b) Generate the ISO file.

```
$ mkisofs -R -relaxed-filenames -joliet-long -iso-level 3 -l -o node-svm.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-svm" is the VM hostname.

2. Download the Crosswork Network Controller VM qcow2 tar file and extract it twice (the file is double-compressed).

```
tar -xvf cnc-advantage-single-node-deployment-7.2.1-qcow2.tar.gz
```

This command extracts the tar file (`cnc-advantage-single-node-deployment-7.2.1-123-qcow2.tar.gz`) and related checksum files.

```
tar -xvf cnc-advantage-single-node-deployment-7.2.1-qcow2.tar.gz
./
./cnc-advantage-single-node-deployment-7.2.1-123-qcow2.tar.gz
./cnc-advantage-single-node-deployment-7.2.1-123-qcow2.tar.gz.signature
./CW_RELEASE.cer
./cisco_x509_verify_release.py3
./cisco_x509_verify_readme_gen.py3
./README
```

Extract it again: `tar -tvf cnc-advantage-single-node-deployment-7.2.1-123-qcow2.tar.gz`

This command creates three qcow2 files:

```
cnc-advantage-single-node-deployment-7.2.1-123_dockerfs.qcow2
cnc-advantage-single-node-deployment-7.2.1-123_extrafs.qcow2
cnc-advantage-single-node-deployment-7.2.1-123_rootfs.qcow2
```

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

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

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

```
ls -1
cw_dockerfs.qcow2
cw_extrafs.qcow2
cw_rootfs.qcow2
disk3
disk4
disk6
```

4. Install the Crosswork Network Controller VM using network bridge or SRIOV.

   In this example, "node-svm" is the host name of the Cisco Crosswork Network Controller VM.

   - Using network bridges:

   ```
   virt-install --boot uefi --boot hd,cdrom --connect qemu:///system
   --virt-type kvm --name node-svm --ram 98304 --vcpus 12 --os-type linux
   --disk
   path=cnc-advantage-single-node-deployment-7.2.1-123_rootfs.qcow2,format=qcow2,bus=scsi
    --disk
   path=cnc-advantage-single-node-deployment-7.2.1-123_dockerfs.qcow2,format=qcow2,bus=scsi
    --disk path=disk3,format=qcow2,bus=scsi --disk
   path=disk4,format=qcow2,bus=scsi --disk
   path=cnc-advantage-single-node-deployment-7.2.1-123_extrafs.qcow2,format=qcow2,bus=scsi
    --disk path=disk6,format=qcow2,bus=scsi
   --disk=node-svm.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-svm --ram 98304 --vcpus 12 --cpu
   host-passthrough --disk path=cw_rootfs.qcow2,format=qcow2,bus=scsi --disk
    path=cw_dockerfs.qcow2,format=qcow2,bus=scsi --disk
   path=disk3,format=qcow2,bus=scsi --disk path=disk4,format=qcow2,bus=scsi
   --disk path=cw_extrafs.qcow2,format=qcow2,bus=scsi --disk
   path=disk6,format=qcow2,bus=scsi --disk=node-svm.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 &
   ```

5. **Post-deployment**:
   a) After deployment, log in to Crosswork Network Controller, and ensure the VM is healthy.
   b) **Import inventory:** Download the inventory sample (.tfvars file) from the Crosswork UI and update it with information about the VM, along with the data center parameters. For sample .tfvars templates, refer to Crosswork Network Controller deployment templates for KVM on page 535. Then, import the file back into the Crosswork UI.

Post installation tasks (for Premier tier installation)

Describes how to deploy the NSO function pack after a Premier-tier single-VM installation.

Deploy the NSO function pack after installing Crosswork Network Controller Premier on a single VM.

Premier-tier installation on a single VM. After installing Crosswork Network Controller Premier on a single VM, deploy the NSO function pack from the Crosswork Network Controller UI.

Procedure

1. Open the **NSO Deployment Manager**.

   From the main menu, navigate to **Administration** > **Crosswork Manager** > **NSO Deployment Manager**

2. Select the **NSO Function Pack Bundles** tab.

   Select **NSO Function Pack Bundles** and then select **CWM SOLUTIONS FPS**.

3. Deploy the bundle.

   Select **Deploy**.

4. Confirm the deployment status.

   View the **View progress in job history** and confirm that the deployment job completed successfully.

The NSO function pack deployment job completes successfully.

Part VIII

Upgrade Cisco Crosswork Network Controller

Topics:

- Upgrade Cisco Crosswork Network Controller
- Security Patches for Crosswork Network Controller

14 Upgrade Cisco Crosswork Network Controller

Topics:

- Upgrade overview
- Upgrade requirements
- Upgrade using existing hardware
- Upgrade using parallel hardware
- Upgrade Crosswork Network Controller applications (standalone activity)
- Upgrade Crosswork Network Controller on a single VM

Introduces the supported upgrade paths for Cisco Crosswork Network Controller. Summarizes preparation, cluster and single-VM migration, Data Gateway upgrade, and post-upgrade verification.

This chapter describes how to upgrade Cisco Crosswork Network Controller and its related components.

It covers upgrade preparation, existing- and parallel-hardware paths, single-VM migration, Crosswork Data Gateway, and post-upgrade checks.

Upgrade overview

Describes the available Crosswork Network Controller upgrade paths, supported releases, maintenance considerations, and related application upgrade information.

This section provides the high-level overview for upgrading Crosswork Network Controller to the latest version.

Upgrade methods and supporting information

Supported paths to release 7.2.1 sp

Important notes

This includes upgrading Crosswork Network Controller cluster, Crosswork Data Gateway and Crosswork Applications within a single maintenance window.

You can upgrade Crosswork Network Controller in these ways:

1. Upgrade using existing hardware on page 374
2. Upgrade using parallel hardware on page 387

The time taken for the entire upgrade window can vary based on size of your deployment profile and the performance characteristics of your hardware.

> **Note**
>
> While the cluster installation is in progress, you must upgrade NSO and SR-PCE. Refer to the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes* to know the NSO and SR-PCE versions compatible with Crosswork Network Controller. The process to upgrade NSO or SR-PCE is not covered in this document. For install instructions, please refer to the relevant product documentation.
>
> - Cisco NSO: Cisco NSO documentation
> - Cisco SR-PCE: Cisco IOS XRv 9000 Router Installation Guide.

Crosswork applications can be independently updated from the Crosswork Network Controller UI for minor updates or selected patch releases. This rolling upgrade procedure does not apply to Release 7.2.1 SP. For more information, see Upgrade Crosswork Network Controller applications (standalone activity) on page 396.

Crosswork Network Controller Release 7.2.1 Security Patch (SP) supports upgrades from these releases.

> **Note**
>
> Release 7.2.1 SP uses a fresh installation and backup-restore migration workflow; it is not a rolling upgrade.

- 7.1.3 → 7.2.1 SP
- 7.2 → 7.2.1 SP
- 7.2.1 → 7.2.1 SP

> **Note**
>
> If you are using Crosswork Network Controller 7.0 on AWS EC2 and want to upgrade to version 7.2.1 SP, the supported upgrade path is: **7.0 → 7.1 → 7.1.3 → 7.2.1 SP**. You must perform the upgrade in stages, but the overall workflow and procedures for each step remain the same as described in this chapter. For assistance, contact your Cisco account team.

Migration of Crosswork Network Controller from an earlier version has certain limitations:

- Changes introduced in the 7.2.x release may affect the compatibility and operation of custom packages in Crosswork Network Controller. If your deployment uses any custom packages, contact the Cisco Customer Experience team to have the packages validated and any necessary changes be completed.

  For 7.2.1 SP, there is a known change in behaviour for Device Performance Monitoring (DPM) custom device packages where they might not function as expected after migration from previous releases. Contact the Cisco Customer Experience team before migration to have the DPM custom device packages validated and any necessary changes completed.
- Ensure you enter the full FQDN as the hostname for AAA servers in the source version of Crosswork Network Controller.
- License tags are not auto-registered as part of the upgrade operation. You must register them manually after the upgrade.
- Third-party device configuration in Device Lifecycle Management and Cisco NSO is not migrated, and needs to be re-applied on the new Crosswork Network Controller version post migration.
- Custom user roles (Read-Write/Read) created in earlier version of Cisco Crosswork are not migrated, and need to be updated manually on the new version post migration.
- Any user roles with administrative privileges in the earlier version of Cisco Crosswork must be assigned new permissions after the upgrade to continue being administrative users.
- Health Insights KPI alert history is not retrieved as part of the migration.
- During an upgrade, the Health Insights alert checkbox is automatically disabled, and an alarm is raised for any KPI profile that contains a custom KPI.
- Before attempting to load a KPI, ensure that it complies with the requirements of the current release. If you try to load a KPI that was created for a previous release and is not compatible, you will receive an error message.
- After a successful migration, you must perform a hard refresh or browser cache deletion before proceeding to use the system. Failing to do this step can result in data discrepency.
- When a large-interface-density Cisco ASR 9000 series router with 1000 or more interfaces, classified as a Large Wireless Router (LWR) is migrated from Crosswork Network Controller version 7.1 to 7.2.1 SP:
  - The device does not immediately reach **Completed** state
  - GNMI credentials must be added
  - A detailed inventory sync must be performed

  After these steps, the device transitions to the **Completed** state.

Upgrade requirements

Summarizes the preparation activities required before upgrading Cisco Crosswork, including AWS EC2 configuration and network bandwidth management feature packs.

You must perform these tasks prior to upgrading Cisco Crosswork.

The preparation requirements are described in these topics:

- Save AWS EC2 configuration
- Disable network bandwidth management feature packs

Save AWS EC2 configuration

Lists the AWS EC2 configuration values to record for the previous Crosswork cluster and Crosswork Data Gateway deployments.

Record the configuration parameters used in the previous Crosswork cluster and Crosswork Data Gateway deployments. These parameters include:

- VM Node IP addresses
- Elastic Network Interface (ENI) ID mapped to the respective nodes (found under **Networking** tab **> ENI ID**)
- Subnet ID
- Security Group
- Availability Zone (AZ)

Disable network bandwidth management feature packs

Identifies the Local Congestion Mitigation, SR Circuit Style Manager, and Bandwidth on Demand feature packs that require preparation before an upgrade.

This topic identifies the feature packs that require preparation before an upgrade.

If you have enabled Local Congestion Mitigation, SR Circuit Style Manager, or Bandwidth on Demand in your current version of Crosswork and want to upgrade to the latest version, you must complete these tasks prior to upgrading.

Review the related feature-pack information and procedures:

- Disable Local Congestion Mitigation (LCM)
- Existing LCM policy behavior during and after an upgrade
- Disable Bandwidth on Demand (BWoD)
- Disable SR Circuit Style Manager (CSM)

Disable Local Congestion Mitigation (LCM)

Disable LCM and save its configuration options before upgrading Crosswork.

Disable LCM and save its configuration options before upgrading.

Follow these steps to disable LCM and save configuration options:

Procedure

1. Navigate to **Services & Traffic Engineering > Local Congestion Mitigation**.
2. From a LCM domain card, click ··· and select **Configuration**.
3. From the LCM **Configuration** page:
   a) Set the **Delete Tactical SR Policies when Disabled** option to **False**. This task must be done prior to disabling LCM so that tactical polices deployed by LCM remain in the network after the upgrade.
   b) Set the **Enable** option to **False**. If LCM remains enabled, there is a chance that tactical policies may be deleted after the upgrade.
   c) Record all options (Basic and Advanced) in the LCM **Configuration** page so that you can confirm the same configuration has been migrated after the upgrade.
4. Repeat steps 2 and 3 for each LCM domain.
5. Export the current list of interfaces managed by LCM. Confirm the interfaces are valid by reimporting the CSV file without errors. For more information, see "Add Individual Interfaces".

What to do next

After the upgrade, it is important to verify LCM data migration before enabling LCM.

Existing LCM policy behavior during and after an upgrade

Explains how existing LCM policies behave while LCM is disabled during an upgrade and after the upgrade is complete.

This topic describes the behavior of existing LCM policies during and after an upgrade.

Although LCM is in a disabled state during an upgrade, existing LCM policies remain active in the network. Congestion detection for existing or new interfaces will resume after the upgrade completes and LCM is enabled.

Disable Bandwidth on Demand (BWoD)

Disable BWoD and record its configuration options before upgrading Crosswork.

Disable BWoD and save its configuration options before upgrading.

Follow these steps to disable BWoD and save configuration options:

Procedure

1. Navigate to **Services & Traffic Engineering > Bandwidth on Demand**.
2. From the **Configuration** page:
   a) Set the **Enable** option to **False**. If BWoD remains enabled, there is a chance that tactical policies may be deleted after the upgrade
   b) Record all options (Basic and Advanced) in the **Configuration** page so that you can confirm the same configuration has been migrated after the upgrade.

What to do next

After the upgrade, wait until the **Traffic Engineering** page shows all the nodes and links before enabling BWoD.

Disable SR Circuit Style Manager (CSM)

Disable CSM and record its configuration options before upgrading Crosswork.

Disable CSM and save its configuration options before upgrading.

Follow these steps to disable CSM and save configuration options:

Procedure

1. Navigate to **Services & Traffic Engineering > Circuit Style SR-TE**.
2. From the **Configuration** page:
   a) Set the **Enable** option to **False**.
   b) Note all options (Basic and Advanced) in the CSM **Configuration** page so that you can confirm the same configuration has been migrated after the upgrade.

> **Note**
>
> Circuit Style SR-TE policies will go to an operation down (Oper Down) state if CSM is not enabled within 8 hours after disabling.

What to do next

After the upgrade, wait until the **Traffic Engineering** page shows all the nodes and links before enabling CSM.

Upgrade using existing hardware

Outlines the existing-hardware upgrade path and its ordered stages, from shutting down Data Gateway VMs through migration and post-upgrade verification.

This section explains how to migrate to the latest version of Crosswork Network Controller using the existing cluster.

Each stage in this upgrade workflow must be executed in sequence, and is explained in detail in later sections of this chapter. The stages are:

1. Shut down the Data Gateway VMs
2. Create backup and shut down
3. Install the latest version of the cluster
4. Install the applications

> **Note**
>
> You are recommended to download and validate the application CAPP files before starting the actual upgrade process. This will reduce your system downtime as opposed to downloading the CAPP files midway through the upgrade process.

5. Migrate the backup
6. Upgrade Crosswork Data Gateway

7. Post-upgrade checklist

After you complete the upgrade steps, you must install the NSO function packs. For the installation instructions, you can either refer to Install Cisco NSO function pack bundles from Crosswork UI or NSO function pack installation troubleshooting, depending on your preference.

Shut down the Data Gateway VMs

Record the Data Gateway VM details and shut down the VMs as the first stage of the existing-hardware upgrade.

Shut down the Crosswork Data Gateway VMs as the first stage of the existing-hardware upgrade.

This is the first stage of the upgrade workflow.

> **Note**
>
> When the data gateway VMs are shut down, data will not be forwarded to data destinations. Check with the application providers to determine if any steps are needed to avoid alarms or other problems.

Before you begin

Take screenshots of all the tabs in the **Data Gateway Management** page to keep a record of the list of data gateways, **Attached Device Count** in the Cisco Crosswork UI. In the **Pools** tab, for each pool listed here, take a screenshot to make a note of the active, spare, and unassigned VMs in the pool. This information is useful during Upgrade Crosswork Data Gateway.

Procedure

1. Check and confirm that all the VMs are healthy and running in your cluster.
2. Shut down the data gateway VMs.
   a) Log in to the data gateway VM. See Access Crosswork Data Gateway VM from SSH.

      Crosswork Data Gateway launches an Interactive Console after you log in successfully.

   b) Choose **5 Troubleshooting**.
   c) From the **Troubleshooting** menu, choose **5 Shutdown VM** to shut down the VM.
3. **For cloud-based deployment:**
   a. Terminate the Crosswork Data Gateway instances for all the previous Crosswork nodes. Do not delete the associated cluster stack.
   b. Manually delete the associated ENI IDs.

   > **Note**
   >
   > Since you are terminating the EC2 instances and not the stack, you must manually delete the associated ENI IDs.

Create backup and shut down

Create a backup of the current Cisco Crosswork system, verify the backup, and shut down the old cluster.

Create a backup of the current Cisco Crosswork system and shut down the old cluster.

This is the second stage of the upgrade workflow. Creating a backup is a prerequisite when upgrading your current version of Crosswork Network Controller to a new version.

> **Note**
>
> We recommend that you create a backup only during a scheduled upgrade window. Users should not attempt to access Crosswork Network Controller while the backup operation is running.

Before you begin

Follow these guidelines whenever you create a backup:

- Cisco Crosswork will back up the configuration of the system to an external server using SCP. Before you begin you need to have these configuration in place and information about the SCP server available:
  - The hostname or IP address and the port number of a secure SCP server.
  - A preconfigured path on the SCP server where the backup will be stored.
  - User credentials with file read and write permissions to the directory.
  - The SCP server storage requirements will vary slightly but you must have at least 25 GB of storage.
- Ensure that you have configured a destination SCP server to store the backup files. This configuration is a one-time activity.
- After the backup operation is completed, navigate to the destination SCP server directory and ensure that the backup file is created. You will require this backup file in the later stages of the upgrade process.
- Both the Crosswork Network Controller cluster and the SCP server must be in the same IP environment. For example: If Crosswork Network Controller is communicating over IPv6, so must the backup server.
- Keep a record of the list of Crosswork applications you have installed in the current version of Cisco Crosswork, as you can only install those applications after migrating to the new version of Cisco Crosswork.
- If you have onboarded a custom MIB package in the current version of Cisco Crosswork, download a copy of the package to your system. You will need to upload the package after you complete migrating to new version of Cisco Crosswork. See Post-upgrade checklist on page 386 for more information.
- If you have modified the current version of Cisco Crosswork to include third-party device types, you must download the third-party device configuration file, and re-apply it to the new version of Cisco Crosswork. The device configuration file is located on the cluster node (at `/mnt/cw_glusterfs/bricks/brick3/sys-oids.yaml`) and on the pod (at `/mnt/backup/sys-oids.yaml`).
- If Local Congestion Mitigation (LCM), SR Circuit Style Manager (CSM), and Bandwidth on Demand (BWoD) are enabled, you must disable them before proceeding. You must also, if available, export the current list of interfaces managed by LCM (**Service and Traffic Engineering > Local Congestion Mitigation > Domain Identifier <domain_id> > Interface Thresholds > Export**). Follow the steps documented in Upgrade requirements on page 372.

Procedure

1. Check and confirm that all the VMs are healthy and running in your cluster.
2. **Configure an SCP backup server:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel. Make the relevant entries in the fields provided.
   c) Click **Save** to confirm the backup server details.
3. **Create a backup:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Actions** > **Data Backup** to display the **Data Backup** drawer panel with the destination server details prefilled.
   c) Provide a relevant name for the backup in the **Job Name** field.
   d) If any of the VMs or applications are not in **Healthy** state, but you want to create the backup, check the **Force** check box.

   > **Note**
   >
   > The **Force** option must be used only after consultation with the Cisco Customer Experience team.

   e) Uncheck the **Backup NSO** checkbox if you do not want to include Cisco NSO data in the backup.

   > **Note**
   >
   > To use the **Backup NSO** option during backup, you must configure the SSH connectivity protocol in the NSO provider; otherwise, the backup will fail. Follow the instructions given in *Backup Cisco Crosswork with Cisco NSO* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* instead of the instructions here.

   f) Complete the remaining fields as needed.

   If you want to specify a different remote server upload destination: Edit the prefilled **Host Name**, **Port**, **Username**, **Password** and **Remote Path** fields to specify a different destination.

   g) (Optional) Click **Verify Backup Readiness** to verify that Crosswork Network Controller has enough free resources to complete the backup. If the verifications are successful, Crosswork Network Controller displays a warning about the time-consuming nature of the operation. Click **OK**.

   If the verification is unsuccessful, please contact the Cisco Customer Experience team for assistance.

   h) Click **Start Backup** to start the backup operation. Crosswork Network Controller creates the corresponding backup job set and adds it to the job list. The Job Details panel reports the status of each backup step as it is completed.
4. Complete the backup review and upload actions.
   a) To view the progress of a backup job: Enter the job details (such as Status or Job Type) in the search fields in the **Backup and Restore Job Sets** table. Then click on the job set you want.

   The **Job Details** panel displays information about the selected job set, such as the job Status, Job Type, and Start Time. If there's a failed job, hover the mouse pointer over the icon near the **Status** column to view the error details.

> **Note**
>
> After the backup operation is completed, navigate to the destination SCP server directory and ensure that the backup file is created. You will require this backup file in the later stages of the upgrade process.

> **Note**
>
> If you do not see your backup job in the list, refresh the **Backup and Restore Job Sets** table.

b) If the backup fails during upload to the remote server: In the **Job Details** panel, just under the Status icon, click the **Upload backup** button to retry the upload.

> **Note**
>
> Upload can fail due to connectivity problems with the SCP backup server (for example, incorrect credentials, missing directory or directory permissions, missing path and so on). This is indicated by failure of the task `uploadBackupToRemote`). If this happens, check the SCP server details, correct any mistakes and try again. Alternatively, you can use the **Destination** button to specify a different SCP server and path before clicking **Upload backup**.

5. After a successful backup, shut down the Crosswork Network Controller cluster by powering down the VMs hosting each node (start with the Hybrid VMs):
   - **For on-premise deployment:**
     a. Log into the VMware vSphere Web Client.
     b. In the **Navigator** pane, right-click the VM that you want to shut down.
     c. Choose **Power** > **Power Off**.
     d. Wait for the VM status to change to **Off**.
     e. Wait for 30 seconds and repeat the above steps (a to d) for each of the remaining VMs.
   - **For cloud-based deployment:**
     a. Terminate the EC2 instance for all the previous Crosswork nodes. Do not delete the associated cluster stack.
     b. Manually delete the associated ENI IDs.

> **Note**
>
> Since you are terminating the EC2 instances and not the stack, you will need to manually delete the associated ENI IDs.

6. Move Cisco NSO into read-only mode to avoid any unintended updates to Cisco NSO during the upgrade.
   Use these command to move NSO to read-only mode:

```
ncs_cmd -c maapi_read_only
```

Install the latest version of the cluster

Install the latest cluster version after confirming prerequisites and completing a successful backup of the old version.

Install the latest cluster version after the old-version backup succeeds.

After the successful backup of the old version of Crosswork Network Controller, proceed to install the latest version of the Crosswork Network Controller cluster.

> **Note**
>
> The number of VMs installed in the new version of Crosswork Network Controller must be equal or more than the number of VMs in the old version of Crosswork Network Controller.

Before you begin

- Make sure that your environment meets all the installation prerequisites (see Installation requirements overview on page 44).
- During installation, Cisco Crosswork will create a special administrative ID (**virtual machine (VM) administrator**, with the username *cw-admin*, and the default password *cw-admin*). The administrative username is reserved and cannot be changed. The first time you log in using this administrative ID, you will be prompted to change the password. Data center administrators use this ID to log into and troubleshoot the Crosswork application VM. You will use it to verify that the VM has been properly set up.

Procedure

1. **For on-premise deployments:** Install Crosswork Network Controller cluster using the instructions in this document.
   - VMware: Installation overview on page 82
   - KVM: Installation overview on page 170
2. **For cloud-based deployments:** Install Crosswork Network Controller cluster using the template for a single VM (the sample template file (.yaml) can be downloaded from cisco.com). Deploy the template sequentially for each VM in your cluster, ensuring a minimum interval of 10 minutes between each deployment.

> **Note**
>
> - Ensure that you update the parameters (such as VM Node IP addresses, ENI ID, Subnet ID, Security Group, AZ, etc.) in the sample template to match the values used in the previous cluster deployment.
> - The first VM deployed must be configured as the seed node (by setting `IsSeed` as `True`).
> - Each script run should have a unique stack name. Do not use the stack name of the previous cluster. Ensure that you use a unique stack name when deploying each VM.

After the installation is completed for all nodes, verify that the Load Balancer value (**Load balancer** > **Choose any one Listener** > **Registered Target groups**) matches the value in the previous cluster stack.

3. After the installation is completed, log into the Crosswork Network Controller UI. From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**. Click **Crosswork Cluster** tile to view the details of the cluster such as resource utilization by node, the IP addresses in use, whether each node is a Hybrid or Worker, and so on.

   Ensure all the nodes are up and running in the cluster. Check the health status to confirm everything is functioning correctly.

Install the applications

Install the same applications that were backed up on the new cluster and verify cluster health.

Install the backed-up applications on the new cluster and check its health.

After successfully installing the new version of the Crosswork Network Controller cluster, proceed to install the latest version of the Crosswork Network Controller applications.

> **Note**
>
> The Crosswork Network Controller applications that you install must be the same ones that were backed up during Create backup and shut down on page 376.

Procedure

1. Install the Crosswork Network Controller applications using the steps described in Install Crosswork Network Controller applications on page 266.
2. After the applications are successfully installed, check the health of the new Crosswork Network Controller cluster.
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**.
   b) Click **Crosswork Cluster** tile to view the health details of the cluster.

Migrate the backup

Restore the saved backup on the new cluster using the destination SCP server and backup file.

Restore the saved backup on the new cluster.

After successfully installing the new versions of the Crosswork Network Controller applications, proceed to migrate the Crosswork Network Controller backup taken earlier to the new Crosswork Network Controller cluster.

Before you begin

Before you begin, ensure that you have:

- The hostname or IP address and the port number of a secure destination SCP server used in Create backup and shut down on page 376.
- The name and path of the backup file created in Create backup and shut down on page 376.
- User credentials with file read and write permissions to the directory.

Procedure

1. Check and confirm that all the VMs are healthy and running in your cluster.
2. **Configure an SCP backup server:**
   a) From the main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel.
   c) Make the relevant entries in the fields provided.

   > **Note**
   >
   > In the **Remote Path** field, please provide the location of the backup created in Create backup and shut down on page 376.

   d) Click **Save** to confirm the backup server details.
3. Migrate the previous Crosswork Network Controller backup on the new Crosswork Network Controller cluster:
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Actions** > **Data Migration** to display the **Data Migration** drawer panel with the destination server details prefilled.
   c) Provide the name of the data migration backup (created in Create backup and shut down on page 376) in the **Backup File Name** field.
   d) If you want to perform the data migration backup despite any Crosswork Network Controller application or microservice issues, check the **Force** check box.
   e) Click **Start Migration** to start the data migration operation. Crosswork Network Controller creates the corresponding data migration job set and adds it to the **Backup and Restore Job Sets** table. The Job Details panel reports the status of each backup step as it is completed.

> **Note**
>
> If you do not see your job in the list, please wait for a few minutes and refresh the **Backup and Restore Job Sets** table.

f) To view the progress of a data migration job: Enter the job details (such as Status or Job Type) in the search fields in the **Backup and Restore Job Sets** table. Then click on the job set you want.

The **Job Details** panel displays information about the selected job set, such as the job Status, Job Type, and Start Time. If there's a failed job, hover the mouse pointer over the icon near the **Status** column to view the error details.

> **Note**
>
> Crosswork UI might become temporarily unavailable during the data migration operation. When the Crosswork UI is down, you can view the job status in the Grafana dashboard. The Grafana link is available as *View Data Migration Process Dashboard* option on the right side of the Job Details window.

g) If the data migration fails in between, you need to restart the procedure from step 1.

4. After the data migration is successfully completed, check the health of the new Crosswork Network Controller cluster.
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**.
   b) Click **Crosswork Cluster** tile to view the health details of the cluster.

Upgrade Crosswork Data Gateway

Upgrade the Crosswork Data Gateway Base VMs after cluster migration, then verify enrollment, pools, devices, and collection jobs.

Upgrade Crosswork Data Gateway Base VMs after completing cluster migration.

This is the final stage of the upgrade workflow. Ensure that the migration is complete and the new Cisco Crosswork UI is available before you proceed with installing the latest version of Crosswork Data Gateway.

> **Note**
>
> The upgrade procedure is required only for data gateway Base VM upgrade. Upgrade of other components, such as collectors, is performed by Cisco Crosswork.

Crosswork Data Gateway functions as a passive device in the network. The Data Gateway upgrade process consists of these steps replacing all the old Data Gateway VMs with the Data Gateway VMs in the network.

> **Note**
>
> Step 8 in this procedure requires you to log out of Cisco Crosswork and log in again after verifying the deployment and enrollment of the latest data gateway VMs with Cisco Crosswork. After you log in, an **Action to be taken** window appears prompting you to confirm that the upgrade is complete. Do not click **Acknowledge** unless you have completed all the verification steps mentioned in Step 3, Step 4, and Step 5 in the procedure.

Before you begin

The administrators must assign the Crosscluster API to the user roles. This permission is required for a user to access the Data Gateway Management page.

To assign the API permissions, from the Crosswork UI, go to **Administration > Users and Roles > Roles > Global API permissions > Crosscluster Infra**. For more information on roles, refer to the *User Roles, Functional Categories, and Permissions* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

Procedure

1. Log out of Cisco Crosswork and log in again.
2. After you log in, an **Action to be taken** window appears. Close this window and do not click **Acknowledge**.
3. Install new data gateway VMs with the same number and the same information (management interface importantly) as the old data gateway VMs.

   The installation procedure for the new data gateway VMs may vary depending on your data center.

   - **For on-premise deployment:**

     Follow the steps in the Crosswork Data Gateway on KVM on page 200.

   - **For cloud-based deployment:**

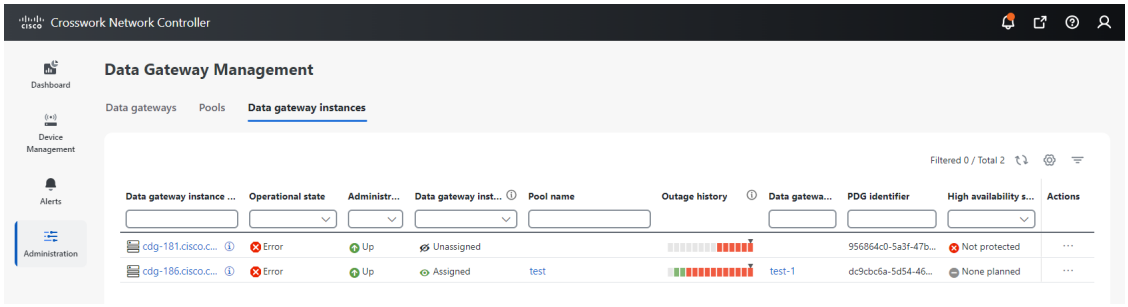
     a. Install the data gateway VM nodes using the sample Crosswork Data Gateway templates (.yaml) which you can download from cisco.com. Deploy the templates sequentially for each data gateway VM, ensuring a minimum interval of 10 minutes between each deployment.

     > **Note**
     >
     > - Ensure that you update the parameters (such as VM Node IP addresses, ENI ID, Subnet ID, Security Group, AZ, etc.) in the sample template to match the values used in the previous cluster deployment.
     > - Each script run should have a unique stack name. Do not use the stack name of the previous cluster. Ensure that you use a unique stack name when deploying each VM.

     b. After a few minutes, verify if the Data Gateway is operational.

4. Wait for about 5 minutes and navigate to **Administration** > **Data Gateway Management**.
5. Check the **Data Gateway Instances** tab to verify that the new data gateway VMs are enrolled with Cisco Crosswork and have the **Admin State** as **Up** and **Operational State** as **Not Ready**.

   Figure 50: Data gateway instances window

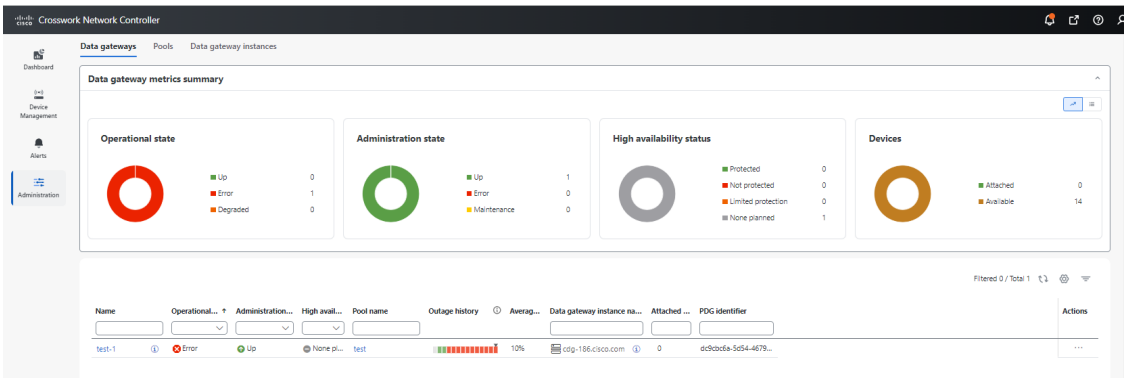


6. After the **Operational State** of the VMs changes to **Ready**, navigate to the **Pools** tab and verify that all the data gateway pools from the previous version of Cisco Crosswork, are listed here. Edit each data gateway pool to verify that the active data gateway is the same as one that you noted in the previous version of Cisco Crosswork. To view the pool details, click the pool name.

> **Note**
>
> If you are upgrading the Crosswork Data Gateway from a single stack deployment to a dual stack deployment, edit the pool and specify the VIP IP address of the additional stack. For example, if the pool was created with an IPv4 address, you must provide the IPv6 address. For information on editing a pool, refer to the *Edit or delete a Crosswork Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

7. Verify that devices are attached to the data gateways in the Cisco Crosswork UI.
   a) Navigate to the **Administration** > **Data Gateway Management** page.
   b) Check the **Attached Device Count** of the data gateway.

   Figure 51: Data gateway window



8. Log out of Cisco Crosswork.

After the upgrade is complete:

- The data gateway VMs are enrolled with Cisco Crosswork.
- All destinations, data gateway pools, device-mapping information can be viewed on the Cisco Crosswork UI with the upgraded data gateway VMs.
- Collection jobs start again automatically with the new data gateway VMs.
- After upgrading the data gateway VM, you must reconfigure the collector resources and the disabled containers. Global Parameter resources that were configured before the upgrade are not retained. To configure the resource

parameters, on the Crosswork UI, navigate to **Administration** > **Data Gateway Global Settings** > **Data Gateway** > **Resource**. For more information on the resources, refer to the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.

What to do next

1. After you log in to the Crosswork Network Controller UI, a window appears prompting for confirmation is displayed. In the pop-up that appears, click **Acknowledge**.

   > **Note**
   >
   > Do not click **Acknowledge** unless you have verified that the VMs are in the **Up**/**Not Ready** state. Doing so results in VMs having the state as **Error**. See Troubleshoot Crosswork Data Gateway Upgrade Issues.

2. To move Cisco NSO out of maintenance or read-only mode, use:

   ```
   ncs_cmd -c maapi_read_write
   ```

Troubleshoot Crosswork Data Gateway upgrade issues

Lists common Crosswork Data Gateway upgrade issues and source-provided actions for identifying and addressing them.

these table lists common problems that might be experienced when upgrading Crosswork Data Gateway, and provides approaches to identifying the source of the problem and solving it.

Issue	Recommended Action
Some of the data gateway VMs are in **Error** or **Degraded** state because you clicked **Acknowledge** before the VMs came to the **Up**/**Not Ready** state	1. Wait for the data gateway VMs to have the state as **Up** or **Not Ready** state. 2. Once the VMs have the state as **Up** or **Not Ready**, delete all data gateway pools and create them again.
Some of the data gateway VMs are in **Error** or **Degraded** state because you clicked **Acknowledge** before the VMs came to the **Up**/**Not Ready** state. The state of the VMs did not change to **Up**/ **Ready** and they are still in **Error**.	1. Delete all data gateway pools. 2. Check if the VMs now have the state as **Up** or **Not Ready**. 3. If the VMs are still in a state of **Error**, manually re-enroll the VMs with the new version of Cisco Crosswork. For more information, refer to *Re-enroll Crosswork Data Gateway* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
The data gateways VMs are stuck in the **Degraded** state with Image manager being in exited state. The list of components for the data gateway either do not show Image manager or show it in an exited state.	1. In the Cisco Crosswork UI, navigate to **Data Gateway Management** > **Virtual Machines**. 2. Click the data gateway that is degraded. 3. Click **Actions** and click **Reboot**.

Post-upgrade checklist

Complete the source-listed checks after the new cluster is healthy, including collection-job, Data Gateway, and AAA verification.

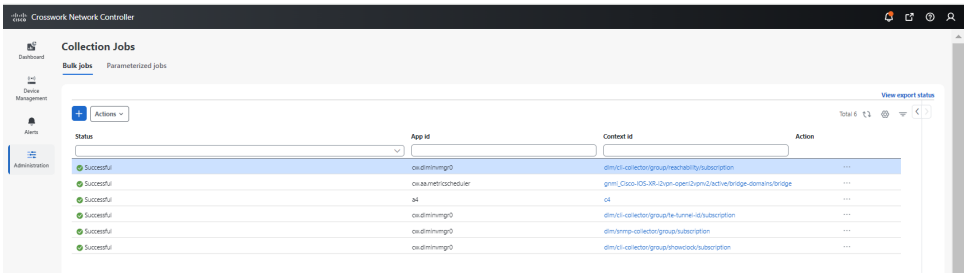
Complete the post-upgrade checks after the new cluster is healthy.

After you upgrade Cisco Crosswork to the latest version, check the health of the new cluster. If your cluster is healthy, perform these activities:

Procedure

1. Complete the initial post-upgrade checks.
   a) Perform a hard refresh or browser cache deletion before proceeding to use the system. Failing to do this step can result in data discrepancy.
   b) Navigate to **Administration** > **Collection Jobs** in Cisco Crosswork UI and delete the duplicate system jobs.

   Figure 52: Collection jobs window

   c) Verify that the collection jobs are running on the data gateway VMs in the **Administration** > **Collection Jobs** page.
   d) Verify the restored AAA data by logging in using default credentials, and configure custom user roles (Read-Write/Read) in the upgraded Cisco Crosswork.
   e) (Optional) Based on your network requirements, download the relevant map files from cisco.com and re-upload them to the upgraded Cisco Crosswork.
   f) (Optional) If any NSO device onboarding policy was set in the previous version of Cisco Crosswork, you must update the policy with new Network Element Drivers (NED) on the NSO.
   g) (Optional) Re-apply any third-party device configurations (used in the previous version of Cisco Crosswork) to the new version of Cisco Crosswork.
   h) (Optional) If you have saved views in your topology, you may need to re-apply certain settings (such as map configuration, table filters, etc.) as this data might not be fully migrated.
2. Complete the remaining post-upgrade checks.
   a) If you are using Crosswork Change Automation, verify that all the stock and custom playbooks are migrated successfully.
   b) If you are using Crosswork Health Insights, verify that the collection to the external destination is working. Also, check if the alert dashboard is displaying the correct data.
   c) For Traffic Engineering, perform these actions:
      - Upgrade the software versions in your devices as per the supported Cisco IOS XE/XR versions documented in the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes*.
      - If you have upgraded Local Congestion Mitigation (LCM), Verify LCM data migration on page 387 before enabling LCM.

- If you have upgraded SR Circuit Style Manager or Bandwidth on Demand, wait until the **Traffic Engineering** page shows all the nodes and links, and confirm configuration options match the ones you recorded prior to the upgrade before enabling feature packs.

What to do next

If you encounter errors in any of the above activities, please contact the Cisco Customer Experience team.

Verify LCM data migration

Verify that previously monitored LCM interfaces and policies migrated before enabling LCM domains.

Verify LCM data migration before enabling LCM domains.

Once the system is stable following an upgrade, and prior to enabling domains for LCM, confirm that the migration of previously monitored interfaces has completed and that each domain has the expected configuration options. Follow these steps to verify LCM data migration:

Procedure

---

1. After the upgrade, wait until the **Traffic Engineering** page shows all the nodes and links.
2. Navigate to **Alerts > Alarm and Events** and enter **LCM** to filter the **Source** column.
3. Look for these event: "Migration complete. All migrated LCM interfaces and policies are mapped to their IGP domains". If this message does not appear wait for the **Congestion Check Interval** period (set in the LCM **Configuration** page), then restart LCM (**Administration > Crosswork Manager > Optimization Engine > optima-lcm > ... > Restart**).
4. Wait until the optima-lcm service changes from Degraded to Healthy state.
5. For each domain, navigate to the **Configuration** page and verify the options have been migrated successfully. If the domain configurations are incorrect, restart LCM (**Administration > Crosswork Manager > Optimization Engine > optima-lcm > ... > Restart**).
6. Check the **Alarms and Events** page for the event mentioned above and the **Configuration** page to verify the options.

   Note

   - If the confirmation message does not appear or domain configuration options are incorrect, then contact Cisco Technical support and provide them with showtech information and the exported Link Management CSV file.
   - You can also manually add missing interfaces that were previously monitored or update domain configuration options *after* the system is stable.

Upgrade using parallel hardware

Summarizes the parallel-hardware upgrade path, its on-premises restriction, resource requirement, and ordered upgrade stages.

This section explains how to migrate to the latest version of Crosswork Network Controller using new hardware. This method relies on installing the new Crosswork Network Controller cluster on new hardware in parallel while the data from the old Cisco Crosswork cluster is being backed up. This method is faster but requires twice the amount of resources for creating the new cluster in parallel.

> **Note**
>
> Upgrade using parallel hardware is only supported for on-premise deployments. For cloud-based deployments, please follow the instructions in Upgrade using existing hardware.

The stages of the parallel upgrade workflow are:

1. Deploy a new cluster
2. Back up the cluster
3. Update DNS server and run migration
4. Add Crosswork Data Gateway to Cisco Crosswork
5. Shut down the old cluster

After you complete the upgrade steps, you must install the NSO function packs. For the installation instructions, you can either refer to Install Cisco NSO function pack bundles from Crosswork UI or NSO function pack installation troubleshooting, depending on your preference.

Deploy a new cluster

Install the latest cluster and applications on new VMs while the existing cluster remains available for migration.

Deploy the latest cluster and applications on a new set of VMs.

Install the latest version of Crosswork Network Controller cluster and applications on a new set of VMs in parallel.

> **Note**
>
> The new Crosswork Network Controller cluster must be installed with the same FQDN and same number of nodes as in the old version of Crosswork Network Controller.

Before you begin

- Make sure that your environment meets all the installation prerequisites (see Installation requirements overview on page 44)

Procedure

1. Install the new Crosswork Network Controller cluster (see Installation overview on page 82).

> **Note**
>
> During installation, Cisco Crosswork will create a special administrative ID (**virtual machine (VM) administrator**, with the username *cw-admin*, and the default password *cw-admin*). The administrative username is reserved and cannot be changed. The first time you log in using this administrative ID, you will be prompted to change the password. Data center administrators use this ID to log into and troubleshoot the Crosswork application VM. You will use it to verify that the VM has been properly set up.

2. After the installation is completed, log into the Crosswork Network Controller UI by navigating to https://<NEW_VIP>:30603.
3. Check if all the nodes are up and running in the cluster.
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**.
   b) Click **Crosswork Cluster** tile to view the details of the cluster such as resource utilization by node, the IP addresses in use, whether each node is a Hybrid or Worker, and so on.
4. Install the applications which were part of the old version of Crosswork Network Controller. For more information, see Install Crosswork Network Controller applications on page 266.
5. After the applications are successfully installed, check the health of the new Crosswork Network Controller cluster.

Back up the cluster

Back up the cluster configuration to an external SCP server before migrating to the new cluster.

Back up the cluster configuration to an external SCP server.

Before you begin

Follow these guidelines whenever you create a backup:

- Cisco Crosswork will back up the configuration of the system to an external server using SCP. Before you begin you need to have these configuration in place and information about the SCP server available:
  - The hostname or IP address and the port number of a secure SCP server.
  - A preconfigured path on the SCP server where the backup will be stored.
  - User credentials with file read and write permissions to the directory.
  - The SCP server storage requirements will vary slightly but you must have at least 25 GB of storage.
- Ensure that you have configured a destination SCP server to store the backup files. This configuration is a one-time activity.
- Both the Crosswork Network Controller cluster and the SCP server must be in the same IP environment. For example: If Crosswork Network Controller is communicating over IPv6, so must the backup server.
- Keep a record of the list of Crosswork applications you have installed in the current version of Cisco Crosswork, as you can only install those applications after migrating to the new version of Cisco Crosswork.
- If you have onboarded a custom MIB package in the previous version of Cisco Crosswork, download a copy of the package to your system. You will need to upload the package after you complete upgrading Cisco Crosswork. See Post-upgrade checklist on page 386 for more information.

- If you have modified the previous version of Cisco Crosswork to include third-party device types, you must download the third-party device configuration file, and re-apply it to the upgraded Cisco Crosswork. The device configuration file is located on the cluster node (at `/mnt/cw_glusterfs/bricks/brick3/sys-oids.yaml`) and on the pod (at `/mnt/backup/sys-oids.yaml`).
- For Traffic Engineering, perform these actions:
  - Upgrade the software versions in your devices as per the supported Cisco IOS XE/XR versions documented in the Traffic Engineering Compatibility Information.
  - If you have upgraded Local Congestion Mitigation (LCM), Verify LCM data migration on page 387 before enabling LCM.
  - If you have upgraded SR Circuit Style Manager or Bandwidth on Demand, wait until the **Traffic Engineering** page shows all the nodes and links, and confirm configuration options match the ones you recorded prior to the upgrade before enabling feature packs.

> **Note**
>
> We recommend that you create a backup only during a scheduled upgrade window. Users should not attempt to access Crosswork Network Controller while the backup operation is running.

Procedure

1. Launch the Cisco Crosswork UI by using a browser and navigating to https://<FQDN>:30603
2. Check and confirm that all the VMs are healthy and running in your cluster.
3. **Configure an SCP backup server:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel. Make the relevant entries in the fields provided.
   c) Click **Save** to confirm the backup server details.
4. **Create a backup:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Actions** > **Backup** to display the **Backup** drawer panel with the destination server details prefilled.
   c) Provide a relevant name for the backup in the **Job Name** field.
   d) If any of the VMs or applications are not in **Healthy** state, but you want to create the backup, check the **Force** check box.

   > **Note**
   >
   > The **Force** option must be used only after consultation with the Cisco Customer Experience team.

   e) Uncheck the **Backup NSO** check box if you do not want to include Cisco NSO data in the backup.

> **Note**
>
> To use the **Backup NSO** option during backup, you must configure the SSH connectivity protocol in the NSO provider; otherwise, the backup will fail. Refer to the *Backup Cisco Crosswork with Cisco NSO* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* instead of the instructions here.

f) Complete the remaining fields as needed.

If you want to specify a different remote server upload destination: Edit the prefilled **Host Name**, **Port**, **Username**, **Password** and **Remote Path** fields to specify a different destination.

g) (Optional) Click **Verify Backup Readiness** to verify that Crosswork Network Controller has enough free resources to complete the backup. Crosswork Network Controller will also confirm that none of the applications are being updated, if the remote destination is correctly defined and if the applications are healthy. If the verifications are successful, Crosswork Network Controller displays a warning about the time-consuming nature of the operation. Click **OK**.

If the verification is unsuccessful, please contact the Cisco Customer Experience team for assistance.

h) Click **Start Backup** to start the backup operation. Crosswork Network Controller creates the corresponding backup job set and adds it to the job list. The Job Details panel reports the status of each backup step as it is completed.

5. Complete the backup review and upload actions.

a) To view the progress of a backup job: Enter the job details (such as Status or Job Type) in the search fields in the **Backup and Restore Job Sets** table. Then click on the job set you want.

The **Job Details** panel displays information about the selected job set, such as the job Status, Job Type, and Start Time. If there's a failed job, hover the mouse pointer over the icon near the **Status** column to view the error details.

> **Note**
>
> If you do not see your backup job in the list, refresh the **Backup and Restore Job Sets** table.

b) If the backup fails during upload to the remote server: In the **Job Details** panel, just under the Status icon, click the **Upload backup** button to retry the upload.

> **Note**
>
> Upload can fail due to connectivity problems with the SCP backup server (for example, incorrect credentials, missing directory or directory permissions, missing path and so on). This is indicated by failure of the task `uploadBackupToRemote`). If this happens, check the SCP server details, correct any mistakes and try again. Alternatively, you can use the **Destination** button to specify a different SCP server and path before clicking **Upload backup**.

Update DNS server and run migration

Update the DNS server and restore the saved backup on the new cluster using the specified SCP server and backup file.

Update DNS and restore the backup on the new cluster.

Before you begin

Before you begin, ensure that you have:

- The hostname or IP address and the port number of a secure SCP server.
- The name and path of the backup file created in .
- User credentials with file read and write permissions to the directory.

Procedure

1. Update the DNS server to point the FQDN of the previous version of Crosswork Network Controller cluster to the <VIP> of the new Crosswork Network Controller cluster.
2. Navigate to the upgraded Cisco Crosswork UI using https://<new_VIP>:30603.
3. **Configure an SCP backup server:**
   a) From the main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel.
   c) Make the relevant entries in the fields provided.

   > **Note**
   >
   > In the **Remote Path** field, please provide the location of the backup created in Back up the cluster on page 389.

   d) Click **Save** to confirm the backup server details.
4. Migrate the old Crosswork Network Controller backup:
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Actions** > **Data Migration** to display the **Data Migration** drawer panel with the destination server details prefilled.
   c) Provide the name of the data migration backup (created in Back up the cluster on page 389) in the **Backup File Name** field.
   d) If you want to perform the data migration backup despite any Crosswork Network Controller application or microservice issues, check the **Force** check box.
   e) Click **Start Migration** to start the data migration operation. Crosswork Network Controller creates the corresponding data migration job set and adds it to the **Backup and Restore Job Sets** table. The Job Details panel reports the status of each backup step as it is completed.

> **Note**
>
> If you do not see your job in the list, refresh the **Backup and Restore Job Sets** table.

f) To view the progress of a data migration job: Enter the job details (such as Status or Job Type) in the search fields in the **Backup and Restore Job Sets** table. Then click on the job set you want.

The **Job Details** panel displays information about the selected job set, such as the job Status, Job Type, and Start Time. If there's a failed job, hover the mouse pointer over the icon near the **Status** column to view the error details.

> **Note**
>
> Crosswork UI and Grafana monitoring might become temporarily unavailable during the data migration operation.

g) If the data migration fails in between, you need to restart the procedure from step 1.

5. After the data migration is successfully completed, check the health of the new Crosswork Network Controller cluster.
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**.
   b) Click **Crosswork Cluster** tile to view the health details of the cluster.

> **Note**
>
> After a successful migration, please perform a hard refresh or browser cache deletion before proceeding to use the system. Failing to do this step can result in data discrepancy.

Add Crosswork Data Gateway to Cisco Crosswork

Enroll the new Data Gateway VMs in Cisco Crosswork and verify their state, pools, attached devices, and post-upgrade configuration.

Enroll the new Crosswork Data Gateway VMs and verify their upgraded configuration.

Ensure that the migration is complete and the new Cisco Crosswork UI is available before you proceed with installing the new version of Crosswork Data Gateway.

Note

This procedure is required only for a data gateway Base VM upgrade. Upgrade of other components, such as collectors, is performed by Cisco Crosswork.

Crosswork Data Gateway functions because a passive device in the network. The Crosswork Data Gateway upgrade process consists of replacing all old data gateway VMs with the new data gateway VMs (latest version) in the network.

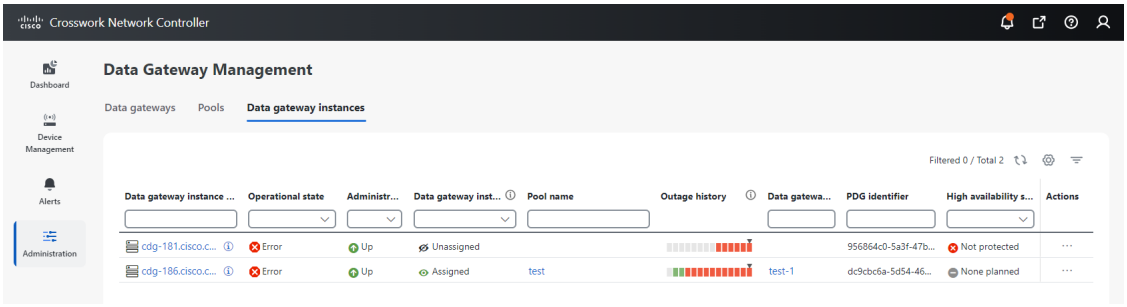
Note

Step 6 in this procedure requires you to log out of Cisco Crosswork and log in again after verifying the deployment and enrollment of the new data gateway VMs with Cisco Crosswork. After you log in, an **Action to be taken** window appears prompting you to confirm that the upgrade is complete. Do not click **Acknowledge** unless you have completed all the verification steps that are mentioned in Step 3, Step 4 and Step 5 in the procedure.

Procedure

1. Enroll and verify the new Crosswork Data Gateway VMs.
   a) Log out of the upgraded Cisco Crosswork and log in again.
   b) After you log in, an **Action to be taken** window appears. Close this window and do not click **Acknowledge**.
   c) Install new data gateway VMs (latest version) with the same number and the same information (management interface importantly) as the old data gateway VMs. Follow the steps in the Crosswork Data Gateway on KVM on page 200.
   d) Wait for about 5 minutes and navigate to **Administration** > **Data Gateway Management**.
   e) Check the **Data Gateway Instances** tab to verify that the new data gateway VMs are enrolled with the new Cisco Crosswork, and have the **Admin State** as **Up** and **Operational State** as **Not Ready**.

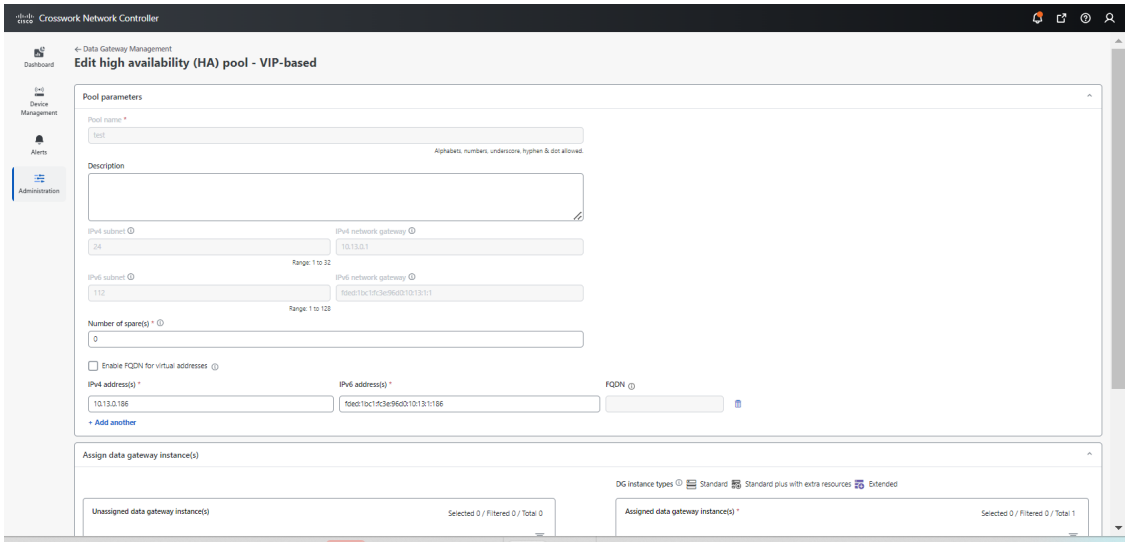
   Figure 53: Data gateway instances window



2. After the **Operational State** of the VMs changes to **Ready**, navigate to the **Pools** tab and verify that all the Crosswork Data Gateway pools from the old Cisco Crosswork, are listed here. Edit each Crosswork Data Gateway pool to verify that the active Crosswork Data Gateway is same as one that you noted in the older version of Cisco Crosswork.

   For example, the Crosswork Data Gateway pool in these image contains two VMs, where the active VM is 172.23.247.78

   Figure 54: Edit ha pool window



3. **Verify that devices are attached to the new data gateways in the upgraded Cisco Crosswork UI.**
   a) Navigate to the **Administration** > **Data Gateway Management** page.
   b) Check the **Attached Device Count** of the data gateway.
4. Log out of Cisco Crosswork and log in again.
5. After you log in, Cisco Crosswork presents you with a window prompting for confirmation that the VMs. In the pop up that appears, click **Acknowledge**.

   Note

   Do not click **Acknowledge** unless you have verified that the VMs are in the **Up/Not Ready** state. Doing so will result in VMs having the state as **Error**. See Troubleshoot Crosswork Data Gateway Upgrade Issues.

6. To move Cisco NSO out of maintenance or read-only mode, use:

   ```
   ncs_cmd -c maapi_read_write
   ```

After the upgrade is complete:

- The new data gateway VMs are enrolled with upgraded Cisco Crosswork.
- All destinations, HA Pools, device mapping information can be viewed on the Cisco Crosswork UI with the upgraded data gateway VMs.
- Jobs start again automatically with the new data gateway VMs.

Shut down the old cluster

Record the old cluster and Data Gateway IP addresses, then shut down the old cluster after migration.

Record the old cluster and Data Gateway IP addresses, then shut down the old cluster.

Before you begin

Gather these information before shutting down the older version of Cisco Crosswork:

- All the IP addresses in the cluster.
- All the IP addresses of the CDGs.

Procedure

1. After a successful backup, shut down the Crosswork Network Controller cluster by powering down the VMs hosting each node (start with the Hybrid VMs):
   a) Log into the VMware vSphere Web Client.
   b) In the **Navigator** pane, right-click the VM that you want to shut down.
   c) Choose **Power** > **Power Off**.
   d) Wait for the VM status to change to **Off**.
   e) Wait for 30 seconds and repeat steps 1a to 1d for each of the remaining VMs.
2. Shut down the data gateway VMs.
   a) Log in to the previous version of the data gateway VM. See Access Crosswork Data Gateway VM from SSH on page 156.

      Crosswork Data Gateway launches an Interactive Console after you login successfully.

   b) Choose **5 Troubleshooting**.
   c) From the **Troubleshooting** menu, choose **5 Shutdown VM** to shut down the VM.
3. (Optional) Move Cisco NSO into read-only mode to avoid any unintended updates to Cisco NSO during the upgrade. Use these command to move NSO to read-only mode:

   ```
   ncs_cmd -c maapi_read_only
   ```

   For more information, please refer to the relevant Cisco NSO documentation.

Upgrade Crosswork Network Controller applications (standalone activity)

Lists the preparation and maintenance-window guidance for upgrading Crosswork Network Controller applications as a standalone activity.

Prepare to upgrade Crosswork Network Controller applications as a standalone activity.

This section explains how to independently update Crosswork Network Controller applications from the UI in case of minor updates or patch releases. This procedure is not part of the upgrade workflow discussed in the earlier sections.

Crosswork Network Controller leverages rolling upgrade functionality to seamlessly update its components while maintaining high availability. It is a deployment strategy that allows applications or services to be updated incrementally, replacing individual instances (e.g., pods or application components) one at a time. This approach ensures that some instances remain operational throughout the upgrade process, minimizing service interruptions and downtime.

Warning

Do not use this rolling upgrade procedure for Release 7.2.1 SP. To move to Release 7.2.1 SP, use the fresh installation and backup-restore migration workflow described in this chapter.

Note

- During a rolling upgrade, `Tyk` will restart and be unable to process requests for approximately 30 seconds to 1 minute. During this time, APIs will be unavailable, and a "Not Found" popup error is displayed in the UI.
- Crosswork Network Controller does not support downgrading an application. To revert to an older version, uninstall the current application and reinstall the desired version. It is recommended to back up your data before performing this operation.

Before you begin

- Take a backup of your data (using the backup/restore functionality) before any critical upgrade.
- Download the latest version of the Crosswork Network Controller tier to your local machine.
- Perform upgrade operations in a maintenance window to minimize disruptions.

Procedure

1. **Download and validate the application files:**
   a) Navigate to cisco.com and locate the files (.tar.gz) that you require.
   b) Hover over the file and copy the MD5 or SHA512 checksum to your clip board.
   c) Download the files to a server that can be reached from the Crosswork server.
   d) Run a tool of your choice to calculate the checksum, and compare the checksum value in your downloaded file with the value you copied in the clip board.

      For example, on a MAC you can use the **md5** command to calculate the MD5 sum on a file:

      ```
      md5 signed-cw-na-coe-7.2.1.tar.gz

      ff47a72ed7dc4fc4be388db3a43fa13f
      ```

      Verify that the result value matches with the posted value on cisco.com.

2. Click on **Administration** > **Crosswork Manager**, and select the **Application Management** tab.

   The Crosswork Platform Infrastructure and any applications that are added are displayed here as tiles.

3. Click **Add new file** > **Upload application bundle (.tar.gz)**.

   The **Add Application Bundle (.tar.gz)** dialog box is displayed.

4. In the dialog box, choose either **URL** or **SCP** as your preferred protocol. Based on your selection, fill in the additional fields with the required information.

   Click **Add** to proceed.

5. To upgrade, click the Upgrade prompt and the new version of the application is installed. Alternately, click ··· on the tile, and select the **Upgrade** option from the drop down list.

   In the Upgrade screen, select the new version that you want to upgrade to, and click **Upgrade**.

6. (Optional) Click on **Job History** to refer to the progress of the upgrade operation.

> **Note**
>
> During an upgrade, only the changed components are installed, with the new version reusing most resources from the older version to ensure a fast and non-disruptive process. However, the application being updated will be degraded until the upgrade is complete, and users will be notified of the upgrade through an alarm.

Upgrade Crosswork Network Controller on a single VM

Describes the single-VM migration path and lists the supported source and destination deployment profiles.

This section describes the migration process for moving data from a Crosswork Network Controller deployed on a single VM to a newer release.

1. Migration prerequisites on page 398
2. Back up the current release data on page 399
3. Install the newer Crosswork Network Controller version on page 399
4. Run data migration on page 400

Supported migration scenarios

Ensure you are not downgrading your profile during the migration. These combinations are supported:

- Essentials to Essentials
- Essentials to Advantage
- Essentials to Premier
- Advantage to Advantage
- Advantage to Premier

Migration prerequisites

Lists the access, backup-server, installation-file, and network prerequisites for single-VM migration.

Before starting the migration, ensure these:

Before starting the migration, ensure these:

- You have administrative access to the Crosswork Network Controller UI and the VM environment.
- A remote server is available to store and retrieve backup files.

- You have recorded all IP addresses of the current Crosswork Network Controller instance, including the Management Virtual IP (VIP) and physical node IP addresses.

Back up the current release data

Configure the backup destination and create a backup of the current single-VM release before continuing.

Create a backup of the current release data before the upgrade.

You must create a backup of your current data before proceeding with the upgrade.

Procedure

1. **Configure the backup destination**:
   a) Log in to the Crosswork Network Controller as a user with admin privileges.
   b) Go to **Administration** > **Backup and Restore**
   c) Click **Destination** to display the **Edit Destination** drawer panel. Make the relevant entries in the fields provided.
   d) Click **Save** to confirm the backup server details.
2. **Perform the data backup**:
   - **Back up from the Crosswork Network Controller UI**:
     a. Click **Actions** > **Data Backup** to display the **Data Backup** drawer panel with the destination server details prefilled.
     b. Provide a relevant name for the backup in the **Job Name** field. Complete the remaining fields as needed.
     c. (Optional) Click **Verify Backup Readiness** to verify there are enough resources to complete the backup. If the verifications are successful, Crosswork Network Controller displays a warning about the time-consuming nature of the operation. Click **OK**.
     d. Click **Start Backup** to start the backup operation. Crosswork Network Controller creates the corresponding backup job set and adds it to the job list. The Job Details panel reports the status of each backup step as it is completed.
     e. Once the status is **Completed**, verify the file on your remote storage server. The filename typically includes the job name, the source version identifier, and a timestamp.
     f. **Note the exact filename**, as it is required for the migration step.
   - **Back up using REST API (data only)**:
     a. Obtain a JWT token using SSO APIs.
     b. Send a POST request to the data backup endpoint:

     ```
     https://<VIP>:30603/crosswork/platform/v1/platform/backup/dataonly
     ```

     c. Use these payload: `{"jobName": "your_job_name", "force": false}`

Install the newer Crosswork Network Controller version

Prepare the environment and install the newer Crosswork Network Controller version after backing up the current data.

Prepare the environment and install the newer Crosswork Network Controller version.

After successfully backing up your data, prepare the environment and install the newer version of Crosswork Network Controller.

Procedure

1. **Shut down the current instance of Crosswork Network Controller:**:
   a) Record all IP addresses from the current Crosswork Network Controller VM, specifically the Management Virtual IP (VIP) and the physical node IP addresses.
   b) Power down the Crosswork Network Controller VM.

      For VMware deployments:

      - In the vCenter UI, go to the **Navigator** pane. Right-click the VM, and choose **Power** > **Power Off**.
      - Wait for the VM status to change to **Off**.

   c) Delete the VM to clear the environment for the newer release installation.
2. **Install the newer release**:
   a) Install the newer version of Crosswork Network Controller on a single VM. For detailed instructions, see Introduction on page 332.

      > **Note**
      >
      > You must use the same IP addresses (VIP and physical IPs) that were used for the previous installation.

   b) Once the installation is complete, log in to the Crosswork UI at `https://<VIP>:30603`.
   c) Wait for the VM and all applications to reach a **Healthy** state.
   d) Verify that the Crosswork inventory is imported. If it is not present, import it manually via the UI.

Run data migration

Restore the previous backup on the new VM and verify that the migrated data and services are available.

Migrate the previous backup to the new VM.

Migrate the previous backup on the new Crosswork Network Controller VM:

Procedure

1. Check and confirm that your new VM is healthy and running.
2. **Configure an SCP backup server:**
   a) From the main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel.
   c) Enter the details of the remote server where the backup file is stored. These details must match the destination used in Back up the current release data on page 399.
   d) Click **Save** to confirm the server details.
3. Migrate the previous Crosswork Network Controller backup on the new Crosswork Network Controller VM:
   a) On the **Backup and Restore** window, click **Actions** > **Data Migration** to display the **Data Migration** drawer panel with the destination server details prefilled.

b) Provide the name of the exact data migration backup created in Back up the current release data on page 399.
c) If you want to perform the data migration backup despite any Crosswork Network Controller application or microservice issues, check the **Force** check box.
d) Click **Start Migration** to start the data migration operation. Crosswork Network Controller creates the corresponding data migration job set and adds it to the **Backup and Restore Job Sets** table. The **Job details** panel reports the status of each backup step as it is completed.
e) Monitor the progress in the **Job details** table. You can also monitor the process via the dashboard link provided in the UI.
f) Wait for the job to complete. This process may take several hours depending on the volume of data.

4. After the data migration is successfully completed, check the health of the new Crosswork Network Controller VM.
a) From the main menu, choose **Administration** > **Crosswork Manager** > **Crosswork Summary**.
b) Click **System summary** tile to view the health details of the VM.
c) Click **Application management** and verify that all applications have returned to a **Healthy** state.
d) Go to **Devie Management** > **Network Devices**. Select **Actions** > **Detailed sync all devices**to synchronize your network devices once they are active and reachable.

15 Security Patches for Crosswork Network Controller

Topics:

- Security patches for OS and core components
- Crosswork Data Gateway patches and patch management

Describes patch types and workflows for Crosswork Network Controller and Crosswork Data Gateway. It also covers applying, monitoring, and verifying patch upgrades.

This chapter describes supported security patch types and patch workflows for Crosswork Network Controller and Crosswork Data Gateway. It also explains how to apply, monitor, and verify patch upgrades.

Security patches for OS and core components

Explains how security patches are packaged and applied to Crosswork Network Controller components, including supported patch types, the update workflow, and multi-cluster geo deployment order.

A security patch or operating system patches are essential updates designed to protect Crosswork Network Controller from vulnerabilities, improve system performance, and maintain compatibility with the latest platform requirements. These updates are applied with minimal disruption, targeting one node at a time to ensure system availability.

Security patches for the Crosswork Network Controller are offered as tar bundles. Each tar bundle contains packages in a unique format specific to Crosswork, known as SPACK (System Package). The naming convention for these tar bundles is: `cw-na-spack-7.2.1-sp.<version>-<release>-<date>.tar.gz`

Supported patch types within the SPACK bundle

Security patch update workflow

- OS security updates
- Kubernetes security updates
- Container Runtime Interface security updates
- Calico security updates
- Robot Orchestrator pod upgrade
- Custom binaries such as `vmexec`

The patching process supports base VM patches and includes inline patching support for OS, Kubernetes, Calico, and Container Runtime Interface security patches.

The patching process facilitates the update of packages on the base VM and other core components. This is a series of steps in the security patch update workflow:

1. **Backup**:It is advised to take a backup before starting the patch update procedure to ensure data safety in case of unforeseen issues.
2. **Maintenance mode**:Maintenance mode will be activated automatically before the patching procedure begins and deactivated after the patching is completed. Alarms will be raised to indicate the activation and deactivation of maintenance mode.
3. **Patch application order**:
   - Each SPACK bundle contains one or more updates such as OS, Kubernetes, and Calico, and these patches are applied in the order defined in the bundle manifest.
   - Patches are applied to all nodes sequentially, starting with the **worker nodes** and then proceeding to the **hybrid nodes**, targeting one node at a time to ensure system availability.
4. **Job tracking**:
   - After applying security patches, individual jobs are created in the **System Management** > **Job History** window for each node in the cluster.
   - Operations performed on each node can be reviewed in the **Job Details** page to understand the actions being executed.
5. **Handling failed jobs**: If a job fails on a particular node, you can retry the job. The retry process will start from the failed node while ignoring nodes that were successfully patched.

6. **Adding new worker nodes**: If a new worker node is added after the patching process, patches must be manually applied to that node. The system does not automatically apply patches to newly added nodes.
7. **Estimated time for completion**: It typically takes 3 to 4 hours to apply patches on a 5-node cluster. It is recommended to schedule all upgrades during a maintenance window to minimize disruptions.
8. **Verification**: Once all jobs on all nodes are completed successfully, you can verify the updated package list by navigating to **Administration** > **Crosswork Manager** > **System Summary** > **<Node-Name>** > **View details** > **Package details** tab.

> **Note**
>
> - All packages are bundled and do not require any external connectivity to fetch packages or configurations.
> - The security patch update workflow applies to both clusters and single VM setups. For Geo HA (High Availability) setups, patches must be applied to both the active and standby sites.
> - There may be node restarts during the patch operation.
> - Rollback is not supported. If an issue occurs, you should retry the patch process or use an updated system patch bundle with fixes.

Add a security OS patch

Add an operating system patch to the Crosswork Network Controller system, monitor the repository and cluster jobs, and optionally verify the updated package details.

Add and install an operating system patch on the Crosswork Network Controller system.

Before you begin

Before installing the patch, verify that all inventory details are visible in the System summary page (**Administration** > **Crosswork Manager** > **System summary**). Patch installation will fail if this information is missing.

Procedure

1. From the main menu, choose **Administration** > **Crosswork Manager**. The **Crosswork Summary** page is displayed with Crosswork Cluster and Crosswork Platform Infrastructure tiles.
2. Click the **System management** tab.
3. Click **Add OS patch**, and the **Add OS Patch** popup windowis displayed. Choose either **URL** or **SCP** as your preferred protocol. Based on your selection, fill in the additional fields with the required information. Click **Add** to proceed.

   Individual jobs are created to add the patch to repository, and to apply the package to each VM in your cluster.
4. Click **System Management** > **Job History** window to monitor each job created. You can review operations performed on each node in the **Job Details** page to understand the actions being executed.

   The newly added OS patch will be displayed under the **System updates** tab.
5. You can upgrade an already installed OS patch by using the **Upgrade** button displayed under the **System updates** tab.
6. (Optional) Once all jobs on all nodes are completed successfully, you can verify the updated package list by navigating to **Administration** > **Crosswork Manager** > **System Summary** > **<Node-Name>** > **View details** > **Package details** tab.

The **Package details** drawer panel will be displayed with details such as package name, version and description.

Events are generated for system patch success or failure:

- Success:

```
System package <version> update completed for <node name>
```

- Failure:

```
System package <package name> update failed for <node name>
```

```
System update failed due to <reason>, please fix the system and retry the
 operation again
```

Crosswork Data Gateway patches and patch management

Explains how Crosswork Data Gateway patch management updates the base virtual machine, supports rolling upgrade scenarios, and handles geo redundancy and HA pool conditions.

Patch management in Crosswork Data Gateway refers to the process of applying minor updates or fixes to the Data Gateway base virtual machine to ensure optimal performance, reliability, and security.

Supported patch types

While patching focuses on maintaining the integrity and functionality of the base VM, the upgrade and management of other components, such as collectors, are handled separately by Crosswork Network Controller. These updates are applied with minimal disruption to data collection activities and are managed through a rolling upgrade process to ensure system availability.

The patching process facilitates base VM upgrades and applies security patches. It also ensures that newly enrolled Data Gateway VMs automatically receive the applicable patches.

Precautions during Data Gateway upgrade

States the precautions for Data Gateway upgrades, including using a unique stack name and matching updated deployment parameters to their previous values.

Important: Parameter mismatch or stack name reuse not allowed.

Apply patch upgrades to Crosswork Data Gateway

Apply a patch to the Crosswork Data Gateway base VM by uploading the patch bundle, then review the resulting enrollment, destinations, pools, and collection state.

Apply a patch upgrade to the Crosswork Data Gateway base VM.

Before you begin

Complete these requirements before you start the Crosswork Data Gateway base VM upgrade:

- Review the available patches and obtain the applicable Crosswork Data Gateway patch file from cisco.com.
- Schedule the patching during a maintenance window to minimize impact on network operations.
- Understand that the upgrade procedure applies only to the Crosswork Data Gateway base VM. Crosswork Network Controller automatically upgrades other components, such as collectors.

- Read the precautions that you must take when performing minor upgrades. For more information, refer to Precautions during Data Gateway upgrade.
- Ensure that you have the Crosscluster Infra API permission to access the **Data Gateway Management** page. For details on user roles and permissions, refer to the *User Roles, Functional Categories, and Permissions* section in *Cisco Crosswork Network Controller 7.2.x Administration Guide*

With the patch upgrade process, minor updates or patch releases are applied to the Crosswork Data Gateway base VM with the less disruptive upgrade process.

Procedure

1. From the main menu, choose the **Administration > Crosswork Manager > Application Management** tab.
2. Click **Add new file > Upload application bundle (.tar.gz)**.
   The **Add Application Bundle (.tar.gz)** dialog box appears.
3. In the Add Application Bundle (.tar.gz) dialog box, configure the patch bundle source:
   - select the protocol as **URL** or **SCP** to access the installation file
   - enter the location (path or address) where the installation file is stored
   - select the **Automatically clean all repository files before adding new one** check box to remove any existing repository files before adding new ones, and
   - select the **Basic Auth** check box if the file source requires authentication.
4. Click **Add**.
   A message appears confirming that the patch installation has started.

After the upgrade is complete:

- The Data Gateway VMs are enrolled with Crosswork Network Controller.
- All the destinations, pools, device-mapping information is available on Crosswork Network Controller UI.
- The upgraded Data Gateway VMs start collecting data from the devices.

What to do next

- If you delete a Data Gateway from the Cisco Crosswork Network Controller after a successful upgrade and re-enroll it later (for example, during maintenance), the patch upgrade status will not be displayed. To re-enroll the Crosswork Data Gateway, refer to the *Re-enroll Crosswork Data Gateway* sections in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*
- Monitor the progress of the installation process. refer to Monitor the Data Gateway patch status.
- Review the events for issues.

Monitor the Data Gateway patch status

Monitor a security patch upgrade in the Data Gateway Management or Events page and inspect gateway state changes, upgrade details, timestamps, and duration.

Monitor the progress of a security patch installed on a Data Gateway VM.

Monitoring the upgrade process is essential to ensure a smooth, reliable, and recoverable patch deployment.

Follow one of these methods to monitor the progress of the security patch installed on Data Gateway VM.

Procedure

1. Monitor the patch status on the Data Gateway Management page.
   a) Go to the **Data Gateway Management** page. The Data Gateway currently being upgraded will appear with the **Admin state** status as in progress.
   b) In the **Admin state** column, hover over the info icon next to the Data Gateway to view details about the package being applied.

      A banner also appears on the **Data Gateway Management** page indicating that the Data Gateway is being upgraded to a specific version.

   c) Go to the **Data Gateway Management** page, click the Data Gateway that is being upgraded. The Data Gateway details page opens.
2. Monitor the patch status on the Events page.
   a) Go to the **Data Gateway Management** page, click the Data Gateway that is being upgraded. The Data Gateway details page opens.
   b) Go to the **Events** pane and review the details, including operational state changes, role changes, a message indicating the reason for the status change, timestamp, and duration.

Patch progress and status tracking

Summarizes the patch status information shown in the Crosswork Network Controller banner, including gateway progress, overall status, and container image versions.

During the patching process, the Crosswork Network Controller UI displays a banner that helps you monitor the upgrade status of each Data Gateway.

The banner provides information about:

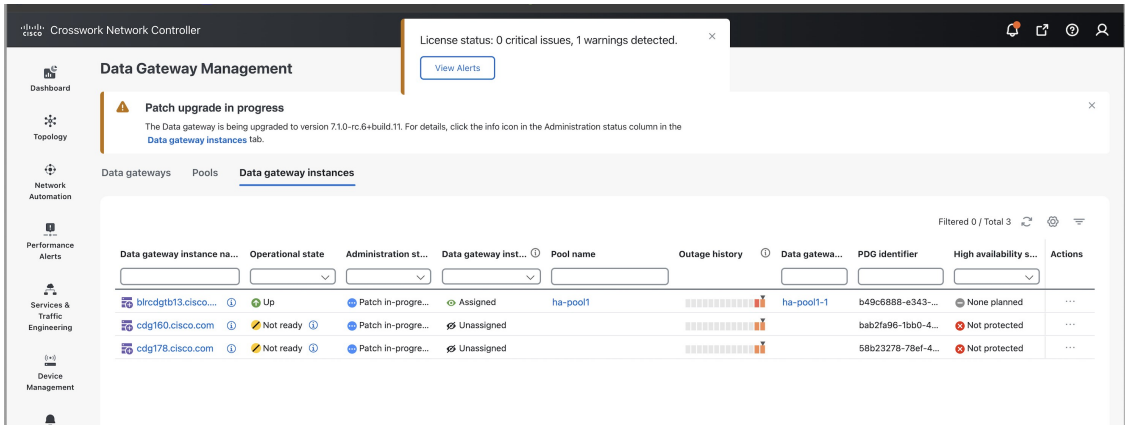
Patch status at Data Gateway level: Track the progress of each individual Data Gateway as the patch is applied.

Overall patch status: View the complete upgrade timeline, from when the DG-manager starts processing the patch to when all Data Gateway VMs are successfully upgraded.

Container image version details: View the previous and current image versions for each container updated by the patch on each Data Gateway.

- Patch status at Data Gateway level: Track the progress of each individual Data Gateway as the patch is applied.
- Overall patch status: View the complete upgrade timeline, from when the DG-manager starts processing the patch to when all Data Gateway VMs are successfully upgraded.
- Container image version details: View the previous and current image versions for each container updated by the patch on each Data Gateway.

Patch status



Verify the patch upgrade status

Verify that the Data Gateway has returned to the UP state, the upgrade banner has cleared, and the patch details show the expected status, version, and build date.

Confirm that the security patches were applied to the Data Gateway VM.

Follow one of these methods to confirm that the security patches were applied to the Data Gateway VM.

Procedure

1. Verify the patch status on the Data Gateway Management page.
   a) Go to the **Data Gateway Management** page.
   b) Confirm that the operational state of the upgraded Data Gateway status is in the UP state.
   c) Verify that the upgrade banner is no longer visible. Its disappearance indicates that the patch process is completed.
2. Verify the patch details.
   a) Go to the **Data Gateway Management** page and hover over the upgraded node.
   The **Enrollment Details** drawer opens.
   b) In the drawer, go to the **Patch Upgrade Details** pane.
   c) Review these values to verify the patch application:
      - Status
      - Installation version, and
      - Build date.

Part IX

Uninstall Cisco Crosswork Network Controller

Topics:

- Uninstall Crosswork Network Controller

16 Uninstall Crosswork Network Controller

Topics:

- Uninstall the Crosswork cluster
- Uninstall Crosswork Data Gateway
- Uninstall Crosswork applications

Describes the available procedures for removing the Crosswork cluster, Crosswork Data Gateway, and Crosswork applications.

This chapter describes how to uninstall the Crosswork cluster, Crosswork Data Gateway, and Crosswork applications.

Uninstall the Crosswork cluster

Lists the Docker installer tool and vSphere UI procedures for uninstalling the Cisco Crosswork cluster, including cleanup after installation failure or changes before redeployment.

You can uninstall the Cisco Crosswork cluster using the Docker installer tool or the vSphere UI.

- Delete the VM using the Docker installer tool
- Delete the VM using the vSphere UI

Delete the VM using the Docker installer tool

Use the Docker installer cleanup operation to remove a cluster deployment after a failed installation or before redeploying when VM settings or data center host information changes.

Clean up a cluster deployment after an installation failure or before redeployment.

Use the Docker installer tool to clean up or delete VMs created by a failed installation. Changes to VM settings or the data center host require cleanup before redeployment.

Procedure

1. Change to the directory that stores the deployment information.

   For example, run `cd ~/cw-cluster`.

2. Run the container on the host.

   Example

   ```
   docker run --rm -it -v `pwd`:/data <cw-installer docker container>
   ```

3. Edit a copy of the template file, such as `deployment.tfvars`, to add the data center access parameters.

   Provide the remaining parameters with dummy values or enter them on the command line when you run the operation.

4. Run the `cw-installer.sh` install script with the `clean` directive and the deployment manifest specified with the `-m` flag.

   Add the `-o` option to remove the Cisco Crosswork image template from the data center.

   For example:

   ```
   ./cw-installer.sh clean -m /data/deployment.tfvars -o
   ```

5. Enter **`yes`** when prompted to confirm the operation.
6. Optional: To clean the cluster quickly without verification, run the installer with this command.

   Example

   ```
   docker run --rm -it -v `pwd`:/data <cw installer docker image> -exec
   './cw-installer.sh clean -m /data/deployment.tfvars'
   ```

The installer cleanup option deletes the cluster deployment based on the inventory in the `/data` directory.

Delete the VM using the vSphere UI

Delete a Cisco Crosswork application VM from vCenter using the VMware vSphere Web Client. This procedure also deletes the application data stored on the VM.

Delete an application VM from vCenter using the VMware vSphere Web Client.

This procedure deletes an application VM from vCenter.

> ⚠ **Caution**
>
> This procedure deletes all application data on the VM.

Before you begin

If you are deleting only a Crosswork Data Gateway VM, complete these steps first:

1. Detach the devices from the Crosswork Data Gateway VM. For more information, refer to *Delete Cisco Crosswork Data Gateway VM from Cisco Crosswork* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
2. Delete the Crosswork Data Gateway VM from Crosswork Network Controller as described in this chapter.

Procedure

1. Log in to the VMware vSphere Web Client.
2. In the **Navigator** pane, right-click the application VM that you want to remove and choose **Power** > **Power Off**.
3. After the VM powers off, right-click it again and choose **Delete from Disk**.

The VM is deleted.

Uninstall Crosswork Data Gateway

Provides the procedure for removing Crosswork Data Gateway by deleting its instance from Cisco Crosswork after it meets the required maintenance and attachment conditions.

Remove Crosswork Data Gateway by deleting its instance from Cisco Crosswork.

- Delete the Data Gateway VM from Cisco Crosswork

Delete the Data Gateway VM from Cisco Crosswork

Delete a Data Gateway instance from the Data Gateway Instances page in Cisco Crosswork after verifying its maintenance state and confirming that it is not part of a pool or attached to a device.

Delete a Data Gateway instance from the Data Gateway Instances list in Cisco Crosswork.

Before you begin

The Data Gateway VM that you want to delete must meet these requirements:

- It must be in maintenance mode.
- It must not be part of a pool or attached to a device.

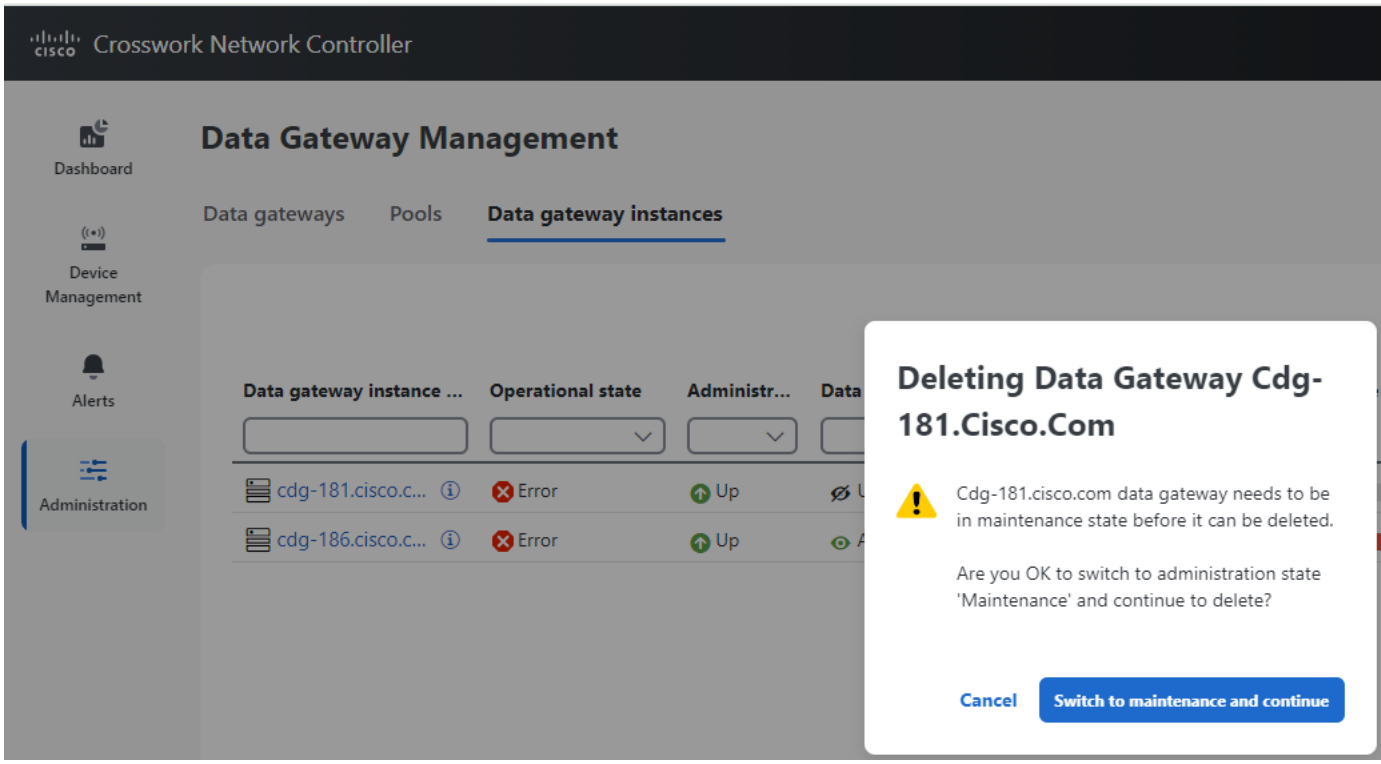
Procedure

1. Log in to the Crosswork Network Controller UI.
2. From the navigation panel, select **Administration** > **Data Gateway Management**.

   Select the **Data Gateway Instances** tab.
3. In the **Data Gateway Instances** list, find the instance that you want to delete and select the ··· in the **Actions** column.

   Select **Delete**.
4. If the Data Gateway instance is not in maintenance mode, Crosswork Network Controller prompts you to switch it to maintenance mode. Select **Switch to maintenance and continue**.

   Figure 55: Switch to maintenance and continue dialog



The Data Gateway instance is deleted.

Uninstall Crosswork applications

Uninstall a Crosswork application and its application-specific menus and associated data. Review application dependencies and adapter worker requirements before you start.

Remove a Crosswork application, its application-specific menus, and associated data.

Use the Crosswork UI to uninstall an application.

Before you begin

- If Crosswork Active Topology is installed, uninstall it before you uninstall Crosswork Optimization Engine.

- Before uninstalling CWMS, remove all adapter workers. The uninstall operation fails if any EMS Adapter, Cisco NSO Adapter, or REST Adapter workers are present. Remove all EMS, Cisco NSO, and REST adapter workers from the dashboard before you uninstall CWMS.

Procedure

1. Choose **Administration** > **Crosswork Manager**, then select the **Application Management** tab.

   The Crosswork Platform Infrastructure and any added applications are displayed as tiles.

2. On the application tile that you want to uninstall, select the ⋯, then select **Uninstall** from the list.
3. Select **Uninstall** to confirm when prompted.

> **Note**
>
> The uninstall operation does not remove the CAPP file from the repository. The file remains visible in the Crosswork UI if you want to install it in the future.

The selected application is uninstalled, and its tile is modified to reflect the change.

What to do next

You can view the uninstallation progress in the **Application Management** > **Job History** window. If the uninstall fails, you can retry it using the relevant options in the same window.

Part X

Enable Geo Redundancy

Topics:

- Geo Redundancy Overview
- Geo Redundancy Requirements
- Deploy the Arbiter VM
- Enable Geo Redundancy Solution
- Upgrade to Geo Redundancy Solution
- Geo Redundancy Switchover
- Geo Redundant Cluster Disaster Recovery

17 Geo Redundancy Overview

Topics:

- Geo redundancy

Introduces geo redundancy for Crosswork Network Controller deployments. Summarizes its benefits, scope, infrastructure considerations, and deployment scenarios.

This chapter introduces geo redundancy for Crosswork Network Controller deployments. It summarizes its benefits, scope, infrastructure considerations, and deployment scenarios.

Geo redundancy

Explains how geo redundancy protects on-premises Crosswork Network Controller deployments against region or data center failures and reduces disruption during system upgrades and overall data loss.

Geo redundancy is an optional high-availability functionality for on-premises Crosswork Network Controller deployments. It supports business continuity during a region or data center failure through geographical or site redundancy.

Geo redundancy provides these benefits:

- It protects against an entire-site failure.
- It reduces disruption during system upgrades.
- It reduces overall data loss.

Support scope

Key factors

Deployment scenarios

The chapters in this part explain the requirements and processes to install or upgrade geo redundancy in the Crosswork Network Controller solution.

> **Note**
>
> Geo redundancy is optional. Contact the Cisco Customer Experience team for assistance.

Geo redundancy involves placing physical servers in geographically diverse availability zones (AZ) or data centers (DC) to safeguard against catastrophic events and natural disasters.

Crosswork Network Controller supports geo redundancy on both cluster and single VM deployments. All requirements and behaviors described in these chapters apply to both cluster and single VM deployments unless stated otherwise.

Consider these factors to ensure geo redundancy:

- *VM node availability:* Ensure that both the active and standby sites have the same number of virtual machines (cluster nodes, NSO, Crosswork Data Gateway, and others) and the same level of network accessibility between the sites and network nodes.
- *Geo availability of nodes:* Physical data centers must not share common infrastructure, such as power and network connectivity. It is recommended that you place them in different availability zones (AZ) or regions to avoid a single point of failure that could affect all VM nodes.
- *Network availability:* To keep the sites synchronized, the network link between the data centers must meet the availability and latency requirements detailed later in this chapter.

Geo redundancy supports these deployment scenarios:

- **Traditional geo redundancy** uses active and standby sites.
- **Arbiter-based geo redundancy** uses active, standby, and arbiter sites.

Caution

You can deploy a geo-HA setup with AZ1 alone, without AZ2 or the arbiter node (AZ3). Although AZ1 starts successfully in this scenario, this setup defeats the purpose of HA, which requires at least AZ1 and AZ2. This configuration significantly limits the geo-HA functionality offered by Crosswork Network Controller and causes these issues:

- Attempting to save the cross-cluster settings triggers an error pop-up that indicates the settings failed to save on the peer clusters.
- The *sync* and *switchover* buttons remain active in the drop-down menu on the cross-cluster UI window, regardless of whether peer sites are present. However, pressing either button to initiate a job fails with error messages indicating the absence of peer clusters.

18 Geo Redundancy Requirements

Topics:

- Crosswork cluster requirements
- Crosswork Data Gateway requirements
- Geo redundancy port requirements
- Unified endpoint requirements
- Data store roles

Lists the infrastructure, port, endpoint, and data-store requirements for geo redundancy in Crosswork Network Controller. It summarizes the details needed to plan a supported deployment.

This chapter lists the infrastructure, port, endpoint, and data-store requirements for geo redundancy in Crosswork Network Controller. It summarizes the details needed to plan a supported deployment.

Crosswork cluster requirements

Summarizes the infrastructure requirements and key configuration notes for on-premises geo redundancy, including synchronization, DNS, dual-stack, inventory, applications, backups, and arbiter resources.

Geo redundancy solution requires double the number of VMs required for a regular Crosswork Cluster installation. For more information, refer to Installation Prerequisites.

Single VM deployments

> **Note**
>
> Crosswork Network Controller deployed on a single VM follows the same infrastructure VM requirements for geo redundancy as cluster deployments. The only exception is that there are no requirements for Data Gateway or NSO, as both are integrated into the single VM build packages.

Important notes

Network and inventory requirements include:

- To ensure synchronization between the clusters, the network connection between the data centers should be set up with a minimum bandwidth of 5 Gbps and a latency of less than 100 milliseconds.
- While preparing a geo inventory file, you must include details of cluster constituents along with connectivity information.
- Configure the DNS server to resolve the unified multi-cluster FQDN (data and management) domain (for example, *.cw.cisco) that you want to use. The DNS server must be reachable from both clusters, the Crosswork Data Gateway, NSO, and SR-PCE. For more information on the DNS setup procedure, refer to the *Cisco Prime Network Registrar Caching and Authoritative DNS User Guide* .
- The DNS server should forward any outside domains to the external DNS servers.
- In case of dual stack deployments, Crosswork cluster and Crosswork Data Gateway deployments should be configured in dual stack mode, with the geo FQDN (data and management) pointing to the dual stack IP addresses.
- You should sequentially bring up the active cluster, standby cluster, and arbiter VM using the existing installer mechanism. Ensure you use the previously identified DNS server during the installation of the Crosswork cluster, Crosswork Data Gateway, and NSO. It is recommended to have multiple DNS servers, with one in each data center.

Additional deployment and operations requirements include:

- Applications should be installed only after enabling geo redundancy and completing the initial on-demand synchronization. Begin by installing applications on the active cluster, followed by the standby cluster. Other configuration information (such as devices, providers, or destinations) must be onboarded only on the active cluster and will be synchronized between the clusters as part of the activation process.
- You should ensure that the Day 0 inventory is onboarded before enabling geo redundancy.
- The NSO in HA mode requires provider connectivity onboarding via an L3 VIP or FQDN in the Crosswork Network Controller.
- Before enabling geo-redundancy mode, you are recommended to make a backup of the active and standby cluster.
- An arbiter VM requires a resource footprint of 8 vCPUs, 48 GB of RAM, and 650 GB of storage. For installation instructions, refer to Deploy the Arbiter VM.

- Before enabling geo redundancy, configure the same SSH and HTTPS passwords on the active and standby Crosswork clusters. Different passwords can cause recurring SSH login test failure events for the peer cluster, even after the inventory YAML file and cross-cluster configuration are updated.

> **Warning**
>
> Once the geo-redundancy mode is set up, it cannot be undone because the certificates are regenerated using common root CA. To revert to non-geo redundancy mode, you must restore the backup made before enabling the geo redundancy mode.

Crosswork Data Gateway requirements

Provides a cross-reference to the minimum Crosswork Data Gateway installation requirements that apply to geo-redundant Crosswork Network Controller deployments in this guide.

Confirm that you have met the minimum requirements outlined for Crosswork Data Gateway installation.

Minimum requirements

For more information, refer to Crosswork Data Gateway VM requirements.

Geo redundancy port requirements

Lists port management and communication requirements for geo-redundant Crosswork deployments, including traffic between sites, Crosswork clusters, Crosswork Data Gateway, NSO, devices, and other components.

This topic lists port requirements and communication flows for Crosswork Network Controller geo-redundancy deployments.

Port management best practices

As a general best practice, any unused ports should be disabled. To identify all open listening ports after installing and activating the applications, log in as a Linux CLI admin user on any Crosswork cluster VM and run the `netstat -aln` command:

> **Note**
>
> All IP addresses, including virtual IP addresses, must be reachable between Crosswork Cluster, Crosswork applications, and Crosswork Data Gateway in both directions.

Additional ports for Crosswork Network Controller cluster

The additional ports, 30180 and 30190, are outside the normal Crosswork cluster deployment. Both ports are used for database replication (PQ binary protocol) and are secured by SSL.

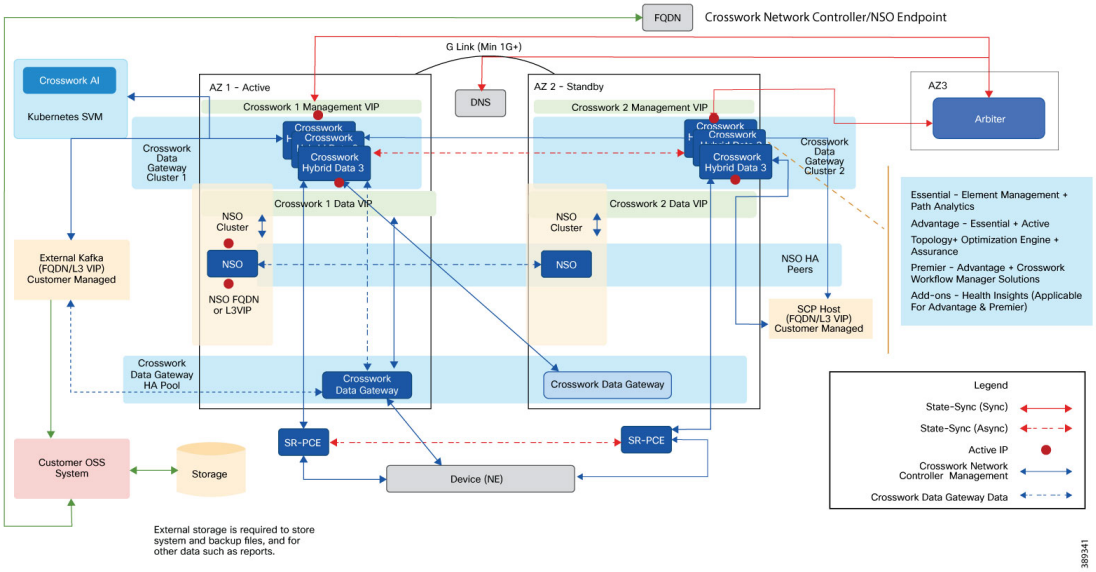
Required TCP or UDP ports

The TCP or UDP port numbers must be allowed through any external firewall or access control list (ACL) deployed by the data center administrator. Depending on the NIC deployment, these ports may be applicable to only one or both NICs.

Traffic types and ports for Crosswork communication with Data Gateway

The diagram shows communication paths between active and standby sites and their Crosswork components.

Figure 56: Communication flow between Crosswork Network Controller and different components



These processes support replication and accessibility between the sites:

- The Crosswork VM intercluster checks are performed, with replication happening over the eth0 management links between both sites.
- NSO replication occurs over the management links between the sites.
- Crosswork cluster replication is asynchronous, while NSO replication is synchronous.
- Devices and the active or standby Crosswork cluster at each site must be able to reach the Crosswork Data Gateway.
- The Crosswork Network Controller NBI is accessible through its DNS FQDN.

The table summarizes communication traffic across sites for two-NIC Crosswork Data Gateway deployments.

Table 89: Two NICs Crosswork communication with Crosswork Data Gateway: comparison of traffic types across sites

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth0 interfaces on both the Active and Standby Crosswork devices exchange bidirectional traffic with each other.	30651, 30652, 30653, 30180, 30190, 22, 30603, 30649, and 30650. The ports are used for these purposes: • 30651: used for external east/west gRPC server communication.. • 30653: used for Raft node to peer node communication. • 22 and 30603: used as TCP ports for intersite communication where 22 for SSH and 30603 for the HTTPS web server. Ensure that these ports are accessible over the eth0 interface on all sites, allowing each site to communicate with every other site through eth0. • 30649 and 30650: used to facilitate communication between internal services across sites.	Crosswork eth0	Crosswork eth0	Crosswork eth0 The Arbiter node uses only ports 30651 and 30653.

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth0 interface on Crosswork Data Gateway (Active and Standby) and eth0 on Crosswork (Active) exchange bidirectional traffic with each other.	The ports are used for these purposes: • 30607: REST API calls to download the collector images, system resource, and custom resource bundle. • 30607: Sending the Crosswork Data Gateway status information to Crosswork. • 30608: Handling the gRPC communication to send Crosswork Data Gateway's heartbeat.	Crosswork Data Gateway eth0 (Management)	Crosswork Data Gateway eth0 (Management)	Not applicable
The eth1 interface on Crosswork Data Gateway (Active) and eth1 Crosswork (Active) exchange bidirectional traffic with each other.	The ports are used for these purposes: • 30607: Handling the gRPC communication with Crosswork's Magellan service. • 30993, 30994, and 30995: Kafka ports that are used to distribute data across Crosswork's Kafka or gRPC destinations.	Crosswork Data Gateway eth1 (Northbound Data)	Crosswork Data Gateway eth1 (Northbound Data)	Not applicable
The eth1 interface on Crosswork Data Gateway (Active) and eth1 Crosswork (Active) exchange bidirectional traffic with each other.	The 30993, 30994, and 30995 ports are used for distributing Icon's data to the standby Crosswork's Kafka.	Crosswork Data Gateway eth1 (Northbound Data)	Crosswork Data Gateway eth1 (Northbound Data)	Not applicable

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth1 interface on Crosswork Data Gateway (Active) and device exchange bidirectional traffic with each other.	The ports are used for these purposes: • 1062: SNMP trap. • 9514, 9898, and 6514: Syslog UDP, TCP, and TLS.	Crosswork Data Gateway eth1 (Southbound Data)	Crosswork Data Gateway eth1 (Southbound Data)	Not applicable

The table summarizes communication traffic across sites when a three-NIC Crosswork Data Gateway is used.

Table 90: Two NICs Crosswork communication with three NIC Crosswork Data Gateway: comparison of traffic types across sites

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth0 interfaces on both the Active and Standby Crosswork devices exchange bidirectional traffic with each other.	30651, 30652, 30180, 30653, 30190, 22, and 30603. The TCP ports are used for these purposes: 22 and 30603: used to facilitate intersite communication where 22 for SSH and 30603 for the HTTPS web server. Ensure that these ports are accessible over the eth0 interface on all sites, allowing each site to communicate with every other site through eth0. 30649 and 30650: used to facilitate communication between internal services across sites.	Crosswork eth0	Crosswork eth0	Crosswork eth0 The Arbiter node uses only 30651 and 30653.

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth0 interface on Crosswork Data Gateway (Active and Standby) and eth0 on Crosswork (Active) exchange bidirectional traffic with each other.	The ports are used for these purposes: • 30607: REST API calls to download the collector images, system resource, and custom resource bundle. • 30607: Sending the Crosswork Data Gateway vitals to Crosswork. • 30608: Handling the gRPC communication to send Crosswork Data Gateway's heartbeat.	Crosswork Data Gateway eth0 (Management)	Crosswork Data Gateway eth0 (Management)	Not applicable
The eth1 interface on Crosswork Data Gateway (Active) and eth1 Crosswork (Active) exchange bidirectional traffic with each other.	The ports are used for these purposes: • 30607: Handling the gRPC communication with Crosswork's Magellan service. • 30993, 30994, and 30995: Kafka ports that are used to distribute data across Crosswork's Kafka or gRPC destinations.	Crosswork Data Gateway eth1 (Northbound Data)	Crosswork Data Gateway eth1 (Northbound and Southbound Data)	Not applicable
The eth1 interface on Crosswork Data Gateway (Active) and Crosswork (Active) exchange bidirectional traffic with each other.	The 30993, 30994, and 30995 ports are used for distributing Icon's data to the standby Crosswork's Kafka.	Crosswork Data Gateway eth1 (Northbound Data)	Crosswork Data Gateway eth1 (Northbound and Southbound Data)	Not applicable

Traffic Type	Ports Used	Site 1 (Active)	Site 2 (Standby)	Site 3 (Arbiter)
The eth2 interface on Crosswork Data Gateway (Active) and device exchange bidirectional traffic with each other.	The ports are used for these purposes: • 1062: SNMP trap port. • 9514, 9898, and 6514: Syslog UDP, TCP, and TLS ports.	Crosswork Data Gateway eth1 (Northbound Data)	Crosswork Data Gateway eth1 (Northbound and Southbound Data)	Not applicable

Unified endpoint requirements

Lists unified endpoint behavior and the DNS and domain-zone requirements for mapping Crosswork component data and management VIPs to FQDNs.

A unified endpoint hides multiple instances for various components in high-availability deployments.

DNS and domain requirements

The unified endpoint requirements include:

- Unified endpoint hides multiple instances in high availability for various components.
- DNS allows endpoints to be referred via FQDN, which should point to the active instance IP (data and management).
- Domain zone provisioning is needed where Crosswork components have data and management VIP mapped to FQDN.
- The DNS authoritative server must have an A or AAAA entry for the IP address from you for the domain zone dedicated to Crosswork components.

Data store roles

Describes data-store roles and expected states in active and standby clusters. It also outlines synchronization operations and the roles they assign to local data stores.

The data-store roles apply to both Geo HA and non-Geo HA deployments when Postgres or Timeseries data stores are operational.

Data store roles

- **Leader**: should be in `running` status.
- **Replica**: indicates a regular standby.
- **Sync Standby**: similar to standby replica, but applicable when synchronous mode is ON in `postgresql.conf`.

For *robot-postgres*, when synchronous mode is ON, the replica is assigned the Sync Standby role. For *cw-timeseries-db*, when synchronous mode is OFF, the replica retains the Replica role. The state will be `streaming` on the respective replica to indicate that local replication is occurring correctly within the PG cluster.

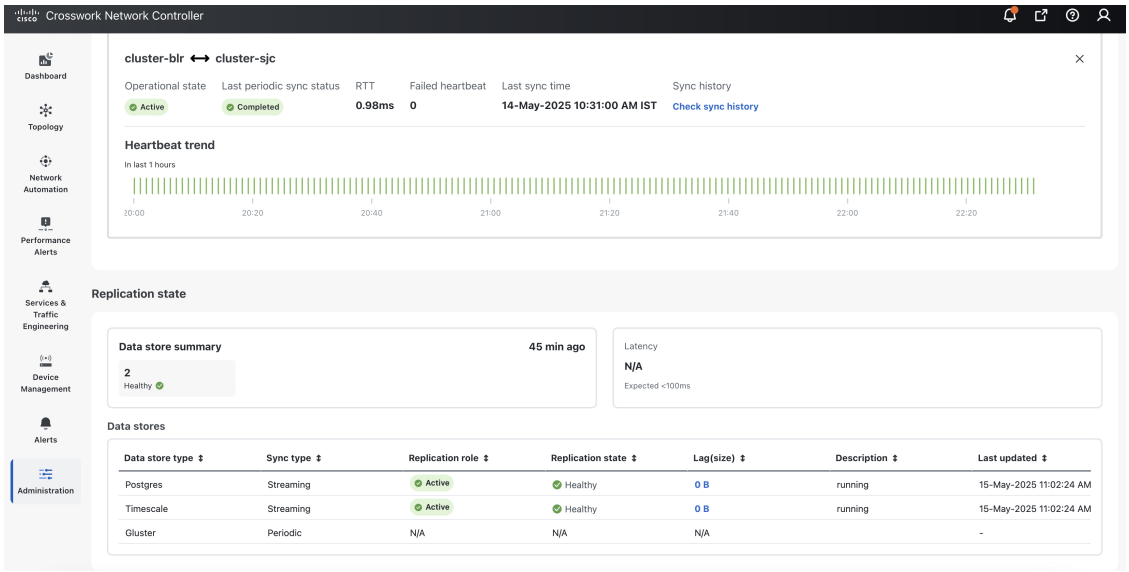
Additionally, in Geo HA deployments, when Cross Cluster replication is enabled and the standby cluster is streaming from a remote active cluster, its leader will assume the **Standby Leader** role, and the status will be `streaming`.

Sync operation

One of the initial steps in a synchronization operation cycle involves setting up asynchronous data replication for local Postgres and Timescale data stores. The system first verifies the health status of the data store. It then aligns the Cross Cluster state and data store state to ensure consistency (for example, if the service *cw-timeseries-db-0* is already active for Timescale DB, the system will assign the ACTIVE_REPLICATION_ROLE to the Timescale data store). After setting the role for the active side, the system will replicate this process for the standby side. Once asynchronous replication is complete, the system initiates backup and restore operations. These operations can also be monitored in the **Backup and restore jobs** window.

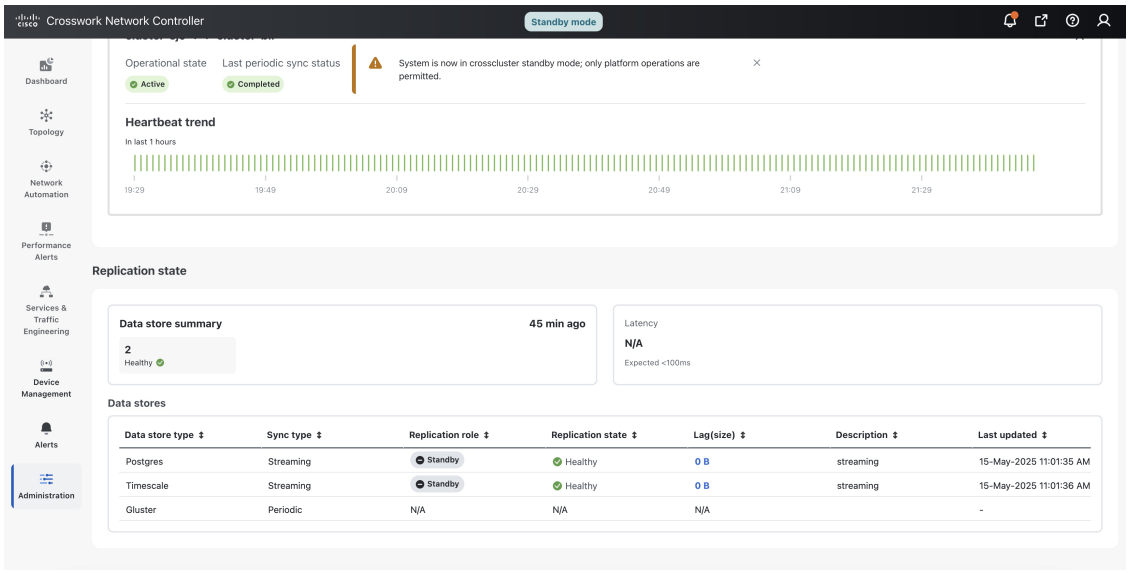
During a sync, the data stores in active and standby clusters have some expected roles. After the sync, the data store roles are verified to confirm a successful sync. The active cluster should contain one Leader with `running` status. The active-cluster figure shows the Postgres and Timescale data-store roles during sync.

Figure 57: Roles in an active cluster during sync



Similarly, the standby cluster should contain one Standby Leader with `streaming` status. The standby-cluster figure shows the Postgres and Timescale data-store roles during sync.

Figure 58: Roles in a standby cluster during sync



19 Deploy the Arbiter VM

Topics:

- Auto-arbitration in Crosswork Network Controller
- Arbiter VM requirements
- Deploy an arbiter VM on VMware
- Deploy an arbiter VM on KVM

Introduces arbiter VM requirements and deployment guidance for VMware and KVM. It also explains how auto-arbitration supports geo-redundant clusters.

This chapter describes the arbiter VM's role in geo-redundant deployments and provides platform-specific deployment guidance.

Auto-arbitration uses the arbiter VM to help coordinate active and standby clusters. The chapter also lists arbiter VM requirements and deployment methods for VMware and KVM.

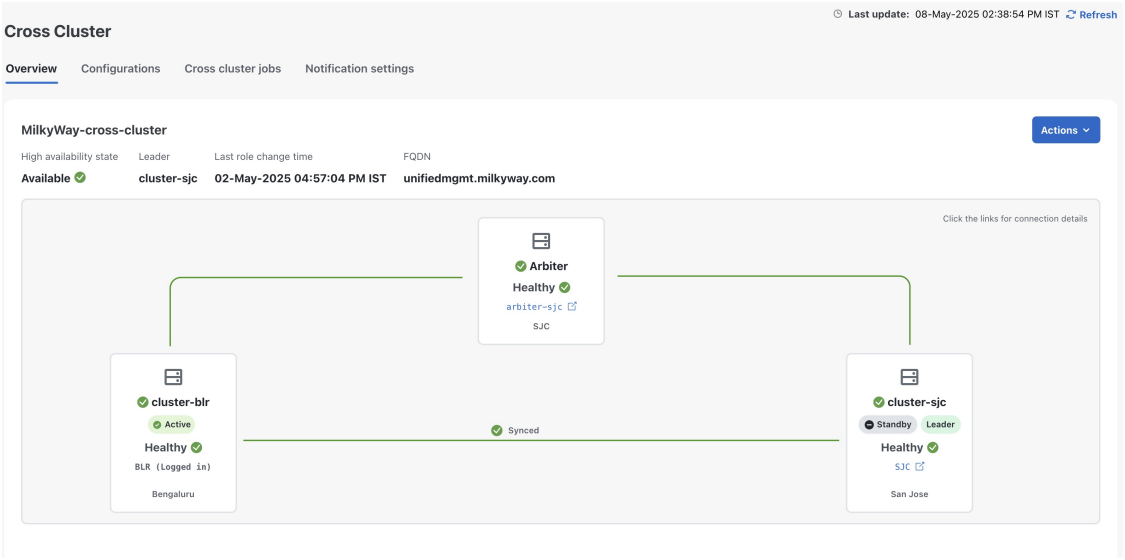
Auto-arbitration in Crosswork Network Controller

Explains how auto-arbitration coordinates switchover between active and standby clusters and how the arbiter VM prevents split-brain scenarios.

Auto-arbitration is a functionality in Crosswork Network Controller that automates the switchover process between active and standby clusters during communication failures, eliminating the need for manual intervention. It introduces an arbiter VM (AZ3), which acts as a mediator to prevent split-brain scenarios and ensure system integrity.

Key features

Figure 59: Cross cluster with active, standby, and arbiter AZs



1. **Split-brain scenario prevention**:
   - In a geo redundant setup, communication failures can cause the standby cluster to mistakenly assume the active cluster is down, leading to both clusters becoming active (split-brain scenario). This can result in data loss if one cluster fails.
   - Auto-arbitration uses an arbiter VM and the leader election protocol to prevent such conflicts by ensuring consensus on the active cluster.
2. **Arbiter VM role**:
   - The arbiter VM is deployed on a single VM with a small resource footprint.
   - It provides the essential quorum needed to form a majority vote for electing a leader among the three AZs.
   - It includes only essential infrastructure components and does not allow application installations after deployment.

> **Note**
>
> By design, the Arbiter VM in the geo HA cluster is a minimal node used primarily for quorum voting. It does not host workload applications like other cluster nodes, but it runs the essential `capp-infra` and `capp-geo` components required for quorum and geo-redundancy orchestration. By default, only the admin user is created on the arbiter node for administrative operations. The arbiter VM does not synchronize AAA users or settings from other nodes. However, you can create and manage additional users independently on the arbiter VM if needed.

3. **Switchover automation**: During a switchover (manual request or network failure), the arbiter node performs these steps.

   When auto arbitration mode is enabled, the elected cluster leader automates the three steps of a switchover, whether triggered manually or by a failure in a network node or link.

   - Updates the active cluster's role to standby.
   - Updates the standby cluster's role to active.
   - Updates the DNS FQDN records for data and management to point to the new active cluster.

   For more information, refer to How auto-arbitration works on page 514.

4. **Day 0 and day N deployments**:
   - **Day 0**: All three clusters (active, standby, and arbiter) are deployed and configured sequentially on day 0. For more information, refer to Geo redundancy workflow (Day 0) on page 443.
   - **Day N**: You can add an arbiter VM to an existing two-cluster geo-redundant model (active and standby) by reimporting the updated cross-cluster inventory file, which includes new parameters for the arbiter VM. For more information, refer to Geo redundancy workflow (Day N) on page 466.

Auto-arbitration is a critical enhancement for ensuring reliable and fail-safe operations in distributed environments, enabling seamless cluster management with minimal manual effort.

Arbiter VM requirements

Lists the deployment location, profile, resource, and inventory requirements for an arbiter VM.

This topic lists requirements for deploying the arbiter VM.

- The arbiter VM is deployed as a single VM using a **small** profile.
- The arbiter VM requires a neutral third site with a different subnet from the two workload sites.
- An arbiter VM requires a resource footprint of **8 vCPUs**, **48 GB of RAM**, and **650 GB of storage**.
- After deploying the arbiter VM, ensure that the Crosswork inventory is onboarded before proceeding.

Deploy an arbiter VM on VMware

Describes the VMware vCenter, OVF tool, and Docker installer methods for deploying an arbiter VM.

Install the Arbiter VM using the arbiter OVA file to reduce manual steps and simplify the installation process.

Before you begin

Prepare the arbiter OVA file for deployment.

Procedure

1. Deploy the arbiter VM through the vCenter UI.

   For the installation procedure, refer to Install Crosswork Network Controller using the vCenter vSphere UI on page 342.

   Set these parameters during deployment:

   - Select the virtual disk format as **Thin provision**.
   - Set **Datafs Disk Size** to **100**.
   - Select **VM type** as **Hybrid**.
   - Set **Cluster seed node** to **True**.
   - Set **Initial node count** to **1**.
   - Set **Initial leader node count** to **1**.
   - Set **Enable Skip Auto Install Feature** to **True**.

   After deployment, right-click the VM and select **Edit Settings**. The **Edit Settings** dialog box appears. Under the **Virtual Hardware** tab, update these attributes:

   - Change **CPU** to **8** for the `Small` profile.
   - Change **Memory** to **48 GB** for the `Small` profile.

2. Install the arbiter VM with the OVF tool.

   For installation instructions, refer to Install Crosswork Network Controller via the OVF Tool on page 348.

   Sample OVF tool template:

```
env variables

numberOfCpus=8
memorySize=49152
EnableSkipAutoInstallFeature="True"
ddatafs=100

ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv
--datastore="${vmDatastore[$i]}" --diskMode="${vmDiskMode[$i]}" \
    --numberOfCpus:"*"=$numberOfCpus  --memorySize:"*"=$memorySize \
    --coresPerSocket:"*"=1 \
   --viCpuResource=-1:$viCpuResource:-1 --viMemoryResource=-1:$memorySize:-1
 \
    --overwrite --noSSLVerify --allowExtraConfig \
    --name="${vmName[$i]}" \
    --prop:"ManagementIPv4Address=${vmMngIP[$i]}" \
    --prop:"ManagementIPv4Gateway=$mngGw" \
    --prop:"ManagementIPv4Netmask=$mngMask" \
    --prop:"ManagementVIP=$mngVIP" \
    --prop:"ManagementVIPName=$mngVIPName" \
    --prop:"DataIPv4Address=${vmDataIP[$i]}" \
    --prop:"DataIPv4Gateway=$dataGw" \
    --prop:"DataIPv4Netmask=$dataMask" \
```

```
    --prop:"DataVIP=$dataVIP" \
    --prop:"DataVIPName=$dataVIPName" \
    --net:"Management Network=$mngNet" \
    --net:"Data Network=$dataNet" \
    --prop:"DNSv4=$dns" \
    --prop:"NTP=$ntp" --prop:"Domain=$domain" \
    --prop:"Disclaimer=Cisco..." \
    --prop:"ddatafs=$ddatafs" --prop:"logfs=$logfs" \
    --prop:"Timezone=$timezone" \
    --prop:"EnableSkipAutoInstallFeature=$EnableSkipAutoInstallFeature" \
    --prop:"CWUsername=cw-admin" \
    --prop:"CWPassword=$pass" \
    --prop:"VMType=${vmType[$i]}" --prop:"IsSeed=$isSeed"
--prop:"InitNodeCount=$nodes" --prop:"InitMasterCount=$masterCount" \
    --prop:"IgnoreDiagnosticsCheckFailure=False" \
    ${CW_IMAGE} vi://"${vcenter_user}:${vcenter_pass}"@"$url"
```

3. Install the arbiter VM with the Docker installer.

   Set these parameters in the tfvars file:

```
ManagerDataFsSize = 100
VMSize = "Small"
EnableSkipAutoInstallFeature = "True"
```

   You can modify the other parameters to match your setup. For complete installation instructions, refer to Install Crosswork Network Controller using the Docker installer tool on page 351.

   Docker deployment command example:

```
docker run --rm -it -v `pwd`:/data <cnc-arbiter-docker-deployment-7.2.1> -exec
 'yes yes | ./cw-installer.sh install -p -m /data/svm_tfvar.txt -o
https://software.cisco.com/crosswork/release-7.2.1/<cnc-arbiter-7.2.1>.ova
-y -s'
```

Deploy an arbiter VM on KVM

Provides the arbiter tar bundle requirement and a sample `virt-install` command for deploying an arbiter VM on KVM.

Deploy the arbiter VM on KVM by using the arbiter tar bundle.

Before you begin

Obtain the arbiter tar bundle.

For installation guidance, refer to Install Crosswork Network Controller cluster using CLI on page 178.

Procedure

Deploy the arbiter VM with the sample `virt-install` command. In this example, `cnc-arbiter-7.2.1` is the arbiter tar bundle.

Example

```
virt-install --boot uefi --boot hd,cdrom --connect qemu:///system --virt-type
 kvm --name arbiter
--ram 58982 --vcpus 8 --os-type linux --disk
path=cnc-arbiter-7.2.1_rootfs.qcow2,
format=qcow2,bus=scsi --disk path=cnc-arbiter-7.2.1_dockerfs.qcow2,
format=qcow2,bus=scsi --disk path=disk3,format=qcow2,bus=scsi --disk path=disk4,
format=qcow2,bus=scsi --disk path=cnc-arbiter-7.2.1_extrafs.qcow2,
format=qcow2,bus=scsi --disk path=disk6,format=qcow2,bus=scsi
--disk=arbiter.iso,device=cdrom,bus=scsi
--import --network bridge=intCWMgmt1,model=virtio --network
bridge=intData,model=virtio --noautoconsole
--os-variant ubuntu22.04 --graphics vnc,listen=0.0.0.0
```

20 Enable Geo Redundancy Solution

Topics:

- Geo redundancy inventory file guidelines
- Geo redundancy workflow (Day 0)
- Geo redundancy workflow (Day N)
- Update or add an application after enabling geo redundancy
- Geo redundancy scenarios
- Standby cluster behavior
- Arbiter VM behavior

Introduces the geo redundancy enablement workflows and related configuration guidance.

This chapter describes Day 0 and Day N workflows for enabling geo redundancy.

It also covers inventory file guidance, cluster settings, application updates, and cluster behavior.

Geo redundancy inventory file guidelines

You must ensure accuracy while preparing the inventory file (.yaml) for geo redundancy activation.

You must ensure accuracy while preparing the inventory file (.yaml) for geo redundancy activation.

You must ensure accuracy while preparing the inventory file (.yaml) for geo redundancy activation.

Here are some guidelines to assist you.

1. Add these fields to the inventory file for each cluster's unified endpoint:

```
data_vip: <>
data_vip_mask: <>
management_vip: <>
management_vip_mask: <>
```

2. When geo enabling the setup with `is_skip_peer_check_enabled` set to `true`, update each cluster's unified endpoint connectivity details based on the deployment model (IPv4, IPv6, or dual stack). For dual stack deployments, IPv6 addresses are given preference.

   If the cluster's unified endpoint type is set to **FQDN**, ensure that these fields are populated in the inventory file according to the setup configuration.

```
management_fqdn
data_fqdn
```

   If the cluster's unified endpoint type is set to **IP**, ensure that these fields are populated in the inventory file according to the setup configuration.

   Table 91: Unified endpoint connectivity details

Field	Supported deployment model
`data_vip_ipv4: <>`	IPv4, dual stack
`data_vip_ipv4_mask: <>`	IPv4, dual stack
`management_vip_ipv4: <>`	IPv4, dual stack
`management_vip_ipv4_mask: <>`	IPv4, dual stack
`data_vip_ipv6: <>`	IPv6, dual stack
`data_vip_ipv6_mask: <>`	IPv6, dual stack
`management_vip_ipv6:<>`	IPv6, dual stack
`management_vip_ipv6_mask: <>`	IPv6, dual stack

3. For **IPv4 deployment**, ensure the values of these fields match in the cluster's unified endpoint:
   - `data_vip` must match `data_vip_ipv4`
   - `data_vip_mask` must match `data_vip_ipv4_mask`
   - `management_vip` must match `management_vip_ipv4`
   - `management_vip_mask` must match `management_vip_ipv4_mask`

4. For **dual stack deployment**, ensure the values of these fields match in the cluster's unified endpoint:
   - `data_vip` must match `data_vip_ipv6`
   - `data_vip_mask` must match `data_vip_ipv6_mask`
   - `management_vip` must match `management_vip_ipv6`
   - `management_vip_mask` must match `management_vip_ipv6_mask`

Geo redundancy workflow (Day 0)

This topic explains the workflow for enabling geo redundancy on Day 0. It provides a high-level view of the tasks involved in installing and enabling geo redundancy in Crosswork Network Controller.

This topic explains the workflow for enabling geo redundancy on Day 0. It provides a high-level view of the tasks involved in installing and enabling geo redundancy in Crosswork Network Controller.

Geo redundancy workflow (day 0) for cluster deployments

> **Note**
>
> The recommended Day 0 setup for enabling geo redundancy is an empty Crosswork cluster (without any applications, devices or data gateways onboarded).

Here is the geo redundancy workflow when Crosswork Network Controller is deployed on a cluster.

Table 92: Geo redundancy workflow for cluster deployments (Day 0)

Step	Action
1. Install the active cluster in AZ1. 2. Install the standby cluster in AZ2.	Install using your preferred method: • *VMware vCenter:* : Install Cluster on VMware vCenter on page 81 • *KVM*: Install Crosswork Cluster on KVM on page 169 Verify if the installation was successful, and log into the Cisco Crosswork UI.

Step	Action
3. (Optional) Install the arbiter VM in AZ3. **Note** Please skip this step if you do not wish to use the auto arbitration functionality.	Follow the instructions in Deploy the Arbiter VM on page 435.
4. (Optional) Import the Crosswork inventory.	If you want to perform node-related activities (such as adding or removing a node) from the Crosswork UI, you must manually import a cluster inventory (.yaml) file into the Crosswork UI. For more information, refer to Import cluster inventory on page 122.
5. Create a backup of your Crosswork clusters.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide*. **Note** Importing the cross cluster inventory template cannot be undone if there is no pre-existing backup of the system before the template is loaded.
6. Perform the connectivity checks.	Follow the instructions in Perform connectivity checks on page 450 topic.

Step	Action
7. Prepare the cross cluster inventory file and enable geo redundancy. **Note** You can build the inventory YAML manually, or edit and use an existing YAML file. refer to Geo redundancy scenarios on page 477 for example scenarios that fit your requirement.	Follow the instructions in Enable geo redundancy on page 452 topic.
8. Install and enroll Crosswork Data Gateway, and onboard devices.	

Step	Action
	1. Choose the deployment profile for the Data Gateway VM. refer to VM requirements on page 44. **Note** If you are redeploying the same Data Gateway with Crosswork Network Controller, delete the previous Data Gateway entry from the **Virtual Machine** table under **Data Gateway Management**. For information on how to delete a Data Gateway VM, refer to Delete the Data Gateway VM from Cisco Crosswork on page 415. 2. Review the installation parameters to ensure that you have all the required information to install the Data Gateway. refer to Crosswork Data Gateway parameters and deployment scenarios on page 131. **Note** Use an FQDN, for example, geomanagement.cw.cisco, as the unified multi-cluster domain name. Ensure this FQDN is reachable from both clusters and points to the Active Crosswork VIP in the Geo-HA DNS server. If these conditions are not met, Data Gateway instance enrollment will fail. 3. Install the Data Gateway using your preferred method: • Install Crosswork Data Gateway using vCenter vSphere Client on page 144 • Install Crosswork Data Gateway via the OVF Tool on page 154 4. Verify the Data Gateway enrollment. refer to Crosswork Data Gateway authentication and enrollment on page 158. **Note** Use an FQDN such as geomanagement.cw.cisco that is reachable from both clusters and points to the Active Crosswork VIP in the Geo-HA DNS server; otherwise, the enrollment will fail.

Step	Action
	5. After the installation is complete, perform the postinstallation procedure. refer to Crosswork Data Gateway post-installation tasks on page 159. 6. Repeat steps 1 to 5 in the workflow to install Data Gateways on both standby and active sites. 7. After you verify that the Data Gateway VM has enrolled with Crosswork Network Controller, create a common Data Gateway pool for active and standby sites. For more information, refer to the *Create a pool in the geo redundancy-enabled sites* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. 8. Assign the Data Gateways to AZ2. For more information. refer to the *Assign Data Gateways to geo redundancy-enabled sites* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*. 9. Edit the pool to add the new Data Gateways. For more information, refer to the *Edit or delete a Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*.
9. Configure the cross cluster settings	Follow the instructions in Configure cross cluster settings on page 458.
10. Complete an on-demand sync operation successfully.	On the **Cross Cluster** window, select **Actions** > **Synchronize** to initiate the sync operation.

Step	Action
11. Install the Crosswork Applications on the active cluster 12. Install the Crosswork Applications on the standby cluster **Note** Applications on the standby site remain in a degraded state until the first sync completes.	Follow the instructions in Install Crosswork Network Controller applications on page 266 topic. Once geo redundancy is enabled, a **Geo Redundancy** tile is added to the **Application management** window. This tile is built-in and cannot be upgraded, uninstalled, or deactivated. **Warning** • Parallel installation of applications on the active and standby clusters should be avoided. Complete the installation on the active cluster before proceeding with the installation on the standby cluster. • Applications should not be installed during a periodic or on-demand sync operation. Ensure there is sufficient time for the installation to complete before initiating a sync, and verify that no sync operation is in progress before installing an application. It is recommended to temporarily disable periodic sync when installing applications.
13. Verify that the geo redundancy was successfully enabled on all the AZs.	Perform these checks: 1. In the Cross Cluster Health Status, ensure the operational state is **Connected**. 2. In the Cross Cluster Health Status, ensure that active cluster state is **Healthy**. 3. In the Cross Cluster Health Status, ensure that standby cluster state is **Healthy**. 4. In the Cross Cluster Health Status, ensure that arbiter VM state is **Healthy**. 5. In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. 6. Verify that the heartbeat count between the clusters is incrementing and that no failures are observed over a 30-minute period. 7. Confirm the completion of one successful sync between the clusters. For more information, refer to View cross cluster status on page 463 topic.

Geo redundancy workflow (day 0) for single VM deployments

Here is the geo redundancy workflow when Crosswork Network Controller is deployed on a single VM.

Table 93: Geo redundancy workflow for single VM deployments (Day 0)

Step	Action
1. Install the active VM in AZ1. 2. Install the standby VM in AZ2.	Install using the instructions in Install Crosswork Network Controller on a Single VM on page 331. Verify if the installation was successful, and log into the Cisco Crosswork UI.
3. (Optional) Install the arbiter VM in AZ3. **Note** Please skip this step if you do not wish to use the auto arbitration functionality.	Follow the instructions in Deploy the Arbiter VM on page 435.
4. (Optional) Import the Crosswork inventory.	If you want to perform node-related activities (such as adding or removing a node) from the Crosswork UI, you must manually import a cluster inventory (.yaml) file into the Crosswork UI. For more information, refer to Import cluster inventory on page 122.
5. Create a backup of your Crosswork VM.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide*. **Note** Importing the cross cluster inventory template cannot be undone if there is no pre-existing backup of the system before the template is loaded.
6. Perform the connectivity checks.	Follow the instructions in Perform connectivity checks on page 450 topic.

Step	Action
7. Prepare the cross cluster inventory file and enable geo redundancy. **Note** You can build the inventory YAML manually, or edit and use an existing YAML file. refer to Geo redundancy scenarios on page 477 for example scenarios that fit your requirement.	Follow the instructions in Enable geo redundancy on page 452 topic.
8. Configure the cross cluster settings	Follow the instructions in Configure cross cluster settings on page 458.
9. Complete an on-demand sync operation successfully.	On the **Cross Cluster** window, select **Actions** > **Synchronize** to initiate the sync operation.
10. Verify that the geo redundancy was successfully enabled on all the AZs.	Perform these checks: 1. In the Cross Cluster Health Status, ensure the operational state is **Connected**. 2. In the Cross Cluster Health Status, ensure that active VM state is **Healthy**. 3. In the Cross Cluster Health Status, ensure that standby VM state is **Healthy**. 4. In the Cross Cluster Health Status, ensure that arbiter VM state is **Healthy**. 5. In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. 6. Verify that the heartbeat count between the VMs is incrementing and that no failures are observed over a 30-minute period. 7. Confirm the completion of one successful sync between the VMs. For more information, refer to View cross cluster status on page 463 topic.

Perform connectivity checks

Lists connectivity checks to perform before enabling geo redundancy.

Perform connectivity checks.

Perform these connectivity checks before enabling geo redundancy:

Note

- Static routes are not required for cross cluster connectivity.
- Mesh connectivity is required between Crosswork Network Controller, Crosswork Data Gateway, NSO, and data interface components across the Availability Zones (AZ).
- L2/L3 connectivity is supported.

Procedure

1. Copy (using SCP) a file from Availability Zone 1 (AZ1) to Availability Zone 2 (AZ2), and from AZ2 to AZ1 in corresponding Crosswork VMs and Crosswork Orchestrator pods to ensure connectivity between both clusters.

   SCP connectivity.

```
# Perform these actions from AZ1 to AZ2 and AZ2 to AZ1

cw-admin@192-168-6-101-hybrid:~$ sudo su
[sudo] password for cw-admin:
root@192-168-6-101-hybrid:/home/cw-admin# kubectl exec -it -n=kube-system
robot-orch-76856487-562w6 -- bash
robot-orch-76856487-562w6:~# touch t.txt
robot-orch-76856487-562w6:~# scp t.txt
cw-admin@YOUR_PEER_CLUSTER_MGMT_VIP:/home/cw-admin/
(cw-admin@192.168.5.100) Password:
t.txt

robot-orch-76856487-562w6:~# scp t.txt
cw-admin@YOUR_PEER_CLUSTER_DATA_VIP:/home/cw-admin/
(cw-admin@192.168.5.100) Password:
t.txt
```

2. Test the DNS resolution on system wide DNS server.

   DNS connectivity.

   Note

   Ensure validation of IPv4 and/or IPv6 address mapping based on the deployment type (IPv4, IPv6, or dual stack).

```
### Internal Authortative  resolution

dig @your_dns_server_ip  your_name.cw.cisco

; <<>> DiG 9.10.6 <<>> @172.28.122.84 geomanagement.cw.cisco
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8167
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1
```

```
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;your_name.cw.cisco.        IN  A

;; ANSWER SECTION:
your_name.cw.cisco. 5   IN  A   192.168.6.100

;; Query time: 126 msec
;; SERVER: 172.28.122.84#53(172.28.122.84)
;; WHEN: Fri Jun 30 23:47:51 PDT 2023
;; MSG SIZE  rcvd: 67

### External forwarding and resolution

 dig @your_dns_server_ip  ntp.esl.cisco.com

; <<>> DiG 9.10.6 <<>> @172.28.122.84 ntp.esl.cisco.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 43986
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;ntp.esl.cisco.com.     IN  A

;; ANSWER SECTION:
ntp.esl.cisco.com.  1   IN  A   171.68.38.66

;; Query time: 311 msec
;; SERVER: 172.28.122.84#53(172.28.122.84)
;; WHEN: Fri Jun 30 23:46:37 PDT 2023
;; MSG SIZE  rcvd: 62
```

Verify if the DNS TTL in your VM is lesser than 60 seconds (< 60s).

```
cw-user@admin-M-C2EM ~ %  dig +nocmd +noall +answer @your_dns_server_ip
your_fqdn
geomanagement.cw.cisco. 60  IN  A   192.168.6.100

For ipv4 check

````dig +nocmd +noall +answer @your_dns_server_ip  your_fqdn ```

For ipv6 check

````dig +nocmd +noall +answer AAAA @your_dns_server_ip  your_fqdn ```
```

Enable geo redundancy

Explains how to build or edit the inventory YAML file and activate geo redundancy.

Enable geo redundancy.

Before you begin

- Ensure you have met all the requirements specified in Geo Redundancy Requirements on page 425.
- Ensure to provide relevant information for active, standby, and arbiter clusters and the unified cross cluster.

This topic explains the steps to build or edit the inventory YAML file and activate geo redundancy.

Geo redundancy can be configured at any time after the active cluster is built. The process does vary slightly if the standby cluster is not built and activated within 6 hours of the active cluster. These variations are clearly noted in this procedure.

Note

For dual-stack geo setups, the interface requires both IPv6 and IPv4 connectivity details, including the IPv4 subnet mask. The validation error *"Invalid IPv4 subnet mask format (must be between 0-32)"* appears when the IPv4 mask is missing or invalid, and this behavior is expected. Earlier releases allowed IPv6-only inputs in some cases, but in 7.2.x both IPv6 and IPv4 information are required to support inventory creation and DNS updates. The geo activation API still accepts IPv6-only inventory for dual-stack setups.

Note

Click on **How it works?** link to view a visual representation of how geo redundancy is enabled.

Procedure

1. Log in to the Crosswork cluster that will function as the active cluster. From the main menu, choose **Administration** > **Geo Redundancy Manager**.

   The **Geo Redundancy Manager** window is displayed.

2. Click **Get started with geo enablement**, and the **Geo Enablement** dialog box is displayed. Choose the desired option.
   - **Build inventory YAML manually**: Select this to create the inventory YAML file manually using the setup details. Click **Proceed** to open the **Cross cluster details** page. Go to step 3.
   - **Import and edit inventory YAML file**: Select this to use and modify an existing YAML file. Browse for the file on your machine and select it. Click **Proceed** to open the file on the **Validate** page. Go to step 7.

3. The **Geo Enablement** window appears, with the first step, **Cross cluster details**, highlighted. Provide relevant values for these fields:
   - **Cross cluster details**:
     - **Intended cluster role**: Choose **Active**.
     - **Cross cluster name**: Enter a relevant name without spaces.
   - **Crosscluster unified endpoint details**: Enter the unified FQDN to update the DNS entry for the active cluster.
     - **Management host name** and **Domain name**
     - **Data host name** and **Domain name**
   - **Global attributes**:

- **Secret string**: Enter a string to encrypt and decrypt the root certificate key exchanged between clusters. The string must be at least 10 characters long and include uppercase and lowercase letters, numbers, and at least one special character.
- **Post migration - Geo enabled**: Select if geo redundancy was enabled on the system after migration, rather than during a fresh deployment.
- **Peer cluster deployed**: Toggle to indicate if peer cluster was deployed during the initial setup (Day-0) or later (Day-N). For Day-N deployments, you must keep the cluster node details ready, and set the **is_skip_peer_check_enabled** flag in the inventory file to `true` to skip application and compute resource equivalency checks on the peer cluster.
- **Arbiter cluster deployed**: Toggle to indicate whether the arbiter cluster was added during initial setup (Day-0) or later (Day-N).

4. Click **Next**. The **Geo Enablement** window is refreshed, with **Active cluster details** highlighted. Provide relevant values for these fields:
   - **Cluster id**
   - **Cluster name**
   - **Management VIP IP** and **Subnet mask**
   - **Data VIP IP** and **Subnet mask**
   - **Site location**: Each cluster must have a unique location.
   - **HTTP Username** and **Password**
   - **SSH Username** and **Password**
5. Click **Next**. The **Geo Enablement** window is refreshed, with **Standby cluster details** highlighted.
   a) Enter the required values for the standby cluster, following the same guidelines used for the active cluster. Refer to the previous step for field descriptions.

   > **Note**
   >
   > If you set **Peer cluster deployed** slider to OFF in the previous step, you must enter the **Node details** manually.

   b) Click **Add node** and enter the relevant IP address and subnet mask for Management and Data networks.
6. (Optional) Click **Next**. The **Geo Enablement** window is refreshed, with **Arbiter cluster details** highlighted. Enter the required values for the standby cluster, following the same guidelines used for the active cluster. Refer to step 4 for field descriptions.

   The arbiter cluster page appears only if you set **Arbiter cluster deployed** to ON in step 4.
7. Click **Next**. The **Geo Enablement** window is refreshed, with **Validate** highlighted. Resolve any errors listed in the clusters.
8. Click **Next**. The **Geo Enablement** window is refreshed, with **Preview and activate** highlighted. Review the details you entered.
   a) (Optional) Click **Export inventory** to export the completed YAML file to your machine.
   b) Click **Activate geo cross cluster** to activate geo redundancy on the cluster.

   The progress can be viewed from the **Jobs** window, or by clicking the ⓘ icon.

c) (Optional) On the

9. Configure the standby and arbiter clusters.
   a) Log in to the Crosswork cluster that will function as the standby cluster.
   b) On the **Cross cluster details** page, select **Intended cluster role** as **Standby**.

      In case of a re-import: When you use the **Re-import** option to build the YAML file, the cluster role is hard-coded as "standby," and you must click **Fetch active cluster data** to proceed.

   c) The **Fetch Active Cluster Inventory** dialog box is displayed. Enter the login credentials for the active cluster to retrieve its inventory details.
      - Management VIP
      - Username
      - Password

   d) Click **Fetch** to collect the active cluster inventory details.

      After the standby cluster is configured, the **Job status** will be displayed as **Completed** on both clusters. Once the inventory upload is successfully completed on both clusters, the status will be updated in the **Geo Redundancy Manager** window.

      > **Note**
      >
      > Enable **Pairing mode** if the standby cluster is activated more than six hours after the active cluster. **Pairing mode** remains active for six hours from the time it is enabled, after which it automatically disables. You cannot disable it manually.

   e) (Optional) Log in to the Crosswork VM that will function as the arbiter VM.
   f) On the **Cross cluster details** page, click on **Fetch active cluster data**.
   g) The **Fetch Active Cluster Inventory** dialog box is displayed. Enter the login credentials for the active cluster to retrieve its inventory details.
      - Management VIP
      - Username
      - Password

   h) Click **Fetch** to collect the active cluster inventory details.

      After the arbiter VM is configured, the **Job status** will be displayed as **Completed** on all clusters. Once the inventory upload is successfully completed on all clusters, the status will be updated in the **Geo Redundancy Manager** window.

What to do next

To continue with the activation, refer to

View cross cluster status on page 463.

DNS traffic redirect binary requirements

Describes the requirements for developing a DNS traffic redirect binary for Geo HA arbitration.

Use these requirements when developing the binary that updates DNS records during Geo HA arbitration.

Scope

Use this topic when developing the DNS traffic redirect binary that you provide for Geo HA. In the Geo HA configuration, the UI refers to this binary as the traffic redirect script.

The requirements apply whenever Geo HA executes the configured binary. Implement these requirements before you provide the binary. You do not repeat these development steps before each Test traffic redirect or switchover. This topic describes binary compatibility requirements and the arbitration sandbox policy.

> **Note**
>
> Test traffic redirect is available only when an arbiter VM is deployed as part of the Geo HA setup.

Execution context

The binary runs in the arbitration sandbox under this execution context:

- The binary runs as a separate, nonprivileged, non-login system user with no supplementary groups.
- `PR_SET_NO_NEW_PRIVS` is set before execution. Neither the binary nor any child process that it starts can gain elevated privileges.
- The default seccomp action for a blocked system call is `EPERM` (operation not permitted). The binary receives an error return from the system call instead of being killed.

Execution image requirements

The arbitration execution image must include `python3` and `libseccomp`. If either package is missing, the arbitration sandbox cannot initialize.

Allowed operations

The sandbox allows these operations:

- Read-only file access. `open()` and `openat()` are allowed with `O_RDONLY`. The binary can read its own ELF, shared libraries, `/proc/self/*`, time-zone data, `/etc/resolv.conf`, `/etc/hosts`, and any file readable by user ID 5555.
- Read-only operations including `stat()`, `fstat()`, `access()`, `readlink()`, `getdents64()`, `lseek()`, and `pread64()`.
- Network operations including `socket`, `connect`, `bind`, `listen`, `accept`/`accept4`, `send*`/`recv*`, `shutdown`, `getsock*`/`setsock*`, `epoll_*`, `eventfd`/`eventfd2`, `poll`, and `select`. The binary can open connections to any host and port.
- System calls required by a standard Go binary on Alpine and musl, including `mmap`/`brk`/`mprotect`/`madvise`, `clone`/`clone3`/`futex`, signals (`rt_sig*`/`tgkill`), the `getpid`/`getuid` family, `clock_*`/`timerfd_*`, `getrandom`, `getcwd`, `set_robust_list`, `rseq`, `execve`, and `exit`/`exit_group`.

> **Note**
>
> Secrets intended for internal use cannot be read.

Blocked operations

The sandbox blocks these operations:

- File writes. `open()` and `openat()` with `O_WRONLY`, `O_RDWR`, or `O_CREAT` return `EPERM`. The binary cannot create or modify files.
- Writes to `/tmp`. The binary must not use `os.CreateTemp()`, `ioutil.TempFile()`, or another pattern that creates a file.
- File deletion. `unlinkat()` is not allowed and returns `EPERM`.
- `chmod()`, `chown()`, and `ptrace()`. These operations are not allowed and return `EPERM`.
- Credential-file access. Credential files are blocked at the filesystem layer because they are mode 0400 and owned by root.
- Privilege gain. `PR_SET_NO_NEW_PRIVS` prevents privilege gain.
- Any system call that is not listed as allowed. The default action is `EPERM`.

DNS update and input requirements

- Provide an independent, precompiled binary for the DNS traffic redirect operation.
- Perform DNS updates entirely over network communication using TCP, UDP, or HTTPS.
- Do not create, modify, or delete files during execution. This includes temporary files and intermediate TSIG key, zone, or `nsupdate` input files.
- Keep key material and configuration in memory. Do not read TLS keys or other credentials from the pod filesystem.
- Accept runtime parameters as JSON on the command line.
- The seccomp filter is inherited by every child process that the binary starts.
- Print a clear error line to `stderr` before exiting with a nonzero code. The error text is included verbatim in the alarm raised to the operator.

Local testing

Use the Test traffic redirect action in the Cross Cluster UI to test the binary.

Exit codes

Use the exit codes in this table to identify the result reported by the binary or sandbox.

Table 94: DNS traffic redirect binary exit codes

Exit code	Meaning
`0`	Success.

Exit code	Meaning
`1`	The binary reported an application error. Check `stderr` for details.
`2`	Sandbox or interpreter initialization failure caused by image configuration.
`127`	The sandbox is missing from the container image.
`225`	A seccomp violation occurred in kill mode because the binary attempted a blocked system call.
`137`	The process was killed by SIGKILL or an out-of-memory condition.
`143`	The process was killed by SIGTERM.

Compatibility quick reference

Use this table when reviewing the binary implementation.

Table 95: Sandbox operation compatibility

Operation	Compatibility
Read files with `O_RDONLY`	Allowed.
Create, modify, or delete files	Not supported.
Use temporary or intermediate files	Not supported.
TCP or UDP network operations	Allowed.
Read credentials from `/etc/*`	Blocked by filesystem permissions.
`ptrace`, `chmod`, or `chown`	Blocked.
Privilege gain	Blocked by `PR_SET_NO_NEW_PRIVS`.

Configure cross cluster settings

Describes the traffic redirect script requirements for auto-arbitration.

Configure cross cluster settings.

Before you begin

- If you choose to enable auto-arbitration, you must have the traffic redirect script ready for upload. The redirect script must be an independent, pre-compiled binary that does not rely on external libraries. For requirements that apply when developing the traffic redirect binary, refer to DNS traffic redirect binary requirements on page 456.
- You must log in as an admin user to configure cross cluster settings. Only admin users have the necessary permissions to make these changes. Users with a local role, even if they have read, write, and delete permissions, cannot configure cross cluster settings.

Configuring cross cluster settings is important to ensure secure data transfer between clusters, to facilitate reliable backups and recovery, and data compliance.

This topic explains how to configure the cross cluster settings.

> **Note**
>
> The default values shown in the Cross Cluster Configuration UI are the recommended settings.

Procedure

1. From the main menu, choose **Administration** > **Cross Cluster**. The **Cross Cluster** window is displayed. Click on the **Configurations** tab.
2. The **Configurations** window is displayed, with the first step, **1 – Storage settings**, highlighted. Fill all the fields provided for the SCP Host server.

Note

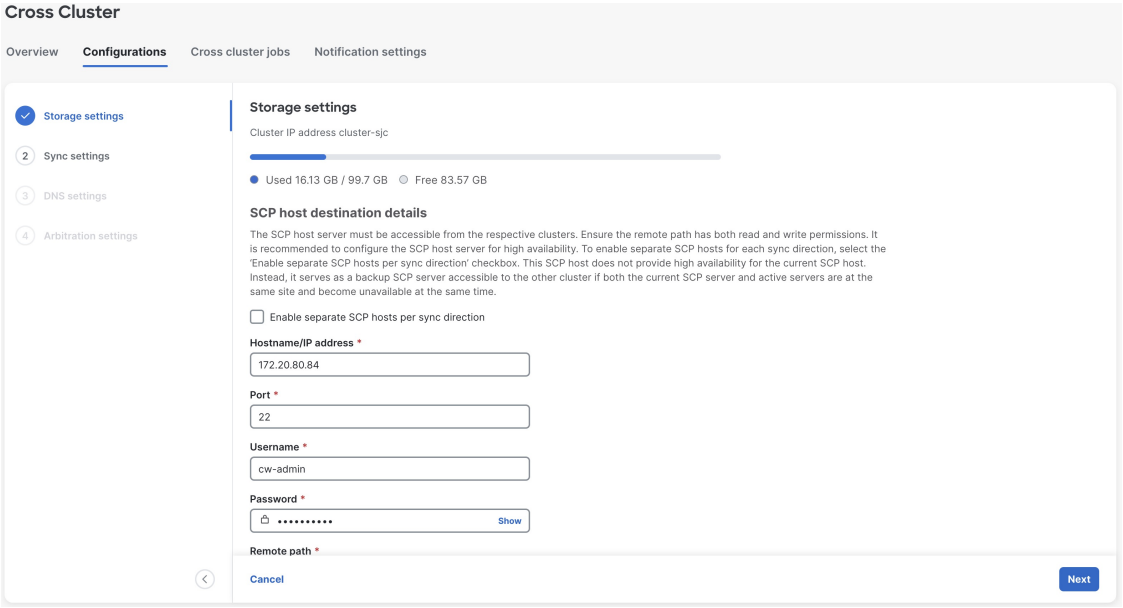
Select the **Additional SCP host** checkbox only if the standby cluster requires its own local SCP server for storage. This option does not provide high availability (HA) between SCP hosts. Each SCP host operates independently and must maintain its own HA setup. If the primary SCP host becomes unavailable, geo-backup and geo-sync operations will fail, even if an additional SCP host exists.

- When the primary SCP host is added during the initial Cross Cluster setup, it must be reachable from both the active and standby clusters. If both can access it, the primary SCP host becomes available to both clusters.
- Adding an additional SCP host means it is local to the standby cluster, while the primary SCP host is local to the active cluster. The primary and additional SCP hosts operate independently and do not provide high availability for each other.
- During a sync operation (from the active cluster to the standby cluster), the active cluster creates the geo-backup and stores it on the primary SCP host. During a geo-restore, the standby cluster retrieves the backup from the same primary SCP host. In short, the active cluster writes the backup to its local SCP host, and the standby cluster restores from that same host.
- If the primary SCP host becomes unavailable, geo-backup and geo-sync operations will fail. The additional SCP host does not provide HA for the primary SCP host. It only serves as an alternate SCP host for the standby cluster that can be used if both the active cluster and its primary SCP host are unavailable (for example, during a site failure).
- Each SCP host must maintain its own HA setup independently. Adding an additional SCP host allows the peer (standby) cluster to use its local SCP host for storage, but it does not create HA between SCP hosts or prevent geo-sync failure if the primary SCP host is down.

Note

After a SCP host is configured, you can view the used and free space available in the server.

Figure 60: Storage settings

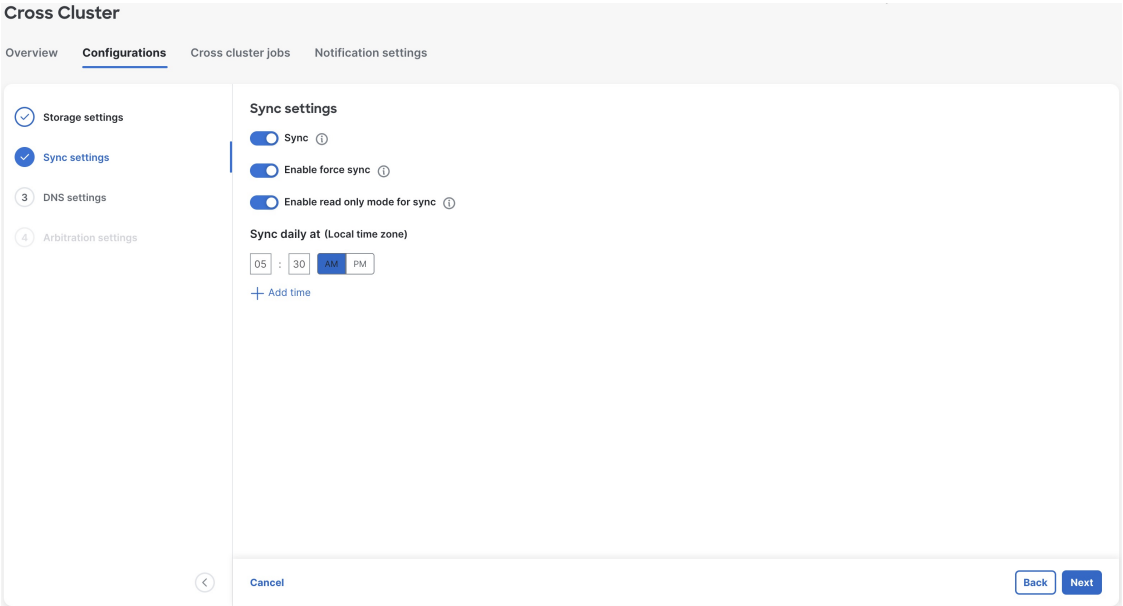


3. Click **Next**. The **Configurations** window is displayed, with the next step, **2 - Sync settings**, highlighted. Data synchronizarion ensures high availability, consistency, load balancing, and data compliance between geo redundant clusters.

   Enable the **Sync** slider button to set an auto-sync schedule, and set the sync times. Optionally, you can enable **Enable force sync** and **Enable read only mode for sync** as per your requirement. Click the tooltip icon next to each option to learn more.

> **Note**
>
> It is recommended to sync at least once every 8 hours.

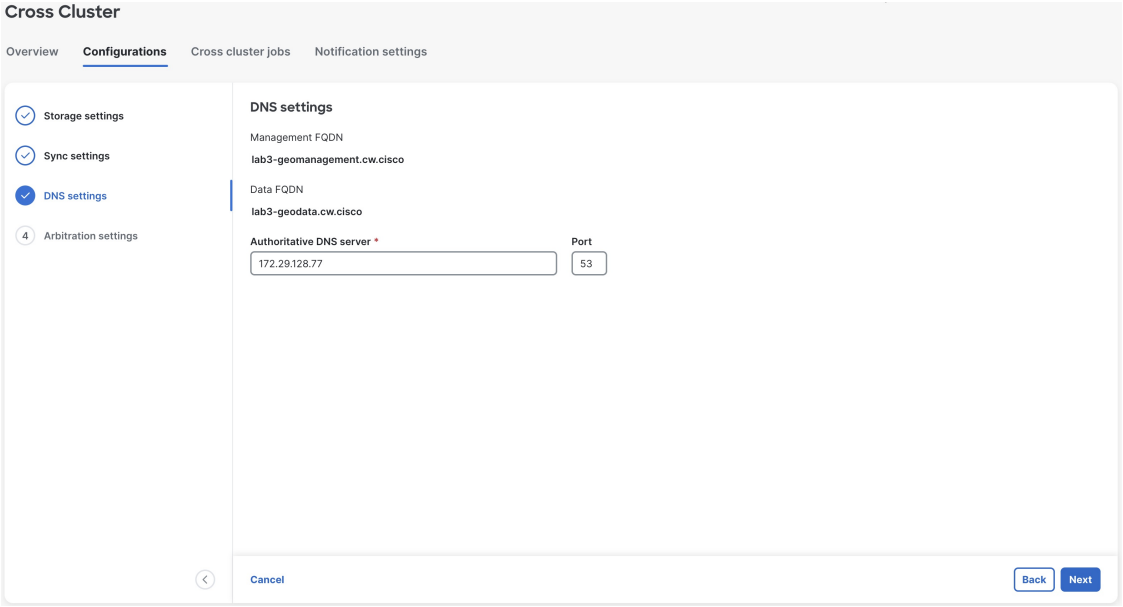
Figure 61: Sync settings



4. Click **Next**. The **Configurations** window is displayed, with the next step, **3 - DNS settings**, highlighted. Add the details for the **Authoritative DNS server** and **Port**.

> **Note**
>
> - The DNS server should be configured with the same management FQDN and data FQDN displayed on the UI.
> - The DNS record TTL for FQDN must be lesser than 60 seconds (< 60s).

Figure 62: DNS settings

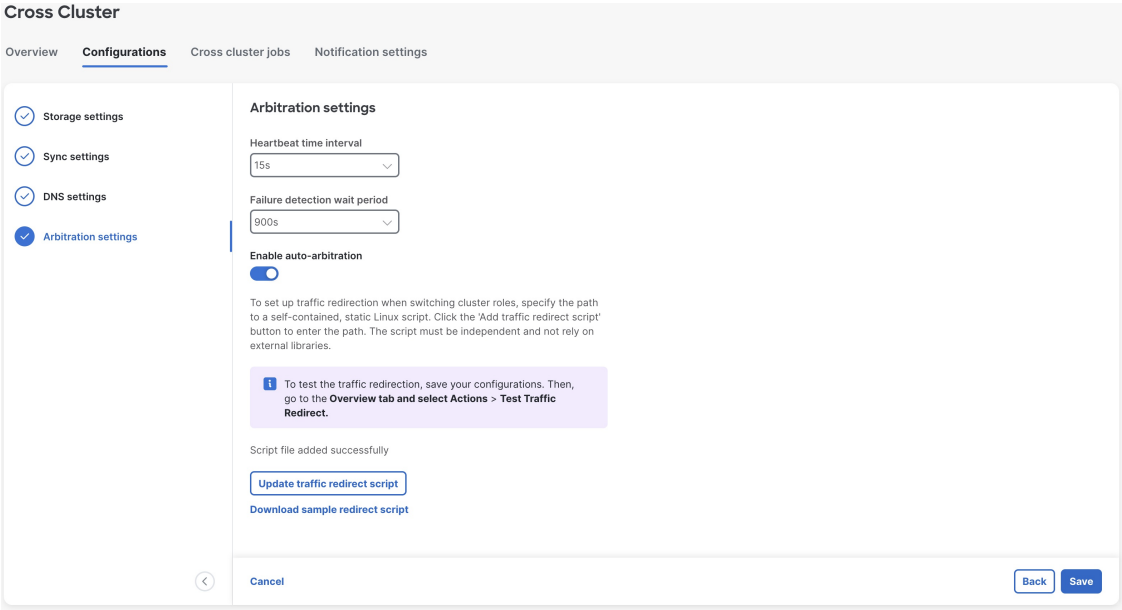


5. Click **Next**. The **Configurations** window is displayed, with the final step, **4 - Arbitration settings**, highlighted. Set relevant values for the **Heartbeat time interval** and **Failure detection wait period** fields.
6. Activate the **Enable auto-arbitration** slider (recommended) to start the leader election process and update the leader information. Click **Add traffic redirect script** to add the script used to update the DNS entry for the specified FQDN.

> **Note**
>
> The **Enable auto-arbitration** slider is available only if the arbiter VM was deployed as part of the geo HA setup.

In the **Add file** dialog box, choose either **URL** or **SCP** as your preferred protocol. Based on your selection, fill in the additional fields with the required information. Click **Add** to proceed.

Figure 63: Traffic redirect script added



7. Click **Save** to save the changes. After saving, the changes replicate to the other clusters in the geo HA setup. You can apply these settings from any cluster in the geo HA setup.

Configure cross cluster notification settings

This topic explains how to configure the cross cluster notification settings.

Configure cross cluster notification settings.

This topic explains how to configure the cross cluster notification settings.

Procedure

1. From the main menu, choose **Administration** > **Cross Cluster**. The **Cross Cluster** window is displayed.
2. Click on the **Notification settings** tab, and **Add notification destination** window is displayed.
3. Click **Add Criteria**, and you will be navigated to the **Create Notification Policies** window.

   Enter relevant values for these fields, and save the policy.

   - Policy name
   - Criteria
   - Notification destination

   For more information, refer to *Create Notification Policy for System Events* in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*

View cross cluster status

This topic explains how to view the cross cluster status after successfully enabling geo redundancy.

View cross cluster status.

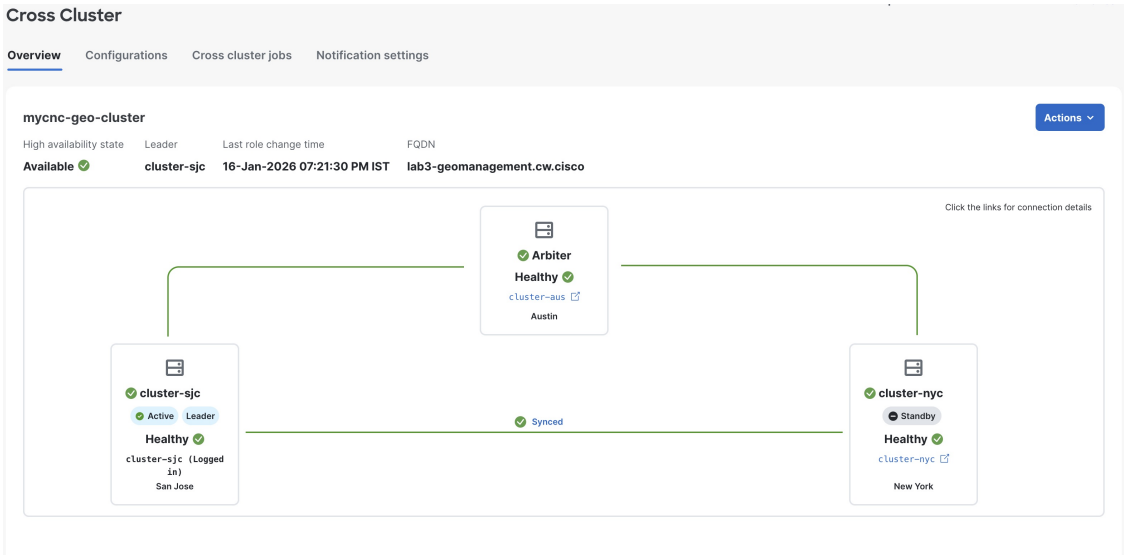
This topic explains how to view the cross cluster status after successfully enabling geo redundancy.

Procedure

1. From the main menu, choose **Administration** > **Cross Cluster**. The **Cross Cluster** window is displayed.

   The cross cluster health status is displayed along with the cluster leader status, high availability state, heartbeats round trip time, failed heartbeats, and last active cluster change time. You can also view the status of the active and standby clusters along with the operational state and last sync status.

   Figure 64: Cross cluster with active, standby, and arbiter AZs



> **Note**
>
> By design, the Arbiter name is not displayed on the Cross Cluster page. The Arbiter VM is consistently shown as "Arbiter," regardless of the value of the *cluster-id* parameter.

Scroll further down to refer to the data store replication states. The **Data stores** table will show all the different data stores that have been replicated along with the lag info for them.

- Each **Data store type** has a corresponding **Sync type**. Postgres and Timescale data stores support live synchronization, while gluster is updated via periodic sync.

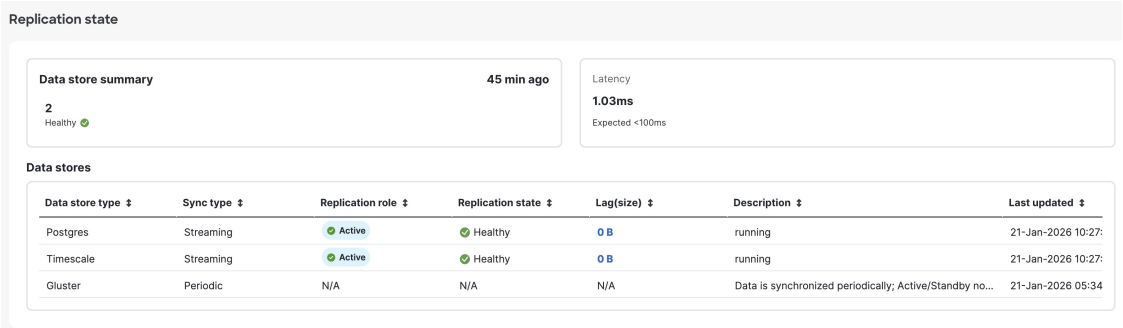
  Table 96: Data store types

Data store type	Sync type
Postgres	Streaming
Timescale	Streaming
Gluster	Periodic

- The **Replication role** displays the role of the cluster. For example, the value for an active cluster will be *Active*.
- The **Replication state** displays the status of the data store.

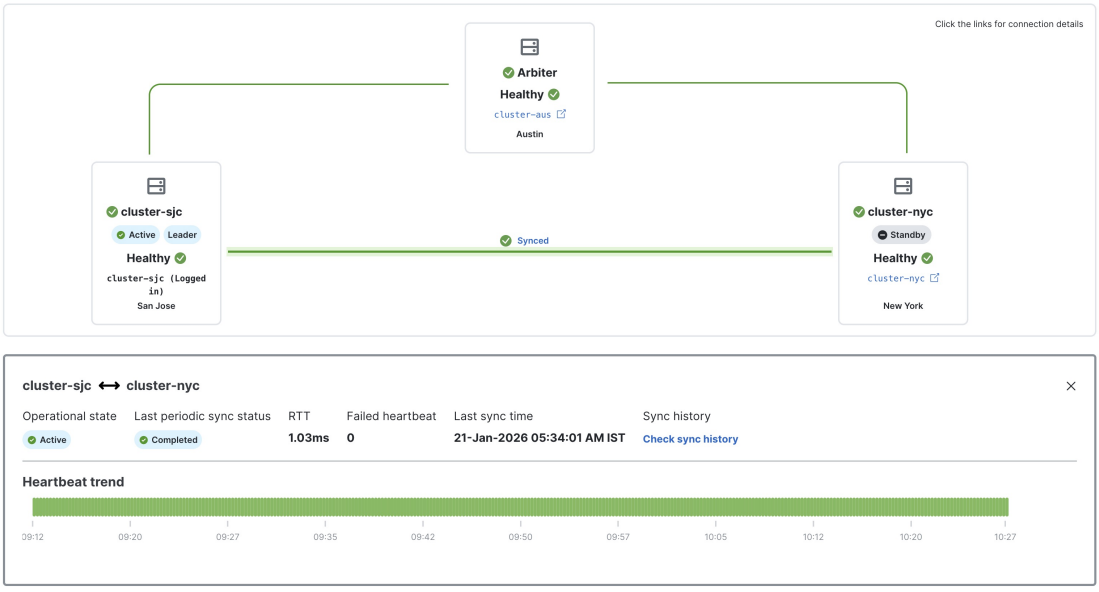
- The **Lag(size)** value indicates the lag between the active and standby clusters.

Figure 65: Cross Cluster window - Replication state



2. You can perform these operations on the **Cross Cluster** window.
   a) Click the link between clusters to view a visual representation of the heartbeat trend.

Figure 66: Cross Cluster window - heartbeat trend



   b) Click on the name of the active and standby clusters to view the cluster details.
   c) Click on **Lag(size)** for a data store to view a detailed graph of the replication summary.

3. You can select and perform these optional operations from the **Actions** drop-down menu.
   a) Click **Actions** > **Set cluster role** to override the switchover process and assign the cluster roles manually. For more information, refer to Perform switchover manually on page 515.
   b) Click **Actions** > **Showtech request**, and the **Showtech Request** pop-up window is displayed. Enter the relevant SCP host details and click **Export** to download the showtech logs.
   c) Click **Actions** > **Synchronize** to initiate an on-demand sync operation.

> **Note**
>
> Do not perform the **Synchronize** operation before completing the other configurations (such as storage, DNS, and sync settings). Once a sync is initiated, it cannot be stopped midway.

d) Click **Actions** > **Repair system** to address any sync issues that require the system to kickstart the replication or bootstrap the standby cluster from the active side. This operation will attempt to automatically repair the applications and database.

e) Click **Actions** > **Test traffic redirect** to execute the traffic redirect script. This operation creates a job to validate the arbitration executable. The redirect script is used to update the DNS entry for the FQDN provided.

This option is available only if the arbiter VM was deployed as part of the geo HA setup.

f) Click **Actions** > **Switchover** to initiate the switchover process. This is a one-click control that performs the three steps of the switchover process. It automates setting the two roles and updating the DNS in a single function. For more information, refer to Perform switchover manually on page 515.

Geo redundancy workflow (Day N)

This topic outlines the high-level workflow for the tasks required to enable geo redundancy on Day N on Crosswork Network Controller operating in a standalone site.

This topic outlines the high-level workflow for the tasks required to enable geo redundancy on Day N on Crosswork Network Controller operating in a standalone site.

A Crosswork site is considered as a"Day N" scenario if it has configured the below options:

- AZ1 is deployed.
- Application are installed.
- Devices are configured.
- Crosswork Data Gateway for AZ1 is enrolled.
- Providers are added.
- Collection jobs are running.

A backend check is conducted to verify the site's eligibility for enabling geo redundancy. If any checks fail, geo redundancy cannot be enabled, and alarms will be generated to notify you.

Once geo redundancy is enabled on Day N, please perform a sync from the Active to the Standby site before installing any applications.

Geo redundancy workflow (day N) for cluster deployments

Here is the geo redundancy workflow when Crosswork Network Controller is deployed on a cluster.

Table 97: Geo redundancy workflow (Day N)

Step	Action
1. Install the standby cluster in AZ2.	Install using your preferred method: • *VMware vCenter:* : Install Cluster on VMware vCenter on page 81 • *KVM*: Install Crosswork Cluster on KVM on page 169 Verify if the installation was successful, and log into the Cisco Crosswork UI.
(Optional) 2. Install the arbiter VM in AZ3. **Note** Please skip this step if you do not wish to use the auto arbitration functionality.	Follow the instructions in Deploy the Arbiter VM on page 435.
3. (Optional) Import the Crosswork inventory.	If you want to perform node-related activities (such as adding or removing a node) from the Crosswork UI, you must manually import a cluster inventory (.yaml) file into the Crosswork UI. For more information, refer to Import cluster inventory on page 122.
4. Create a backup of your Crosswork clusters.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide* **Note** Importing the cross cluster inventory template cannot be undone if there is no pre-existing backup of the system before the template is loaded.
5. Perform the connectivity checks.	Follow the instructions in Perform connectivity checks on page 450 topic.

Step	Action
6. Prepare the cross cluster inventory file and enable geo redundancy. **Note** You can build the inventory YAML manually, or edit and use an existing YAML file. refer to Geo redundancy scenarios on page 477 for example scenarios that fit your requirement.	Follow the instructions in Enable geo redundancy on page 452 topic.
7. Install and enroll Crosswork Data Gateway, and onboard devices.	

Step	Action
	You can choose one of these approaches for AZ2 setup: • Deploy new Data Gateways by following the workflow in this table. • Assign existing Data Gateways to AZ2. For more information. refer to the *Assign Data Gateways to geo redundancy-enabled sites* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* Follow this workflow to deploy a new Data Gateway. **1.** Choose the deployment profile for the Data Gateway VM. refer to VM requirements on page 44. **Note** If you are redeploying the same Data Gateway with Crosswork Network Controller, delete the previous Data Gateway entry from the **Virtual Machine** table under **Data Gateway Management**. For information on how to delete a Data Gateway VM, refer to Delete the Data Gateway VM from Cisco Crosswork on page 415. **2.** Review the installation parameters to ensure that you have all the required information to install the Data Gateway. refer to Crosswork Data Gateway parameters and deployment scenarios on page 131. **Note** Use an FQDN, for example, geomanagement.cw.cisco, as the unified multi-cluster domain name. Ensure this FQDN is reachable from both clusters and points to the Active Crosswork VIP in the Geo-HA DNS server. If these conditions are not met, Data Gateway instance enrollment will fail. **3.** Install the Data Gateway using your preferred method: • Install Crosswork Data Gateway using vCenter vSphere Client on page 144 • Install Crosswork Data Gateway via the OVF Tool on page 154 **4.** Verify the Data Gateway enrollment. refer to Crosswork

Step	Action
	Data Gateway authentication and enrollment on page 158. **Note** Use an FQDN such as geomanagement.cw.cisco that is reachable from both clusters and points to the Active Crosswork VIP in the Geo-HA DNS server; otherwise, the enrollment will fail. 5. After the installation is complete, perform the postinstallation procedure. refer to Crosswork Data Gateway post-installation tasks on page 159. 6. Repeat steps 1 to 5 in the workflow to install Data Gateways on both standby sites. 7. Assign the Data Gateways to standby site. For more information. refer to the *Assign Data Gateways to geo redundancy-enabled sites* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* 8. Add the new Data Gateways to the existing pool. For more information, refer to the *Edit or delete a Data Gateway pool* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide*
8. Configure the cross cluster settings.	Follow the instructions in topics below: • Cross Cluster settings: Configure cross cluster settings on page 458 • (Optional) Notification settings: Configure cross cluster notification settings on page 463
9. Complete an on-demand sync operation successfully.	On the **Cross Cluster** window, select **Actions** > **Synchronize** to initiate the sync operation.

Step	Action
10. Install the Crosswork Applications on the active cluster. 11. Install the Crosswork Applications on the standby cluster. **Note** Applications on the standby site remain in a degraded state until the first sync completes.	Follow the instructions in Install Crosswork Network Controller applications on page 266 topic. Once geo redundancy is enabled, a **Geo Redundancy** tile is added to the **Application management** window. This tile is built-in and cannot be upgraded, uninstalled, or deactivated. **Warning** • Parallel installation of applications on the active and standby clusters should be avoided. Complete the installation on the active cluster before proceeding with the installation on the standby cluster. • Applications should not be installed during a periodic or on-demand sync operation. Ensure there is sufficient time for the installation to complete before initiating a sync, and verify that no sync operation is in progress before installing an application. It is recommended to temporarily disable periodic sync when installing applications.
12. Verify that the geo redundancy was successfully enabled on the active and standby clusters.	Perform these checks: 1. In the Cross Cluster Health Status, ensure the operational state is **Connected**. 2. In the Cross Cluster Health Status, ensure that active cluster state is **Healthy**. 3. In the Cross Cluster Health Status, ensure that standby cluster state is **Healthy**. 4. In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. 5. Verify that the heartbeat count between the clusters is incrementing and that no failures are observed over a 30-minute period. 6. Confirm the completion of one successful sync between the clusters. For more information, refer to View cross cluster status on page 463 topic.

Onboard the arbiter VM

Step	Action
13. Update the cross cluster inventory file, and enable geo redundancy on the arbiter VM.	Follow the instructions in Enable geo redundancy on page 452 topic. **Note** refer to Geo redundancy scenarios on page 477 for relevant example scenarios that fit your requirement. Perform these steps: **1.** Import the active cluster inventory. **2.** Import the standby cluster inventory. **3.** Enable geo redundancy on the arbiter VM.
14. Configure the cross cluster settings	Follow the instructions in topics below: • Cross Cluster settings: Configure cross cluster settings on page 458 • (Optional) Notification settings: Configure cross cluster notification settings on page 463
15. Verify that the geo redundancy was successfully enabled on all the AZs. **Note** During a Day-N arbiter reimport on the active cluster, the Cross Cluster page may temporarily show the link to the standby cluster as disconnected and the standby state as unknown. This is expected behavior. Arbiter reimport is a multi-step process that involves component restarts, and the cluster states remain in transition during this period. The system converges automatically, and there is no functional impact.	Perform these checks: **1.** In the Cross Cluster Health Status, ensure the operational state is **Connected**. **2.** In the Cross Cluster Health Status, ensure that active cluster state is **Healthy**. **3.** In the Cross Cluster Health Status, ensure that standby cluster state is **Healthy**. **4.** In the Cross Cluster Health Status, ensure that arbiter VM state is **Healthy**. **5.** In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. **6.** Verify that the heartbeat count between the clusters is incrementing and that no failures are observed over a 30-minute period. **7.** Confirm the completion of one successful sync between the clusters. For more information, refer to View cross cluster status on page 463 topic.

Geo redundancy workflow (day N) for single VM deployments

Here is the geo redundancy workflow when Crosswork Network Controller is deployed on a single VM.

Table 98: Geo redundancy workflow (Day N)

Step	Action
1. Install the standby cluster in AZ2.	Install using the instructions in Install Crosswork Network Controller on a Single VM on page 331. Verify if the installation was successful, and log into the Cisco Crosswork UI.
(Optional) 2. Install the arbiter VM in AZ3. **Note** Please skip this step if you do not wish to use the auto arbitration functionality.	Follow the instructions in Deploy the Arbiter VM on page 435.
3. (Optional) Import the Crosswork inventory.	If you want to perform node-related activities (such as adding or removing a node) from the Crosswork UI, you must manually import a cluster inventory (.yaml) file into the Crosswork UI. For more information, refer to Import cluster inventory on page 122.
4. Create a backup of your Crosswork clusters.	Follow the instructions in *Manage Backups* chapter in *Cisco Crosswork Network Controller 7.2.x Administration Guide* **Note** Importing the cross cluster inventory template cannot be undone if there is no pre-existing backup of the system before the template is loaded.
5. Perform the connectivity checks.	Follow the instructions in Perform connectivity checks on page 450 topic.

Step	Action
6. Prepare the cross cluster inventory file and enable geo redundancy. **Note** You can build the inventory YAML manually, or edit and use an existing YAML file. refer to Geo redundancy scenarios on page 477 for example scenarios that fit your requirement.	Follow the instructions in Enable geo redundancy on page 452 topic.
7. Configure the cross cluster settings.	Follow the instructions in topics below: • Cross Cluster settings: Configure cross cluster settings on page 458 • (Optional) Notification settings: Configure cross cluster notification settings on page 463
8. Complete an on-demand sync operation successfully.	On the **Cross Cluster** window, select **Actions** > **Synchronize** to initiate the sync operation.
9. Verify that the geo redundancy was successfully enabled on the active and standby clusters.	Perform these checks: 1. In the Cross Cluster Health Status, ensure the operational state is **Connected**. 2. In the Cross Cluster Health Status, ensure that active VM state is **Healthy**. 3. In the Cross Cluster Health Status, ensure that standby VM state is **Healthy**. 4. In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. 5. Verify that the heartbeat count between the VMs is incrementing and that no failures are observed over a 30-minute period. 6. Confirm the completion of one successful sync between the VMs. For more information, refer to View cross cluster status on page 463 topic.

Onboard the arbiter VM

Step	Action
10. Update the cross cluster inventory file, and enable geo redundancy on the arbiter VM.	Follow the instructions in Enable geo redundancy on page 452 topic. **Note** refer to Geo redundancy scenarios on page 477 for relevant example scenarios that fit your requirement. Perform these steps: **1.** Import the active VM inventory. **2.** Import the standby VM inventory. **3.** Enable geo redundancy on the arbiter VM.
11. Configure the cross cluster settings	Follow the instructions in topics below: • Cross Cluster settings: Configure cross cluster settings on page 458 • (Optional) Notification settings: Configure cross cluster notification settings on page 463

Step	Action
12. Verify that the geo redundancy was successfully enabled on all the AZs. **Note** During a Day-N arbiter reimport on the active cluster, the Cross Cluster page may temporarily show the link to the standby VM as disconnected and the standby state as unknown. This is expected behavior. Arbiter reimport is a multi-step process that involves component restarts, and the VM states remain in transition during this period. The system converges automatically, and there is no functional impact.	Perform these checks: 1. In the Cross Cluster Health Status, ensure the operational state is **Connected**. 2. In the Cross Cluster Health Status, ensure that active VM state is **Healthy**. 3. In the Cross Cluster Health Status, ensure that standby VM state is **Healthy**. 4. In the Cross Cluster Health Status, ensure that arbiter VM state is **Healthy**. 5. In the Cross Cluster Health Status, ensure the High Availability state is **AVAILABLE**. 6. Verify that the heartbeat count between the VMs is incrementing and that no failures are observed over a 30-minute period. 7. Confirm the completion of one successful sync between the VMs. For more information, refer to View cross cluster status on page 463 topic.

Update or add an application after enabling geo redundancy

This topic describes the procedure that you must perform to update the application after enabling geo redundancy.

Update or add an application after enabling geo redundancy.

This topic describes the procedure that you must perform to update the application after enabling geo redundancy.

Procedure

1. 1. Complete an on-demand sync operation successfully.

   On the **Cross Cluster** window, select **Actions > Synchronize** to initiate the sync operation.

2. 2. Disable the sync operation.

   From the **Sync settings** window, drag the **Sync** slider button to disable the sync operation.

3. 3. Install the application updates on the Active cluster.

   Follow the instructions in the Install Crosswork Applications topic.

   Once geo redundancy is enabled, a Geo Redundancy tile is added to the Application management window. This tile is built in and cannot be upgraded, uninstalled, or deactivated.

Warning

- Parallel installation of applications on the active and standby clusters should be avoided. Complete the installation on the active cluster before proceeding with the installation on the standby cluster.
- Applications should not be installed during a periodic or on-demand sync operation. Ensure that there is sufficient time for the installation to complete before initiating a sync operation, and verify that no sync operation is in progress before installing an application. We recommend to temporarily disable periodic sync when installing applications.

4. 4. Install the application updates on the Standby cluster.
5. 5. Enable the sync operation in Crosscluster

   From the **Sync settings** window, enable the **Sync** slider button to set an auto-sync schedule, and set the sync times.

6. 6. Perform the sync operation

   After geo redundancy is enabled on both the Active and Standby clusters, update the sync settings and perform the first sync either manually or allow it to occur at the scheduled time. Any further application files should be installed only after the first sync is completed.

Geo redundancy scenarios

There are many scenarios with expected system behaviors that you should be aware of when geo redundancy is enabled.

There are many scenarios with expected system behaviors that you should be aware of when geo redundancy is enabled.

Application installation
Table 99: Application installation scenarios

Scenario	Expected system behavior
Application, version or cluster compute resource (such as nodes, CPU, memory. and disk) mismatch between active and standby clusters prior to enabling geo redundancy.	An equivalency check done prior to the geo redundancy enablement will identify any mismatch between the active and standby clusters (in terms of applications or versions), and prevent enablement. To proceed, please ensure that applications and versions match on both clusters.
Application or version mismatch between active and standby cluster after enabling geo redundancy.	Any configured sync operation will fail until the mismatch is corrected.
Installing an application or patch while a sync is in progress.	A sync operation can be configured as a periodic event or initiated on demand. While a sync operation is in progress, application installation will not be allowed.
Installing an application or patch when sync is not happening.	When sync is not happening, application installation is allowed.

Backup and restore

Table 100: Backup and restore scenarios

Scenario	Expected system behavior
Taking a backup on the active crosswork cluster.	This operation is allowed. Please take backup of the active cluster to have a point-in-time backup to roll back to in case the data sync is corrupted.
Taking a data only backup on the standby crosswork cluster.	This operation is not permitted.
Perform restore operation on the standby cluster.	This operation is not permitted.
Perform restore operation on the active cluster.	Restore operation is not allowed when auto-arbitration is enabled. You must first disable the auto-arbitration before performing the restore operation. If you want to restore a previous backup, perform the restore on the active cluster. The standby cluster will sync on the next sync cadence.
Configure the remote VM destination for backup and restore.	The remote destination for backup and restore must be re-configured after the switchover, as the destination settings are specific to the cluster.

Password update

When changing the SSH or HTTPS password on a geo-redundant deployment:

1. Update the password on the active cluster.
2. Wait for the synchronization operation to complete so that the password update is propagated to the standby cluster.
3. Verify that the active and standby clusters use the same SSH and HTTPS passwords. Do not configure a different password on the peer cluster.
4. Update the inventory YAML file with the new credential values and re-import it if required.
5. Verify that the Cross Cluster Health Status is **Connected**, both clusters are **Healthy**, and no SSH login test failures are reported.

Standby cluster behavior

Lists the accessibility of infrastructure and application features in the standby cluster.

The standby cluster has these accessibility characteristics for infrastructure and application features.

Accessibility of infrastructure and application features in the standby cluster is as listed here:

Note

Read operations may be unavailable during synchronization.

Table 101: Standby cluster behavior

Service/Feature	Behavior	Notes
Dashboard	Not supported	Menu option is disabled.
Platform (Crosswork Manager)	APIs are read-only, except where noted.	Application Management, Node Lifecycle, and Showtech support both read and write operations.
Platform (System management)	Read and write operations are supported.	
Geo redundancy manager	Read and write operations are supported.	
Cross Cluster	Read and write operations are supported.	
Alarms and Events	Read and write operations are supported for system alarms/events.	Only system alarms/events are displayed. Network and Device alarms are not listed.
DG manager	Only read operations are supported.	Data Collector Global Settings UI is accessible, but Data destinations UI is disabled.
Notification Policies	Not supported	Menu option is disabled.
Backup and Restore	Not supported	Menu option is disabled.
Audit Log	Not supported	Menu option is disabled.
Certificate Management	Supported	
Users and Roles	Supported (users, roles, and active sessions)	Synchronization supported from Active to Standby
System settings	Not supported	Menu option is disabled.
AAA	Login/logout operations only.	Menu option is disabled. AAA UI is disabled.
Network inventory	Not supported	Menu option is disabled.
Collection	Not supported	Menu option is disabled.
Topology	Not supported	Menu option is disabled.
Grouping (Port group, Devic group)	Not supported	Menu option is disabled.

Service/Feature	Behavior	Notes
Device management ZTP	Not supported	Menu option is disabled.
CLMS	Not supported	Menu option is disabled.

Geo HA synchronization behavior

In a Geo HA deployment, the infrastructure synchronizes data written to the shared persistence layers across the active and standby clusters. Specifically, data written to **Gluster**, **PostgreSQL**, and **TimescaleDB** is synchronized between the two sites. Whether a configuration is reflected immediately on the standby cluster depends on how the feature stores and applies its data.

Note

Some configuration changes require Geo HA synchronization before they are reflected or usable on the standby cluster. For example:

- Changes to the **Pre-login disclaimer** are reflected on the standby cluster only after **Geo HA synchronization** is performed.
- A **newly created role** is usable on the standby cluster only after **Geo HA synchronization** is performed.
- **Cluster configuration** changes are reflected across clusters after you save them and do not require a separate Geo HA synchronization step.
- There is no difference in synchronization behavior between **manual/on-demand synchronization** and **periodic/automatic synchronization**. The difference is only when synchronization occurs.

Arbiter VM behavior

Lists the accessibility of infrastructure and application features in the arbiter VM.

The arbiter VM has these accessibility characteristics for infrastructure and application features.

Accessibility of infrastructure and application features in the arbiter VM is as listed here:

Note

Read operations may be unavailable during synchronization.

Table 102: Arbiter VM behavior

Service/Feature	Behavior	Notes
Crosswork Manager	Read and write operations are supported.	Application management, Node lifecycle, System management, and Showtech support both read and write operations.

Service/Feature	Behavior	Notes
Alarms and Events	Read and write operations are supported for system alarms.	Only system alarms are displayed. Network and Device alarms are not listed.
Notification Policies	Read and write operations are supported.	
Backup and Restore	Supported	
Audit Log	Supported	
Certificate Management	Supported	
Users and Roles	Supported	Supported only for Arbiter-specific users, roles, and API permissions.
AAA	Supported	
Geo Redundancy Manager	Supported	
Cross Cluster	Supported	
System settings	Supported	Settings supported for alarms and events, maintenance mode, notification destination, pre-login disclaimer, and customer satisfaction survey.

21 Upgrade to Geo Redundancy Solution

Topics:

- Cluster: upgrade from non-geo redundant to geo redundant deployment
- Cluster: upgrade an existing geo redundant deployment
- Single VM: upgrade from non-geo redundant to geo redundant deployment

Summarizes the cluster and single VM upgrade workflows for geo redundancy.

This chapter covers cluster upgrades from non-geo redundant and existing geo redundant deployments, as well as a single VM upgrade path. It also includes the Data Gateway and Arbiter VM procedures associated with these workflows.

Cluster: upgrade from non-geo redundant to geo redundant deployment

Summarizes the day 0 workflow for upgrading a cluster from a non-geo redundant deployment to a geo redundant deployment.

This topic provides a high-level description of the tasks required to upgrade from Crosswork Network Controller version 7.1.x (non-geo redundant) to version 7.2.x (geo redundant).

Note

Any day N activity will yield the system ineligible to migrate to a geo redundant solution. You will need to re-install the Crosswork cluster to enable geo redundancy. For more information about enabling geo redundancy on day N, refer to Geo redundancy workflow (Day N) on page 466.

Upgrade workflow

Note

While the cluster installation is in progress, you must upgrade NSO and SR-PCE. Please refer to the *Release Notes for Crosswork Network Controller, Release 7.2.1 SP* to know the NSO and SR-PCE versions compatible with Crosswork Network Controller. The process to upgrade NSO or SR-PCE is not covered in this document. For install instructions, please refer to the relevant product documentation.

- Cisco NSO: Cisco NSO documentation
- Cisco SR-PCE: Cisco IOS XRv 9000 Router Installation Guide.

Table 103: Cluster: upgrade from non-geo redundant to geo redundant deployment (day 0)

Step	Action
1. Convert single instance NSO to NSO HA	Follow the instructions in Convert single instance NSO to NSO HA on page 485 topic.
2. Deploy SR-PCE	Deploy SR-PCE in a Point of Presence (PoP) site closer to the Crosswork's Availability Zone. For more information, refer to the relevant install instructions in the Cisco IOS XRv 9000 Router Installation Guide. After you have upgraded to the compatible version of SR-PCE, configure gRPC on SR-PCE. Ensure that your credentials in the CSV file include those for a gRPC connection, which should match your SSH credentials. In your provider settings, include gRPC as a connection option for SR-PCE
3. Upgrade to the latest version of Crosswork Network Controller.	Follow the upgrade instructions in Upgrade Cisco Crosswork Network Controller on page 369.

Step	Action
4. Deploy the standby cluster and enable geo redundancy.	Follow the instructions in Geo redundancy workflow (day 0) for cluster deployments.
5. Update the providers.	Follow the instructions in Update providers on page 486 topic.
6. Upgrade Crosswork Data Gateway 7.1 to 7.2.x with Geo Redundancy	Follow the instructions in Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy on page 486 topic.
7. Configure the Controller IP for Crosswork Data Gateway	Follow the instructions in Configure controller IP for Crosswork Data Gateway on page 487 topic.
8. Complete the geo enablement operation.	Follow the instructions in Complete geo redundancy enablement on page 488 topic. Optionally, you can complete the geo enablement operation for a three cluster inventory (active, standby, and arbiter). In this case, the arbiter cluster is added as a day-N scenario. For more information, refer to Geo redundancy workflow (Day N) on page 466.

Convert single instance NSO to NSO HA

Lists the steps to convert a single-instance NSO system to NSO HA.

Convert a single-instance NSO system to NSO HA.

This topic explains the procedure to convert a single instance NSO to NSO HA (High Availability). For detailed instructions, please refer to the NSO Administration Guide on HA.

> **Note**
>
> Make a backup and upgrade your NSO setup to the compatible version before executing the below steps.

Follow the below guidelines to create a HA setup from a standalone NSO.

Procedure

1. Determine the High Availability topology to follow: L2 or L3
2. Make a backup of the original NSO system.
3. Clone the original NSO to a new instance.
4. Install the hcc package on both nodes.
5. Configure the high availability and hcc as per the selected network topology.
6. Request to enable high availability on both nodes.
7. Verify the changes made.

Update providers

Explains how to update NSO provider settings and related authentication and certificate information.

Update the NSO provider configuration for the geo-redundant deployment.

After enabling geo redundancy on the active cluster, update the providers.

> **Note**
>
> Skip this step if you are not planning to enable geo redundancy.

Procedure

1. Add the RBAC JWT token on the Cisco NSO VMs.
2. Upload and update the JWT package on the Cisco NSO High Availability VMs.
3. Reload the NCS packages on both VMs.
4. Update the **JWT auth file** with *geo-CW FQDN cnc-host* value on both VMs.
5. Update the *cert.pem* on both VMs.
6. Update NSO with unified cluster endpoint in the **Manage Providers** window.
7. (Optional) Update SR-PCE IP address in the **Manage Providers** window.
8. (Optional) While upgrading from a non-HA setup to geo redundant mode, Crosswork Data Gateway will end with multiple VIPs for southbound devices. These devices need to be set up for syslogs, traps and MDT. In case of MDT, you can use admin DOWN/UP to push the configuration changes to the devices.

   > **Note**
   >
   > Any other external destination needs to be in HA mode with its own unified endpoint in the form of VIP or FQDN.

Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy

Explains how to redeploy Data Gateways and update their high-availability pools for a 7.2.x geo redundant deployment.

Upgrade Crosswork Data Gateway to 7.2.x as part of geo redundancy enablement.

Before you begin

Ensure that you are aware of the following:

- After Crosswork is upgraded, the Data Gateways, virtual Data Gateways, HA pool, and device-mapping configuration are restored.
- The Data Gateway Manager automatically assigns the active Crosswork site as the default site for all the existing Data Gateways.
- The Data Gateway must be enrolled using the FQDN. Enrolling with a Virtual IP (VIP) address can cause the Data Gateway to enter an error state after a Crosswork upgrade, due to a mismatch in enrollment details.

You must upgrade the Data Gateway when transitioning Crosswork Network Controller from a non-geo-redundant to a geo-redundant deployment.

Procedure

1. Redeploy the Data Gateway by uninstalling the existing instance from vCenter and installing a new instance. During redeployment, specify the unified management FQDN as the ControllerIP value in the OVF deployment script.

   Note

   - If you redeploy the Data Gateway by using the same VM name and hostname specified in the OVF deployment script, DG-Manager treats it as an existing gateway and automatically enrolls it with the upgraded Crosswork during migration.
   - After you add a new device or deploy a new gateway, perform a sync operation to improve data accuracy. For more information, refer to View cross cluster status on page 463.

2. Modify the high availability Data Gateway pools:
   - If a new Data Gateway instance is added to a high availability pool from the Standby site, which is currently the Active site, and a switchover occurs. The Data Gateway's role changed from spare to assigned.
   - By default, the existing pools will be tagged as imbalanced, as there are no Data Gateways connected to the standby site. For preserving the Data Gateway balance, deploy new Data Gateways on the standby site.
   - The SBConfig is configured to the Shared option. You must configure it to be Site-specific.
   - Configure the VIP or FQDNs for the standby site.

3. Migrate Data Gateway from a single stack that is IPv4 or IPv6 to a dual stack:
   - Update the pool information by including the dual stack configuration's VIP IPs and gateway details. For instance, if you have configured IPv4 address when creating a pool, you must add the IPv6 information.

   At this point, the Data Gateways in both Site 1 and Site 2 must be upgraded to release 7.2.x.

4. Accept an upgrade acknowledgment message that appears on the Crosswork UI when all the Data Gateways with the Assigned role are in the UP state and the spare Data Gateways in the NOT_READY state.

Data Gateways with the Assigned role start the data collection.

What to do next

If the Data Gateways cannot connect with the active cluster, reenroll the gateway from the interactive menu. Refer to the *Reenroll Crosswork Data Gateway* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* for more information.

Configure controller IP for Crosswork Data Gateway

Explains how to configure the controller address and signing certificate on a Data Gateway VM.

Configure the controller IP or FQDN for the Data Gateway VM.

This topic explains the procedure for configuring the controller IP or FQDN for the data gateway after enabling the geo redundancy feature.

When a data gateway is deployed with an invalid controller IP, it may get stuck in the enrollment process. To address this, reconfigure the controller IP. Also, if a data gateway is enrolled to a Crosswork and there is a change in controller virtual IP or the IP is changed to FQDN due to the enabled geo redundancy feature, it must be reconfigured.

To configure the controller IP for a new enrollment or change the controller IP of an existing Crosswork that the data gateway is enrolled with:

Navigate to the data gateway on the active cluster before the geo redundancy feature is enabled.

Procedure

1. Log in to the data gateway VM on the active cluster before the geo redundancy feature is enabled.
2. In the data gateway VM interactive menu, select **3 Change Current System Settings**.
3. Select **Configure Controller IP/FQDN**.
4. Enter the SCP URI for the controller signing certificate file.
5. Enter the SCP passphrase or the SCP user password for the controller signing certificate file.
6. Enter the controller IP.

   A message appears to confirm that Crosswork has updated the controller's IP or FQDN, and the VM is rebooted.

The data gateway connects to Crosswork and progresses to the UP state. If the data gateways are in the Assigned state with devices attached, they resume data collection.

Complete geo redundancy enablement

Lists the tasks to complete geo redundancy enablement on the standby and active clusters.

Complete the remaining geo redundancy enablement tasks.

Before you begin

Ensure inventory file is prepared as per the instructions in Geo redundancy inventory file guidelines. For reference, refer to Day 0: Geo inventory for active cluster deployed with peer clusters (standby and arbiter).

- Since both clusters are already deployed, set `is_skip_peer_check_enabled` parameter to **false**
- Since this is a migration flow, set `is_post_migration_activation` paramater to **true**

After updating the providers, activate geo redundancy on the standby cluster.

Optionally, you can complete the geo enablement operation for a three cluster inventory (active, standby, and arbiter). In this case, the arbiter cluster is added as a day-N scenario. For more information, refer to Geo redundancy workflow (Day N)

> **Note**
>
> Skip this step if you are not planning to enable geo redundancy.

Procedure

1. Log in to Site 2 (set as standby) and enable geo redundancy on the standby cluster. For detailed instructions, refer to Enable geo redundancy.

2. Configure the cross cluster settings. For more information, refer to Configure cross cluster settings.
3. Perform a on-demand sync to sync the data from active to standby cluster.

   Once geo redundancy is enabled, a **Geo Redundancy** tile is automatically added to the **Application management** window. This tile is built-in and cannot be upgraded, uninstalled, or deactivated.
4. (Optional) Add an arbiter VM and enable geo redundancy. For more information, refer to Geo redundancy workflow (Day N).

Cluster: upgrade an existing geo redundant deployment

Summarizes the two-site workflow for upgrading an existing geo redundant deployment.

This topic provides a high-level description of the tasks required to upgrade the geo redundant clusters in Crosswork Network Controller version 7.1.x to version 7.2.1 SP (geo redundant). This migration is performed in geo redundant clusters across two sites while ensuring reduced migration downtime.

> **Note**
>
> Any day N activity will yield the system ineligible to migrate to a geo redundant solution. You will need to re-install the Crosswork cluster to enable geo redundancy. For more information about enabling geo redundancy on day N, refer to Geo redundancy workflow (Day N) on page 466.

Prerequisites

- You have installed Crosswork Network Controller version 7.1.x and enabled geo redundancy on both active and standby cluster sites. For more information about the 7.1.x release, please refer to the release notes.
- Geo synchronization must be disabled and avoided during the upgrade from version 7.1.x to 7.2.1 SP. It should only be performed after all clusters have been upgraded to version 7.2.1 SP and geo enablement has been fully completed.
- **Arbiter VM prerequisite**: If the geo-redundant deployment includes an Arbiter VM running release 7.1.x, remove the Arbiter VM before starting the two-site migration. Re-add the Arbiter VM after Site 1 and Site 2 are upgraded to release 7.2.1 SP and the two-site geo-redundant deployment is healthy. For the procedure to deploy and enable the Arbiter VM after migration, refer to Arbiter VM migration on page 506.

Upgrade workflow

Table 104: Upgrade from Crosswork Network Controller 7.1.x (geo redundant) to 7.2.1 SP (geo redundant)

Step	Action
1. Upgrade and deploy SR-PCE	Deploy SR-PCE in a Point of Presence (PoP) Site closer to the Crosswork's Availability Zone. For more information, refer to the relevant install instructions in the Cisco IOS XRv 9000 Router Installation Guide. After you have upgraded to the compatible version of SR-PCE, configure gRPC on SR-PCE. Ensure that your credentials in the CSV file include those for a gRPC connection, which should match your SSH credentials. In your provider settings, include gRPC as a connection option for SR-PCE
2. Upgrade the secondary NSO (Site 2).	If your NSO is running an older version, perform these steps: **1.** Set the primary NSO (Site 1) to read-only mode. **2.** Disable HA on the secondary NSO (Site 2). **3.** Upgrade the secondary NSO (Site 2). Follow instructions on converting single instance NSO to NSO HA, refer to Convert single instance NSO to NSO HA on page 485 topic.
3. (Optional) Create backup of the Crosswork Network Controller 7.1.x cluster.	Follow the instructions in Create backup of the Cisco Crosswork cluster on page 493 topic.
4. Shut down the Crosswork Network Controller 7.1.x cluster on Site 2.	Shut down the Crosswork Network Controller 7.1.x cluster by powering down the VMs hosting each node (start with the Hybrid VMs). **1.** Gather following information before shutting down the cluster. • All IP addresses of the cluster. • All IP addresses of the Crosswork Data Gateways **2.** Shut down the VMs of the Crosswork Network Controller cluster. For vcenter shutdown all the VMs using vcenter UI **3.** Log into the VMware vSphere Web Client. In the **Navigator** pane, right-click the VM that you want to shut down. **4.** Choose **Power** > **Power Off**. Wait for the VM status to change to **Off**. **5.** Wait for 30 seconds and repeat the steps for each of the remaining VMs. **6.** (Optional) Put NSO in read-only mode using `ncs_cmd -c maapi_read_only` command.

Step	Action
5. Install the Crosswork Network Controller 7.2.1 SP cluster and applications on Site 2.	Follow the instructions in Install the Crosswork Network Controller cluster and applications on page 496 topic. **Note** Site 2 cluster will be non-geo redundant after installation of Crosswork Network Controller 7.2.1 SP. After the cluster installation, ensure that the Crosswork inventory is onboarded before proceeding with the upgrade.
6. Trigger the multi-cluster upgrade from the Site 2 cluster.	Follow the instructions in Multi cluster upgrade on page 497 topic.
7. Upgrade Crosswork Data Gateway on Site 2 to version 7.2.1 SP.	Follow the instructions in Upgrade Crosswork Data Gateway to 7.2.x on the standby site on page 498 topic.
8. Update the providers.	Follow the instructions in Update providers on page 486 topic.
9. Enable geo redundancy on the Site 2 cluster.	Follow the instructions in Enable geo redundancy on the Site 2 cluster on page 500 topic.
(Optional) 10. Shut down Crosswork Data Gateway on Site 1.	Follow the instructions in Shut down Crosswork Data Gateway in the active site on page 501 topic.
11. Perform switchover from Site 1 to Site 2.	For the instructions in the Perform switchover between sites on page 502 topic. **Note** Do not perform a sync operation from Site 1 to Site 2 during the switchover.
12. Upgrade Crosswork Data Gateway on Site 1 to version 7.2.1 SP.	Follow the instructions in Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy on page 486 topic.

Step	Action
13. Upgrade the primary NSO (Site 1).	Since the primary NSO (Site 1) is running an older version, perform these steps: **1.** Disable HA on the primary NSO (Site 1) before HA is enabled on the upgraded secondary NSO (Site 2). **2.** Upgrade the primary NSO (Site 1). **3.** After primary NSO (Site 1) is upgraded, enable HA with node as secondary. This will sync all the data from the primary. Follow instructions on converting single instance NSO to NSO HA, refer to Convert single instance NSO to NSO HA on page 485 topic.
14. Shut down the Crosswork Network Controller 7.1.x cluster on Site 1.	Shut down the Crosswork Network Controller 7.1.x cluster by powering down the VMs hosting each node (start with the Hybrid VMs). **1.** Gather following information before shutting down the cluster. • All IP addresses of the cluster. • All IP addresses of the Crosswork Data Gateways **2.** Shut down the VMs of the Crosswork Network Controller cluster. For vcenter shutdown all the VMs using vcenter UI **3.** Log into the VMware vSphere Web Client. In the **Navigator** pane, right-click the VM that you want to shut down. **4.** Choose **Power** > **Power Off**. Wait for the VM status to change to **Off**. **5.** Wait for 30 seconds and repeat the steps for each of the remaining VMs. **6.** (Optional) Put NSO in read-only mode using `ncs_cmd -c maapi_read_only` command.
15. Install the Crosswork Network Controller 7.2.1 SP cluster and applications on Site 1.	Follow the instructions in Install the Crosswork Network Controller cluster and applications on page 496 topic. **Note** Site 1 cluster will be non-geo redundant after installation of Crosswork Network Controller 7.2.1 SP. After the cluster installation, ensure that the Crosswork inventory is onboarded before proceeding.

Step	Action
16. Complete the geo enablement operation on the Site 1 cluster.	Follow the instructions in Enable geo redundancy on the Site 1 cluster on page 505
17. Enable and initiate sync in Site 2 (sync from Site 2 to Site 1)	Perform these steps: **1.** Log in to the Site 2 cluster (currently, the active cluster). **2.** Verify that the geo mode status is completed and the cross-cluster operational status is healthy. **3.** Click on **Configurations** > **Sync settings**, and enable the **Sync** slider button. **4.** Initiate a sync or wait for a configured periodic sync to complete. **5.** After completing the sync, verify that both Site 1 and Site 2 are healthy and running Crosswork Network Controller version 7.2.1 SP.
(Optional) 18. Perform the second switchover (Site 2 to Site 1)	If you need to make the original site active again, perform a switchover from **Site 2** to **Site 1**. Follow the instructions in Perform switchover between sites on page 502, but ensure that you reverse the site roles accordingly.
(Optional) 19. Complete the geo enablement operation with the arbiter node. **Note** You can skip this step if you do not wish to deploy an arbiter node.	Follow the instructions in Geo redundancy workflow (Day N) on page 466.

Create backup of the Cisco Crosswork cluster

Explains how to create and monitor a backup of the Crosswork cluster.

Create a backup of the Cisco Crosswork cluster.

Before you begin

Follow these guidelines whenever you create a backup:

- Cisco Crosswork will back up the configuration of the system to an external server using SCP. Before you begin you need to have the following configuration in place and information about the SCP server available:
  - The hostname or IP address and the port number of a secure SCP server.
  - A preconfigured path on the SCP server where the backup will be stored.
  - User credentials with file read and write permissions to the directory.
  - The SCP server storage requirements will vary slightly but you must have at least 25 GB of storage.

- Ensure that you have configured a destination SCP server to store the backup files. This configuration is a one-time activity.
- After the backup operation is completed, navigate to the destination SCP server directory and ensure that the backup file is created. You will require this backup file in the later stages of the upgrade process.
- Both the Crosswork Network Controller cluster and the SCP server must be in the same IP environment. For example: If Crosswork Network Controller is communicating over IPv6, so must the backup server.
- Keep a record of the list of Crosswork applications you have installed in the current version of Cisco Crosswork, as you can only install those applications after migrating to the new version of Cisco Crosswork.
- If you have onboarded a custom MIB package in the current version of Cisco Crosswork, download a copy of the package to your system. You will need to upload the package after you complete migrating to new version of Cisco Crosswork.
- If you have modified the current version of Cisco Crosswork to include third-party device types, you must download the third-party device configuration file, and re-apply it to the new version of Cisco Crosswork. The device configuration file is located on the cluster node (at `/mnt/cw_glusterfs/bricks/brick3/sys-oids.yaml`) and on the pod (at `/mnt/backup/sys-oids.yaml`).
- If Cisco Crosswork Optimization Engine has feature packs (Local Congestion Mitigation (LCM), Bandwidth Optimization (BWOpt), and Bandwidth on Demand (BWoD)) that are enabled, you must disable them before proceeding. You must also, if available, export the current list of interfaces managed by LCM or BWOpt (**Service and Traffic Engineering > Local Congestion Mitigation > Domain Identifier <domain_id> > Interface Thresholds > Export** OR **Service and Traffic Engineering > Bandwidth Optimization > Interface Thresholds > Export** icon).

Creating a backup is a prerequisite when upgrading your current version of Crosswork Network Controller to a new version.

> **Note**
>
> We recommend that you create a backup only during a scheduled upgrade window. Users should not attempt to access Crosswork Network Controller while the backup operation is running.

Procedure

1. Login to the Crosswork UI by navigating to https://<VIP>:30603.

   The VIP refers to the Management Virtual IP of the cluster.
2. Check and confirm that all the VMs are healthy and running in your cluster.
3. **Configure an SCP backup server:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Destination** to display the **Edit Destination** drawer panel. Make the relevant entries in the fields provided.
   c) Click **Save** to confirm the backup server details.
4. **Create a backup:**
   a) From the Crosswork Network Controller main menu, choose **Administration** > **Backup and Restore**.
   b) Click **Actions** > **Backup** to display the **Backup** drawer panel with the destination server details prefilled.
   c) Provide a relevant name for the backup in the **Job Name** field.
   d) If any of the VMs or applications are not in **Healthy** state, but you want to create the backup, check the **Force** check box.

> **Note**
>
> The **Force** option must be used only after consultation with the Cisco Customer Experience team.

e) Uncheck the **Backup NSO** checkbox if you do not want to include Cisco NSO data in the backup.

> **Note**
>
> To use the **Backup NSO** option during backup, you must configure the SSH connectivity protocol in the NSO provider; otherwise, the backup will fail. Follow the instructions given in the *Backup Cisco Crosswork with Cisco NSO* section in *Cisco Crosswork Network Controller 7.2.x Administration Guide* instead of the instructions here.

f) Complete the remaining fields as needed.

If you want to specify a different remote server upload destination: Edit the prefilled **Host Name**, **Port**, **Username**, **Password** and **Remote Path** fields to specify a different destination.

g) (Optional) Click **Verify Backup Readiness** to verify that Crosswork Network Controller has enough free resources to complete the backup. If the verifications are successful, Crosswork Network Controller displays a warning about the time-consuming nature of the operation. Click **OK**.

If the verification is unsuccessful, please contact the Cisco Customer Experience team for assistance.

h) Click **Start Backup** to start the backup operation. Crosswork Network Controller creates the corresponding backup job set and adds it to the job list. The Job Details panel reports the status of each backup step as it is completed.

> **Note**
>
> You can also perform a data backup (backup using rest-api). The data backup is faster as it does not include application binaries. To perform, do the following:
>
> - Get JWT (using sso apis)
> - API to take the data backup (https://<VIP>:30603/crosswork/platform/v1/platform/backup/dataonly)
> - Payload for the api {"jobName": "jobname","force": false}

i) To view the progress of a backup job: Enter the job details (such as Status or Job Type) in the search fields in the **Backup and Restore Job Sets** table. Then click on the job set you want.

The **Job Details** panel displays information about the selected job set, such as the job Status, Job Type, and Start Time. If there's a failed job, hover the mouse pointer over the icon near the **Status** column to view the error details.

> **Note**
>
> After the backup operation is completed, navigate to the destination SCP server directory and ensure that the backup file is created. You will require this backup file in the later stages of the upgrade process.

> **Note**
>
> If you do not refer to your backup job in the list, refresh the **Backup and Restore Job Sets** table.

a) If the backup fails during upload to the remote server: In the **Job Details** panel, just under the Status icon, click the **Upload backup** button to retry the upload.

> **Note**
>
> Upload can fail due to connectivity problems with the SCP backup server (for example, incorrect credentials, missing directory or directory permissions, missing path and so on). This is indicated by failure of the task `uploadBackupToRemote`). If this happens, check the SCP server details, correct any mistakes and try again. Alternatively, you can use the **Destination** button to specify a different SCP server and path before clicking **Upload backup**.

Install the Crosswork Network Controller cluster and applications

Lists the steps to install the upgraded cluster and its applications.

Install the upgraded Crosswork Network Controller cluster and applications.

Before you begin

- Make sure that your environment meets all the installation prerequisites (refer to Installation Prerequisites on page 43).

This install the latest version of the Crosswork Network Controller cluster and applications.

> **Note**
>
> While the cluster installation is in progress, you must upgrade NSO and SR-PCE. Please refer to the *Cisco Crosswork Network Controller 7.2.1 Security Patch (SP) Release Notes* to know the NSO and SR-PCE versions compatible with Crosswork Network Controller. Please monitor actively to ensure that the NSO leader is in the same site as Crosswork. The process to upgrade NSO or SR-PCE is not covered in this document. For install instructions, please refer to the relevant product documentation.
>
> - Cisco NSO: Cisco NSO documentation
> - Cisco SR-PCE: Cisco IOS XRv 9000 Router Installation Guide.

Procedure

1. Install the new version of the Crosswork Network Controller cluster (refer to Install Cluster on VMware vCenter on page 81) using the same IP addresses and same number of nodes as in old cluster.
2. After the installation is completed, log into the Crosswork Network Controller UI (using https://<VIP>:30603) and check if all the nodes are up and running in the cluster.
3. After the cluster installation, ensure that the Crosswork inventory is onboarded.
4. Install the Crosswork Network Controller applications which were installed in the old cluster. Ensure that you install the latest versions that are compatible with the new version of the cluster. For installation instructions, please refer to the Install Crosswork Applications on page 265 chapter.

   > **Note**
   >
   > The applications binaries and versions are not updated in the migration job.

5. After the applications are successfully installed, check the health of the new cluster. From the Crosswork Network Controller main menu, choose **Administration** > **Crosswork Manager**.
   - Click the **System summary** tile to view the health of the cluster nodes.
   - Click the **Crosswork health** tab to view the health of the pods of all installed applications.

Multi cluster upgrade

Explains how to run and monitor a multi-cluster upgrade.

Upgrade multiple clusters by using the multi-cluster upgrade workflow.

Before you begin

Ensure you have downloaded and prepared the inventory file (.yaml).

This section explains how to perform the multi-cluster upgrade in Crosswork Network Controller.

Procedure

1. From the Crosswork Network Controller main menu, choose **Administration** > **Geo Redundancy Manager**. The **Geo Redundancy Manager** window is displayed.
2. Click on the **Multi-cluster upgrade** tab.
3. Click on **Import upgrade inventory**. The **Multi-Cluster Upgrade** drawer panel is displayed.
4. (Optional) Click **sample file** to download a sample inventory file (.yaml) file for multi-cluster upgrade.
5. Click **Browse** to select and upload the inventory file (.yaml) file from your machine. Ensure the `cluster_id` value in the inventory file matches with the value specified in regular geo inventory.

   Note

   The upgrade job storage path settings in the inventory file must be different from the configured path in the Cross Cluster storage settings.

6. Click the **Jobs** tab to view the progress of the upgrade job.

   The multi-cluster upgrade job creates a data backup backup job on site 1, and a data migration job on site 2. Navigate to the **Backup and Restore** window (**Administration** > **Backup and Restore**) in each site to view the jobs.

   Note

   Once the upgrade job starts on site 2, the UI will be unavailable for some time. An alarm is raised on site 1 with Managed Object as *Upgrade Alarm*. The description of this alarm contains the Grafana dashboard URL. You can launch this URL on the browser to monitor the upgrade job progress.

7. Once the job is completed, check the health of all applications and ensure that they are in a healthy state.

   The data from Crosswork Network Controller 7.1 on site 1 is now migrated to site 2 running Crosswork Network Controller 7.2.x.

Upgrade Crosswork Data Gateway to 7.2.x on the standby site

Explains how to redeploy standby-site Data Gateways and complete the migration acknowledgment.

Upgrade the Data Gateways on the standby site.

Before you begin

Ensure that you are aware of:

- Data Gateways from the standby site must not be in the maintenance mode on the active site.
- Data collection downtime occurs during this procedure. Total downtime includes:
  - the Data Gateway upgrade duration and
  - any delay in receiving jobs from the northbound interface.

Standby-site Data Gateway behavior during upgrade

- When connected to the active site, standby-site Data Gateways have the Spare role.
- Upgrading standby-site Data Gateways does not affect data collection on the active site.
- After redeployment, the upgraded Data Gateways enter the Error state on the active site until reenrollment is completed.

> **Note**
>
> After the first switchover, the redeployed Data Gateways automatically enroll with Crosswork on the standby site.

Use this procedure to upgrade Crosswork Data Gateways on the standby site when upgrading Crosswork Network Controller to release 7.2.x in a geo-redundant deployment.

This procedure upgrades and redeploys the standby-site Data Gateways before switchover.

Procedure

1. Redeploy the standby-site Data Gateway.

   Redeploy the Data Gateway by uninstalling the existing instance from vCenter and installing a new instance. During redeployment, specify the unified management FQDN as the ControllerIP value in the OVF deployment script.

   > **Note**
   >
   > - If you redeploy the Data Gateway by using the same VM name and hostname specified in the OVF deployment script, DG-Manager treats it as an existing gateway and automatically enrolls it with the upgraded Crosswork during migration.
   > - After you add a new device or deploy a new gateway, perform a sync operation to improve data accuracy. For more information, refer to View cross cluster status on page 463.

2. Cancel the migration acknowledgment.

   Log in to the Site 2 Crosswork Network Controller UI. When the **Action to Be Taken** popup appears and asks you to confirm completion of data migration after redeploying the Crosswork Data Gateways, click **Cancel**.

   At this point, the Data Gateways in both Site 1 and Site 2 must be upgraded to release 7.2.x.

   > **Note**
   >
   > Do not acknowledge the popup until the Data Gateways in Site 1 are also upgraded and redeployed. This acknowledgment is required for the gateways to move to the upgraded state. After the Site 1 upgrade is complete, log back in to the Site 2 Crosswork Network Controller UI and acknowledge the popup.

3. If you are migrating from single stack to dual stack, update the pool configuration.

If you are migrating a Data Gateway from a single-stack IPv4 or IPv6 configuration to a dual-stack configuration, update the pool configuration to include the VIP addresses and Data Gateway information for both protocol stacks. For example, if the pool was originally configured with only IPv4 information, add the corresponding IPv6 VIP and Data Gateway details.

What to do next

After the standby site becomes active, the Data Gateway roles, image deployment state, and high-availability pool configuration may change. Review the Post-switchover procedures for standby site on page 505 to understand the expected system behavior and any required post-switchover updates.

Update providers

Explains how to update NSO provider settings and related authentication and certificate information.

Update the NSO provider configuration for the geo-redundant deployment.

After enabling geo redundancy on the active cluster, update the providers.

> **Note**
>
> Skip this step if you are not planning to enable geo redundancy.

Procedure

1. Add the RBAC JWT token on the Cisco NSO VMs.
2. Upload and update the JWT package on the Cisco NSO High Availability VMs.
3. Reload the NCS packages on both VMs.
4. Update the **JWT auth file** with *geo-CW FQDN cnc-host* value on both VMs.
5. Update the *cert.pem* on both VMs.
6. Update NSO with unified cluster endpoint in the **Manage Providers** window.
7. (Optional) Update SR-PCE IP address in the **Manage Providers** window.
8. (Optional) While upgrading from a non-HA setup to geo redundant mode, Crosswork Data Gateway will end with multiple VIPs for southbound devices. These devices need to be set up for syslogs, traps and MDT. In case of MDT, you can use admin DOWN/UP to push the configuration changes to the devices.

   > **Note**
   >
   > Any other external destination needs to be in HA mode with its own unified endpoint in the form of VIP or FQDN.

Enable geo redundancy on the Site 2 cluster

Lists the steps to create the Site 2 standby cluster and verify geo redundancy status.

Enable geo redundancy on the Site 2 cluster.

Before you begin

Kindly ensure

- the connectivity checks work for east-west communication.
- DNS connectivity checks work from both clusters.
- the unified cross-cluster endpoint resolves to the Site 1 cluster (active cluster).
- inventory file is prepared as per the instructions in Geo redundancy inventory file guidelines. For reference, refer to Day N: Geo inventory for active cluster deployed without peer cluster (standby).
  - Set `is_skip_peer_check_enabled` parameter to **true**
  - Set `is_post_migration_activation` paramater to **true**
- the `cluster_id` parameter value matches with the value in the upgrade inventory yaml in Multi cluster upgrade. Any mismatch would cause the geo redundancy pairing to fail.

After updating the providers, activate geo redundancy on the Site 2 cluster.

Procedure

1. Log in to the Site 1 cluster (currently, the active cluster), and navigate to the **Geo Redundancy Manager** window. Enable the **pairing mode** slider. This ensures the active cluster accepts the newly upgraded cluster for multi-cluster upgrade.
2. Create and upload the cluster inventory file on Site 2, to create the standby cluster. For detailed instructions, refer to Enable geo redundancy.
3. Verify that the geo mode status is completed and the cross-cluster operational status is healthy on both Sites.

   At the end of this procedure, Site 1 will serve as the active cluster with Crosswork Network Controller version 7.1, while Site 2 will function as the standby cluster with Crosswork Network Controller version 7.2.x installed.

   Note

   Do not perform a sync operation from Site 1 to Site 2 after this step.

Shut down Crosswork Data Gateway in the active site

Explains how to shut down active-site Data Gateway VMs before a cluster switchover.

Shut down the active-site Data Gateway VMs before switchover.

Before you begin

Ensure that you are aware of the following:

- Shutting down the Data Gateway in the active site interrupts data collections. To minimize the downtime, shutdown Data Gateway only when you are ready to initiate the switchover.
- If you choose not to shut down the Data Gateway:

- After a switchover, the Data Gateway in the active site connects to the standby site.
- When the Data Gateway in the active site downloads the 7.2.x configuration, its operational state changes to NOT_READY.

Before initiating the Crosswork cluster switchover, shut down the Data Gateway on the active site. This action preserves the 7.1 configuration and prevents the automatic download of the 7.2.x version.

Procedure

1. Shut down all the Data Gateway VMs in the active site.
   a) Log in to the Data Gateway VM. For information, refer to Log in and Log out of the Data Gateway VM on page 156. Crosswork Data Gateway launches an Interactive Console after your login.
   b) Choose **5 Troubleshooting**.
   c) From the **Troubleshooting** menu, choose **5 Shutdown VM** to shut down the VM.
2. Perform Crosswork cluster switchover. For more information, refer to Geo Redundancy Switchover on page 511.

   Note

   If an issue occurs on the standby site before completing the switchover, and you need to revert to the active site, which is still running version 7.1:

   - Make the now standby site active again
   - Power on the Data Gateways in that site

   Once powered on, the Data Gateways should resume data collection as expected.

After the switchover to the standby site, the Data Gateways enters the ERROR state. This behavior is expected and has no impact on active data collection, as the affected Data Gateways function only as spares. These Data Gateways are currently running version 7.1 and are re-deployed with version 7.2.x in a later step.

Perform switchover between sites

Lists the steps to switch the active role to Site 2 and update the DNS records.

Switch the active role to Site 2 and update the DNS records.

Before you begin

During migration of a geo-redundant deployment from release 7.1.x to release 7.2.1 SP, use the `SetRole` APIs to perform the first switchover. Do not use the UI **Switchover** action. The UI **Switchover** action may fail during migration and might not restore the original cluster roles. This issue can occur for any user account, including the `admin` user.

Use the `SetRole` APIs to set Site 2 to Active, update the DNS records to point to Site 2, and then set Site 1 to Standby.

This topic explains how to perform the first switchover between the cluster sites.

Procedure

1. Set Site 2 to Active by using the `SetRole` API. Execute the API from the local terminal.

```
curl --insecure --location
'https://<site-2-managementvip>:30603/crosswork/crosscluster/v1/role' --header
 'Content-Type: application/json' --header 'Authorization: <your bearer token>'
--data '{
"cross_cluster_component": {
"id": "<cross cluster id of site2 mentioned in inventory file>",
"name": "<cross cluster name of site2 mentioned in inventory file>",
"component_type": "CROSSWORK_CLUSTER"
},
"preferred_cluster_leadership_state": "ACTIVE",
"force": true
}'
```

> **Note**
>
> You can get cross cluster name and id from the **view inventory** option in the **Geo Redundancy Manager**window.

2. Update the Management FQDN and Data FQDN DNS records to point to Site 2.
3. Set Site 1 to Standby by using the `SetRole` API. Execute the API from the local terminal.

```
curl --insecure --location
'https://<site-1-managementvip>:30603/crosswork/crosscluster/v1/role' --header
 'Content-Type: application/json' --header 'Authorization: <your bearer token>'
--data '{
"cross_cluster_component": {
"id": "<cross cluster id of site1 mentioned in inventory file>",
"name": "<cross cluster name of site1 mentioned in inventory file>",
"component_type": "CROSSWORK_CLUSTER"
},
"preferred_cluster_leadership_state": "STANDBY",
"force": true
}'
```

4. Verify that all services are healthy and that the cross-cluster operational status is healthy.
5. On Site 2, disable the cross cluster sync setting.

Upgrade Crosswork Data Gateway to 7.2.x for geo redundancy

Explains how to redeploy Data Gateways and update their high-availability pools for a 7.2.x geo redundant deployment.

Upgrade Crosswork Data Gateway to 7.2.x as part of geo redundancy enablement.

Before you begin

Ensure that you are aware of the following:

- After Crosswork is upgraded, the Data Gateways, virtual Data Gateways, HA pool, and device-mapping configuration are restored.

- The Data Gateway Manager automatically assigns the active Crosswork site as the default site for all the existing Data Gateways.
- The Data Gateway must be enrolled using the FQDN. Enrolling with a Virtual IP (VIP) address can cause the Data Gateway to enter an error state after a Crosswork upgrade, due to a mismatch in enrollment details.

You must upgrade the Data Gateway when transitioning Crosswork Network Controller from a non-geo-redundant to a geo-redundant deployment.

Procedure

1. Redeploy the Data Gateway by uninstalling the existing instance from vCenter and installing a new instance. During redeployment, specify the unified management FQDN as the ControllerIP value in the OVF deployment script.

   Note

   - If you redeploy the Data Gateway by using the same VM name and hostname specified in the OVF deployment script, DG-Manager treats it as an existing gateway and automatically enrolls it with the upgraded Crosswork during migration.
   - After you add a new device or deploy a new gateway, perform a sync operation to improve data accuracy. For more information, refer to View cross cluster status on page 463.

2. Modify the high availability Data Gateway pools:
   - If a new Data Gateway instance is added to a high availability pool from the Standby site, which is currently the Active site, and a switchover occurs. The Data Gateway's role changed from spare to assigned.
   - By default, the existing pools will be tagged as imbalanced, as there are no Data Gateways connected to the standby site. For preserving the Data Gateway balance, deploy new Data Gateways on the standby site.
   - The SBConfig is configured to the Shared option. You must configure it to be Site-specific.
   - Configure the VIP or FQDNs for the standby site.
3. Migrate Data Gateway from a single stack that is IPv4 or IPv6 to a dual stack:
   - Update the pool information by including the dual stack configuration's VIP IPs and gateway details. For instance, if you have configured IPv4 address when creating a pool, you must add the IPv6 information.

   At this point, the Data Gateways in both Site 1 and Site 2 must be upgraded to release 7.2.x.
4. Accept an upgrade acknowledgment message that appears on the Crosswork UI when all the Data Gateways with the Assigned role are in the UP state and the spare Data Gateways in the NOT_READY state.

Data Gateways with the Assigned role start the data collection.

What to do next

If the Data Gateways cannot connect with the active cluster, reenroll the gateway from the interactive menu. Refer to the *Reenroll Crosswork Data Gateway* section in the *Cisco Crosswork Network Controller 7.2.x Administration Guide* for more information.

Post-switchover procedures for standby site

Describes Data Gateway behavior and follow-up procedures after switchover to the standby site.

To describe the sequence of events and system behaviors that occur after a switchover to the standby site, ensuring that data collection and related operations resume seamlessly during Crosswork 7.2.x deployments.

After switchover to the standby site is complete, these events occur:

- Data Gateways on the standby site move to the Assigned role and take over data collection from the active site.
- Each Data Gateway begins downloading the Crosswork 7.2.x images.
- After the 7.2.x containers are deployed
  - The Data Gateways in the Assigned role resume data collection.
  - Inventory and collection jobs are retransmitted from the northbound interface.

Update high-availability data gateway pools after switchover

After switchover to the standby site, review and update the high-availability Data Gateway pool configuration as needed.

- If a new standby-site Data Gateway is added to a high-availability pool and a switchover occurs, its role changes from Spare to Assigned.
- Existing pools may appear as Imbalanced if no Data Gateways are connected to the standby site.
- To maintain Data Gateway balance across sites, deploy new Data Gateways on the standby site.
- If `SBConfig` is set to Shared, change it to Site-specific.
- Configure the standby-site VIPs or FQDNs as required.

Enable geo redundancy on the Site 1 cluster

Lists the steps to enable geo redundancy on Site 1 after the upgrade.

Enable geo redundancy on the Site 1 cluster.

Before you begin

- Ensure connectivity checks work for east-west communication.
- Ensure DNS connectivity checks work from both clusters.
- Ensure the unified cross-cluster endpoint resolves to the Site 2 cluster (active cluster).
- Ensure you prepare the inventory file as per the instructions in Geo redundancy inventory file guidelines on page 442 and Enable geo redundancy on page 452.
  - Set `is_skip_peer_check_enabled` parameter to **false**
  - Set `is_post_migration_activation` paramater to **false**
  - As Site 2 cluster is already up and in ACTIVE state, set the `initial_preferred_leadership_state` parameter for Site 1 to **STANDBY**

After updating the providers, activate geo redundancy on the Site 1 cluster.

Procedure

1. Log in to the Site 2 cluster (currently, the active cluster), and navigate to the **Geo Redundancy Manager** window. Enable the **pairing mode** slider. This ensures the active cluster accepts the newly upgraded cluster for multi-cluster upgrade.
2. On Site 1, create and upload the geo redundancy inventory by using the migration-specific settings for this phase of the upgrade. For general wizard details, refer to Enable geo redundancy on page 452.

   In this phase of the migration:

   - Set the intended cluster role for Site 1 to **Standby**.
   - Use the active Site 2 cluster details to fetch the peer-cluster data as required by the wizard.
   - Set `is_skip_peer_check_enabled` to **false**.
   - Set `is_post_migration_activation` to **false**.

   Verify that the inventory reflects Site 2 as the active cluster and Site 1 as the standby cluster before you proceed with geo activation.

   Note

   If you accidentally activate Site 1 with `initial_preferred_leadership_state` parameter as `ACTIVE`, perform these recovery steps:

   - Reset the geo mode using the `geo_reset` command from one of the Hybrid nodes in the cluster.

     ```
     geo_reset -h
     Usage: /usr/local/bin/geo_reset -r ORCH_REPLICAS

     Options:
       -r     ORCH_REPLICAS | Specify the number of
     robot-orchestrator replicas between 1 and 2
              : For SVM replicas is 1
              : For Cluster replicas is 2
       -h     Show this help message
     ```

   - Uninstall the geo redundancy CAPP file, and ensure the cross cluster pod is not running.
   - Update the inventory file with the correct configurations, and perform the geo activation again.

3. Verify that the geo mode status is completed and the cross-cluster operational status is healthy on both Sites.

   At the end of this procedure, Site 2 functions as the active cluster and Site 1 functions as the standby cluster. Both sites run Crosswork Network Controller version 7.2.x.

4. Configure the cross cluster settings with SCP host over data network.

Arbiter VM migration

Explains how to back up, redeploy, and restore Arbiter VM data during migration.

Migrate the Arbiter VM to the newer release and restore its data.

Before you begin

Before migrating the Arbiter VM, ensure the following conditions are met:

- Migrate the Active and Standby clusters to version 7.2.1 SP. Alternatively, you may migrate the Arbiter VM independently if required by your specific deployment plan.
- Verify that the three-way topology is visible in the Cisco Crosswork UI and all heartbeat lines are green.
- Perform a full backup of the Arbiter VM.

This section describes how to migrate the Arbiter VM when upgrading a Crosswork Network Controller geo redundant solution to version 7.2.1 SP.

Procedure

1. **Back up the Arbiter VM:**
   a) Log in to the Arbiter VM.
   b) Go to **Administration** > **Backup and Restore**.
   c) Configure the backup server and perform a manual backup of the Arbiter VM. Ensure the backup file is stored in a secure, reachable location.
2. **Remove the old version of Arbiter VM:**
   a) Log in to the data center UI.
   b) Power off the Arbiter VM.
   c) Delete the Arbiter VM from the cluster inventory.
3. **Deploy the newer version of Arbiter VM:**
   a) Install a new Arbiter VM using the Crosswork Network Controller 7.2.1 SP release. For detailed instructions, refer to Deploy the Arbiter VM on page 435.
   b) Power on the VM and wait for the Arbiter VM to reach a healthy state.
4. **Activate geo redundancy:** For detailed instructions, refer to Enable geo redundancy on page 452.
   a) Log in to the Crosswork cluster that will function as the Arbiter VM. From the main menu, choose **Administration** > **Geo Redundancy Manager**. Click **Get started with geo enablement**, and choose **Build inventory YAML manually**.
   b) In the **Cross cluster details** section, verify the cluster role is set to **Arbiter**.
   c) Click **Fetch active cluster data**. Enter the Active cluster credentials to retrieve the inventory from Site 1.
   d) Select the **Post migration - geo enabled** checkbox.
   e) Enter the secret string.

   > **Note**
   >
   > Password and secret fields are masked when fetching data from a peer cluster.

   f) Manually update the password fields for the Active and Standby clusters.
   g) Complete the validation stage and resolve any errors.
   h) Click **Activate geo cross cluster**. Verify that the geo mode status is completed and the cross cluster operational status is **healthy**.
5. **Configure cross cluster settings:** Configure these settings to ensure proper synchronization and arbitration.
   a) **SCP server:** Provide the SCP host server details. Ensure the remote path has read/write permissions and is not shared with other geo redundancy setups.

b) **DNS:** Enter the Management and Data FQDN details.
c) **Synchronization:**
   - Use the Cisco Crosswork APIs (`SetSettings` and `SyncronizeSettings`) to push synchronization configurations to the Active and Standby clusters.
   - Configure on-demand and cadence syncing schedules.
d) **Arbitration:**
   - Define the heartbeat (ping-pong) interval and the failure detection interval.
   - Enable **Auto-Arbitration**.
   - Provide the SCP server details where the DNS update script is located, including the absolute path, IP address, and login credentials.
e) **Notifications:** Configure the required notification settings for cluster events.

6. **Verify geo redundancy health:**
   a) Navigate to the Cross cluster health status dashboard.
   b) Confirm that the Active, Standby, and Arbiter states are healthy.
   c) Verify the three-way topology map displays green heartbeat lines.
   d) If auto-arbitration is enabled, confirm the cluster leader is visible in the UI.
7. **Restore the Arbiter data:**
   a) Go to **Administration** > **Backup and Restore**.
   b) Configure the destination settings to point to your backup location.
   c) Start the migration (restore) job using the backup file created in Step 1.
   d) Wait for the job to complete.

   Note

   This process typically takes 30 to 60 minutes. The Crosswork Network Controller UI will be unavailable during this time.

8. (Optional) **Alarm handling after restore:** If the `Unbalanced applications detected` or `Platform version mismatch detected` alarm is raised after a backup and restore operation:
   a) Compare the `capp-infra` and `capp-geo` versions on the Active, Standby, and Arbiter clusters.
   b) If the versions do not match, complete the patch cycle on the affected cluster.
   c) If the versions match and the alarm remains, treat the alarm as unexpected and contact Cisco Customer Experience for assistance.

   Note

   Do not install workload applications on the Arbiter. Do not document the `Platform version mismatch detected` alarm as expected or instruct operators to suppress it when the required component versions are aligned.

Single VM: upgrade from non-geo redundant to geo redundant deployment

Summarizes the day 0 workflow for upgrading a single VM deployment to geo redundancy.

This topic provides a high-level description of the tasks required to upgrade single VM deployment of Crosswork Network Controller version 7.1.x (non-geo redundant) to version 7.2.x (geo redundant). Refer to *Geo redundancy workflow (day 0) for single VM deployments* in Geo redundancy workflow (Day 0).

> **Note**
>
> Any day N activity will yield the system ineligible to migrate to a geo redundant solution. You will need to re-install the Crosswork cluster to enable geo redundancy. For more information about enabling geo redundancy on day N, refer to Day N: Geo inventory for active cluster deployed without peer cluster (standby).

Upgrade workflow

Table 105: Single VM: upgrade from non-geo redundant to geo redundant deployment (day 0)

Step	Action
1. Upgrade to the latest version of Crosswork Network Controller.	Follow the upgrade instructions in Upgrade Crosswork Network Controller on a single VM.
2. Deploy the standby cluster and enable geo redundancy.	Follow the instructions in *Geo redundancy workflow (day 0) for single VM deployments* section in Geo redundancy workflow (Day 0).

22 Geo Redundancy Switchover

Topics:

- Geo redundancy switchovers
- Auto-switchover triggers
- How auto-arbitration works
- Perform switchover manually
- Crosswork Optimization Engine license count after a switchover

This chapter describes geo-redundancy switchovers, their triggers, and the auto-arbitration process. It also explains how to perform a manual switchover and update Crosswork Optimization Engine license counts.

This chapter describes how geo-redundancy switchovers work and how to perform and verify a manual switchover.

Geo redundancy switchovers

A switchover interchanges the active and standby cluster roles after a failure or by manual action.

A geo-redundancy switchover is the process of interchanging the active and standby cluster roles when a failure occurs.

Switchover behavior and considerations

The system performs preliminary checks, including heartbeat counts, connectivity, HTTP, and SSH login checks. It raises alarms when any of these checks fail.

A switchover can be automatic or manual. Automatic switchover uses the auto-arbitration setting. Before a manual switchover, verify an alarm by checking both clusters.

Note

- If a switchover operation completes on a standby VM before synchronization, tech-support job details do not appear in Publish Details. Tech-support history is written to ETCD, which is not synchronized across geo-redundancy setups. This is expected behavior.
- All services provisioned on AZ1 are continuously synchronized to AZ2 through live asynchronous replication.
- After a switchover, topology is discovered more quickly on AZ2 because it is built through resynchronization rather than from scratch.
- After a switchover, perform an on-demand synchronization. Otherwise, Postgres and Timescale data stores may display inaccurate state information.
- In a single-VM deployment, service availability may be delayed by up to 5 minutes during geo-redundancy switchover if dependent resources such as the Postgres database are unavailable. No secondary HA instance is available to retry service startup unless the service is restarted manually.
- If an application is installed after the last synchronization, synchronize again before starting a switchover.
- L2 link discovery may take up to 2 hours on a scaled system during a geo-redundancy switchover.
- After a switchover, clients using a JWT issued by the previous active cluster receive a 403 (key not authorized) error. They must obtain a new JWT from the current active cluster.

Auto-switchover triggers

Auto-switchover can follow manual intervention or failures in cluster nodes, intercluster links, or Crosswork Data Gateway services.

A switchover can be triggered by manual intervention or system failures. The following sections describe the triggers.

Switchover triggers

A switchover can be triggered by various events, including manual intervention or system failures. These are the key triggers that initiate the switchover process:

- On-demand switchover
- Cluster node failure

- Link failure
- Crosswork Data Gateway failure

On-demand switchover

This occurs when a user manually decides to trigger the switchover process.

Node failure

In a multicluster setup, the switchover is triggered when two or more hybrid nodes fail. In a single-VM setup, it is triggered when any individual VM fails.

Link failure

A switchover may occur because of intercluster link failures, such as a communication failure between the active, standby, or arbiter VMs.

Crosswork Data Gateway failure

> **Note**
>
> Not applicable in a single-VM geo setup.

Crosswork Network Controller continuously monitors and polls the operational state of Data Gateways in the active cluster. If all Data Gateways remain in the error state longer than 30 minutes or the configured arbitration window, Crosswork Network Controller triggers an auto-switchover. This also includes Data Gateways in maintenance mode when they are in the error state.

Crosswork polls Data Gateway status only at scheduled times and may miss or delay detection of failures that occur around a polling cycle. Refer to Data Gateway failure triggers considerations, conditions, and alarms for delays, considerations, and alarms.

The arbitration duration is configurable through the arbitration-window API. Refer to Cisco DevNet for API information.

Data Gateway failure triggers considerations, conditions, and alarms

Data Gateway health is evaluated at polling intervals, which can delay detection and alarms.

Use these considerations to understand Data Gateway failure detection timing, switchover conditions, and alarms.

Key considerations

- The arbitration window is 30 minutes by default.
- Data Gateway health is evaluated only at polling intervals.
- Status changes may not trigger immediate detection or switchover if they occur outside the configured polling interval and arbitration window.
- Data Gateways in maintenance mode are considered for a switchover.

Detection delays during edge cases

Crosswork evaluates Data Gateway health only during scheduled polling intervals. A switchover is triggered when the detection criteria are met, such as all Data Gateways in the active availability zone remaining in the error state throughout the arbitration window. Timing differences between state changes and polling cycles can delay detection.

Data Gateway failure just before polling

- Scenario: Data Gateways transition to an error state just before polling begins.
- Example: Data Gateways enter the error state at 9:59 p.m., and polling is scheduled for 10:00 p.m.
- Outcome: Crosswork does not consider these Data Gateways eligible for switchover because they have not met the 30-minute arbitration window when polling starts.

Data Gateway failure just after polling

- Scenario: Data Gateways transition to an error state shortly after polling begins.
- Example: Polling occurs at 10:00 p.m., and Data Gateways enter the error state at 10:05 p.m.
- Outcome: Crosswork does not detect these Data Gateways as faulty until the next polling cycle because they have not met the 30-minute arbitration window.

Alarms for Data Gateway

A critical alarm indicates that Crosswork Network Controller is about to trigger a geo switchover. Clear the alarms manually.

> **Note**
>
> Alarms are local to each availability zone (AZ) and are not shared across AZs.

How auto-arbitration works

Auto-arbitration coordinates heartbeat and failure detection across peers before triggering a switchover.

This topic describes the high-level stages of an automatic switchover. Refer to Auto-arbitration for more information.

Summary

The process uses these components:

- Clusters or availability zones: Each has two peers that independently monitor the system.
- Peers: Detect an active-cluster failure and initiate switchover requests.
- Leader: Acts as the single decision-making authority and processes the first request.

The process enables auto-arbitration, configures heartbeat and failure-detection intervals, elects a leader, and triggers switchover when an event persists beyond the failure-detection interval.

Workflow

These stages describe how auto-arbitration triggers a switchover.

1. Auto-arbitration is enabled in the arbitration settings. Refer to Configure cross-cluster settings.
2. The heartbeat and failure-detection intervals are configured.

3. The leader-election process starts, and one cluster is elected as leader.
4. When a switchover event occurs, automatic switchover is triggered if the condition persists beyond the failure-detection interval.

> **Note**
>
> When the active cluster fails, both peers detect the event and initiate a request. The leader receives both requests, processes the first, and ignores the second to prevent split-brain. All requests are logged; apparent duplicate requests are expected behavior.

Perform switchover manually

Use the Cross Cluster UI to change cluster roles, update DNS entries, synchronize data, and verify cluster status.

Use the Crosswork Network Controller UI to perform a manual switchover between active and standby clusters.

Before you begin

Ensure the active and standby clusters use the same application versions and resource footprints.

> **Note**
>
> After switchover to the standby site, Data Gateways enter the Error state. This expected behavior does not affect active data collection because these Data Gateways function as spares. They continue to run release 7.1.x until redeployed with release 7.2.x.

This topic explains how to perform a manual switchover between the active and standby clusters from the Crosswork Network Controller UI.

> **Note**
>
> Use this switchover option only if auto-arbitration is not enabled. For migration of a geo-redundant deployment from release 7.1.x to release 7.2.1 SP, do not use the UI Switchover action. Instead, use the `SetRole` API procedure in Perform switchover between sites.

Procedure

1. Log in to the standby cluster.
2. From the main menu, choose **Administration** > **Cross Cluster**. The Cross Cluster window appears.
3. Choose one of the following options to perform the switchover.
   - Choose **Actions** > **Switchover** to initiate a standard one-click switchover. This option changes cluster roles and updates DNS records. Do not use this option during migration from release 7.1.x to release 7.2.1 SP. The

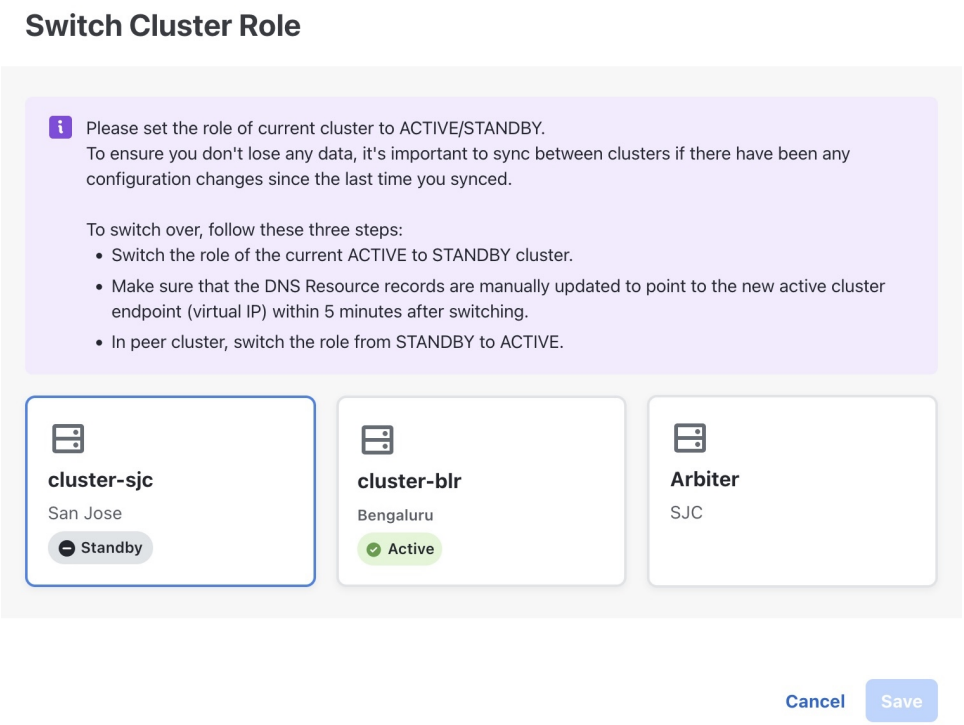
system forwards the request to the leader cluster and creates a switchover job. Monitor progress on the Cross Cluster Jobs tab.

- Choose **Actions** > **Set Cluster Role** to assign the cluster roles manually.

> **Note**
>
> This option is particularly useful if clusters become unresponsive during automatic switchover and both clusters attain the same status.

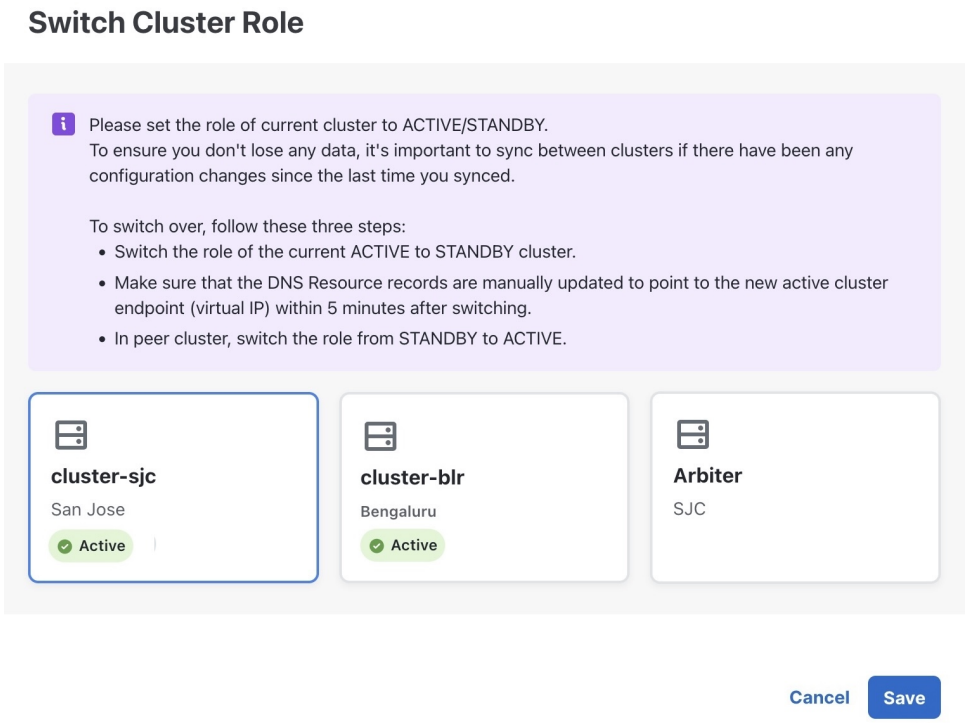
a. The Switch Cluster Role dialog shows the initial cluster states. In this example, cluster-blr is *Active* and cluster-sjc is *Standby*.

Figure 67: Switch cluster role



b. Select the San Jose cluster to change it to *Active*, and choose **Save** to confirm.

Figure 68: Switch standby cluster to active



c. Update the DNS entries for both the cluster endpoint VIP and the Data VIP to point to the new active cluster.

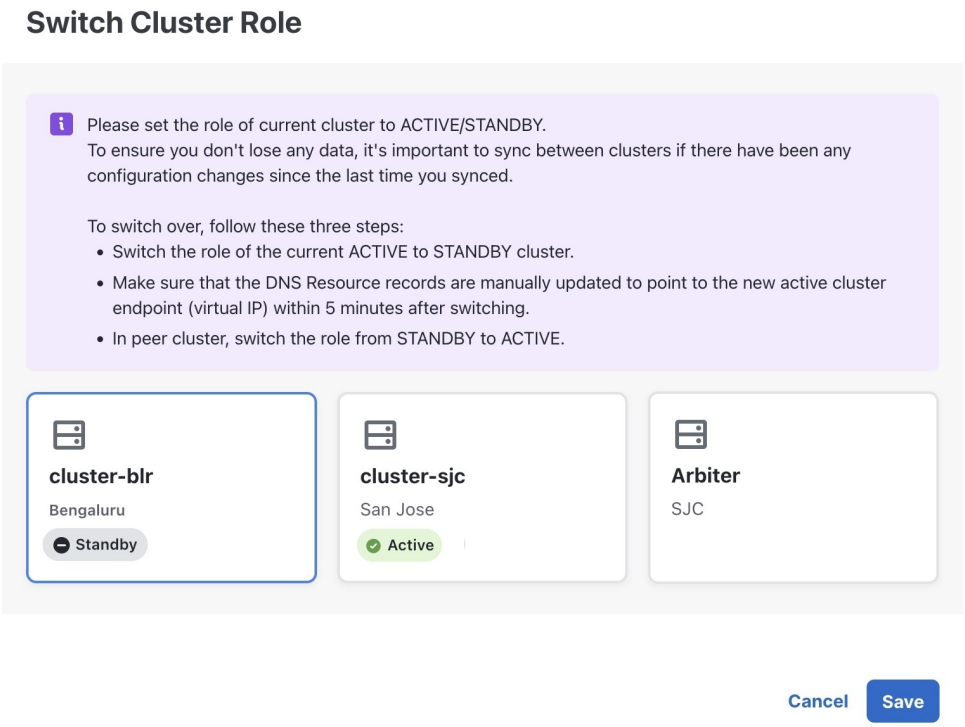
> **Note**
>
> If only the cluster endpoint VIP is updated, services that use the Data VIP may not resolve to the correct site after switchover.

d. Log in to the Bengaluru cluster, which is active. In the Cross Cluster window, choose **Actions** > **Set Cluster Role**.

> **Note**
>
> Until you change the cluster state, both clusters are *Active*.

e. In the Switch Cluster Role dialog, change the cluster to *Standby* and choose **Save**.

Figure 69: Switch active cluster to standby



Note

Wait for device reachability to converge before resuming operations on the standby cluster.

f. After a few minutes, log in to the first cluster. The switchover is complete.

4. Perform an on-demand or periodic synchronization.
5. After the switchover, verify the following:

Note

If you perform several switchovers in a short time, a brief dampening period may occur after each switchover.

a) Verify cluster health and device status.
b) Check Crosswork Data Gateway health.
c) Check the HA pool status.
d) Check collection status and confirm that traffic flows to the newly active cluster.

Crosswork Optimization Engine license count after a switchover

The Smart Licenses page reflects the correct license count after 24 hours or by 1:00 a.m. after a switchover.

For Crosswork Optimization Engine, the Smart Licenses page reflects the correct license count only after 24 hours or by 1:00 a.m. after a switchover.

Force a license update

If you cannot wait 24 hours or until 1:00 a.m., use either of these methods to force a license update:

- Disable or enable feature packs (Bandwidth on Demand, Circuit Style Manager, or Local Congestion Manager).
- Detach devices and add them back.

23 Geo Redundant Cluster Disaster Recovery

Topics:

- Disaster recovery for geo redundant clusters
- Recover the active cluster
- Recover the standby cluster
- Recover after an active cluster disaster restore

This chapter describes when a geo-redundant cluster may need disaster recovery. It provides procedures for recovering after the active or standby cluster is destroyed, and after restoring an active cluster from disaster.

This chapter explains the recovery procedures for geo-redundant clusters.

Disaster recovery for geo redundant clusters

Recover the Crosswork Network Controller system after cluster nodes, services, or applications malfunction, or a disaster destroys the cluster hosts.

Geo-redundant cluster disaster recovery is the recovery of the entire Crosswork Network Controller system after one or more cluster nodes, services, or applications malfunction, or a disaster destroys the hosts for the entire cluster.

When cluster recovery is needed

During normal cluster operations, recovery may be needed because one or more nodes, services, or applications malfunction, or because a disaster destroys the hosts for the entire cluster.

Recover the active cluster

Recover service after the active cluster is destroyed by promoting the standby cluster and rebuilding the destroyed cluster as standby.

Recover the geo-redundant deployment when the active cluster is destroyed.

Use this procedure when the active cluster is destroyed.

Procedure

1. Switch the standby cluster to become the new active cluster. If you use manual arbitration, perform the switchover manually. If you use automatic arbitration, the standby cluster becomes the new active cluster automatically. Refer to Geo redundancy switchover.
2. Redeploy the destroyed cluster, which is now the standby cluster, with the same IP and Virtual IP (VIP) addresses.
3. Install the same application versions as those on the new active cluster.
4. Enable pairing mode on the new active cluster, and update the Cross Cluster inventory file to reflect the changes after the switchover in step 1. You can update the inventory YAML file with the geo wizard. Refer to Enable geo redundancy for details.

   (Optional) You can select **View operational inventory** on the **Geo Redundancy Manager** window to access the existing YAML file.
5. Log in to the old active cluster, which is now standby, and enable geo redundancy with the updated YAML file.
6. Configure Cross Cluster settings on the old active cluster, which is now standby, to match the settings on the new active cluster.
7. In the Cross Cluster window, synchronize data so that the old active cluster, now standby, pulls data from the new active cluster.
8. (Optional) Switch the standby cluster to become the new active cluster if needed.

Recover the standby cluster

Replace a destroyed standby cluster, enable geo redundancy, and synchronize its data from the active cluster.

Recover the deployment when the standby cluster is destroyed.

Before you begin

Refer to Enable geo redundancy for detailed instructions on preparing the YAML file and enabling geo redundancy.

Use this procedure when the standby cluster is destroyed.

Procedure

1. Deploy a new standby cluster to replace the destroyed cluster, using the same IP and Virtual IP (VIP) addresses.
2. Install the same application versions as those on the active cluster.
3. Enable pairing mode on the active cluster.
4. Enable geo redundancy on the new standby cluster.
5. Enter the Cross Cluster settings on the new standby cluster to match the settings on the active cluster.
6. In the Cross Cluster window, synchronize data so that the new standby cluster pulls data from the active cluster.

Recover after an active cluster disaster restore

Use this procedure to recover correct cluster leadership when both clusters show active roles after restoring an Active cluster from disaster.

Recover the deployment and reestablish proper cluster roles when both clusters show as active after a disaster restore.

After restoring an Active cluster from disaster, both clusters may appear as active in the inventory and need reconciliation.

Before you begin

Ensure you have the inventory YAML file; if not, be prepared to build it manually in the geo wizard.

Follow these steps to recover cluster roles:

Procedure

1. Re-deploy the cluster that was destroyed.
2. Install all applications that are present on the current Active cluster.
3. Enable geo redundancy on the rebuilt cluster as Standby by importing the existing inventory YAML.
4. Verify that the leadership roles are correct: the rebuilt cluster must be Standby and the other cluster must be Active.
5. Build the YAML manually using the geo wizard if an inventory file is not available.

   Do not use the **fetch** operation.
6. Activate geo redundancy after the wizard completes a successful validation.

Both clusters now have correct Active/Standby roles and geo redundancy has been restored.

A Sample Deployment Templates

Topics:

- Crosswork Network Controller cluster deployment templates for VMware vCenter
- Crosswork Network Controller deployment templates for KVM
- Crosswork Network Controller VM templates for single VM deployments
- Crosswork Data Gateway deployment templates and scripts
- Sample auto action templates
- Geo redundancy templates

This chapter provides sample deployment manifests, installation scripts, and inventory files for Crosswork Network Controller and Crosswork Data Gateway. The examples cover VMware vCenter, KVM, single VM deployments, auto actions, and geo redundancy.

This chapter provides sample deployment manifests, installation scripts, and inventory files for Crosswork Network Controller and Crosswork Data Gateway.

Crosswork Network Controller cluster deployment templates for VMware vCenter

Provides manifest examples for Crosswork Network Controller cluster deployments on VMware vCenter, including two-host, resource-group, and dual-stack configurations.

This topic contains manifest template examples for various scenarios of Crosswork cluster deployment.

- Example 1: to deploy a cluster (3 hybrid nodes, 2 worker nodes) on 2 hosts
- Example 2: to deploy a cluster (3 hybrid, 2 workers) in a resource group
- Example 3: to deploy a cluster (3 hybrid nodes, 2 worker nodes) on dual stack configuration

> **Note**
>
> In case you are using resource pools, please note that individual ESXi host targeting is not allowed and vCenter is responsible for assigning the VMs to host(s) in the resource pool.

Example 1: to deploy a cluster (3 hybrid nodes, 2 worker nodes) on 2 hosts

This example deploys a Crosswork cluster with specified hosts.

```
#See at the end of the file for a configured sample

/********* Crosswork Cluster Data  *********/

  # The name of the Crosswork Cluster.
  ClusterName      = "CW-Cluster-01"

  # Provide  name of Cw VM image in vcenter or leave empty
  # When empty the image name will be populated from the uploaded image
  Cw_VM_Image = ""    # Line added automatically by installer.

  # The IP stack protocol: IPv4 or IPv6 or DUALSTACK
  ClusterIPStack        = "IPv4"

  # The Management Virtual IP for the cluster
  ManagementVIP     = "x.x.x.x"

  # Optional: The Management Virtual IP host-name
  ManagementVIPName = ""

  # The Management IP subnet in dotted decimal format for ipv4 or prefix length
 for ipv6
  ManagementIPNetmask = "x.x.x.x"

  # The Gateway IP on the Management Network
  ManagementIPGateway = "x.x.x.x"

  # The Data Virtual IP for the cluster. Use 0.0.0.0 or ::0 to disable
  DataVIP           = "x.x.x.x"

  # Optional: The Data Virtual IP host-name
  DataVIPName = ""
```

```
  # The Data IP subnet in dotted decimal format for ipv4 or prefix length for
 ipv6
  # Provied any regular mask when not in use
  DataIPNetmask       = "x.x.x.x"

  # The Gateway IP on the Management Network
  DataIPGateway       = "x.x.x.x"

  #  The IP address of the DNS server
  DNS                 = "x.x.x.x"

  # The domain name to use for the cluster
  DomainName            = "cisco.com"

  # Sets the cw-admin user ssh login password for all VMs in the cluster
  # The password MUST be of min length 8 and strong
  CWPassword            = "**************"

  # Sets the VM size for the cluster. The only supported option is Large.
  VMSize                = "Large"

  # NTP server address or name
  NTP                   = "x.x.x.x"

  # Configuration Manifest schema version
  SchemaVersion         = "7.2.1"

  # Data disk size for Manager/Hybrid nodes in GB. Min 485 Max 8000
  ManagerDataFsSize = 485
  # Data disk size for Worker nodes in GB. Min 485 Max 8000
  WorkerDataFsSize = 485

  // Thin or thick provisioning for all disks. Set to true for thin
provisioning, false for thick
  ThinProvisioned = false

  # Log partition size in GB. Min 20 Max 1000
  LogFsSize = 20

  # Minimum percentage of the data disk space to be used for the size of the
backup partition
  # Note: The final backup partition size will be calculated dynamically. This
 parameter defines the minimum.
  # Valid range 1 - 80
  BackupMinPercent = 35

  # Enforces VM profile reservations as "hard"
  EnableHardReservations = "True"

  # FOR DEMO USE ONLY - NOT TO BE USED IN PRODUCTION DEPLOYMENTS
  # Ram disk size in GB
  RamDiskSize           = 10

  # Vcenter or Non Vcenter deployment. false for Vcenter
  NonVcenter = false

/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 3; maximum: 3
# Number of Worker VMs minimum: 2; maximum: 3

CwVMs = {
```

```
    # Seed VMs' data.
    # IMPORTANT: A VM with id "0" MUST be present in the initial day0 install
 manifest and its role MUST be
    # set to either MASTER or HYBRID.
  "0" = {

    # This VM's name
    VMName                = "CW_Node_0",

    # This VMs' management IP address
    ManagementIPAddress = "x.x.x.x",

    # This VMs' data IP address. Use 0.0.0.0 or ::0 to disable
    DataIPAddress       = "x.x.x.x",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType               = "Hybrid",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },

   # Second VMs' data
  "1" = {
    # This VM's name
    VMName                = "CW_Node_1",

    # This VMs' management IP address
    ManagementIPAddress = "x.x.x.x",

    # This VMs' data IP address
    DataIPAddress       = "x.x.x.x",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType               = "Hybrid",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },

  # Third VMs' data
  "2" = {
    # This VM's name
    VMName                = "CW_Node_2",

    # This VMs' management IP address
    ManagementIPAddress = "x.x.x.x",

    # This VMs' data IP address
    DataIPAddress       = "x.x.x.x",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType               = "Hybrid",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },

# Worker VMs' data
```

```
  "3" = {
    # This VM's name
    VMName                = "CW_Node_3",

    # This VMs' management IP address
    ManagementIPAddress = "x.x.x.x",

    # This VMs' data IP address
    DataIPAddress       = "x.x.x.x",

    # This Cw VM's type - use "Worker"
    NodeType               = "Worker",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },
# Worker VMs' data
  "4" = {
    # This VM's name
    VMName                = "CW_Node_4",

    # This VMs' management IP address
    ManagementIPAddress = "x.x.x.x",

    # This VMs' data IP address
    DataIPAddress       = "x.x.x.x",

    # This Cw VM's type - use "Worker"
    NodeType               = "Worker",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },
}

/********* vcenter Resource Data with Cw VM assignment *********/

VCenterDC = {

  # The vcenter IP or host name
  VCenterAddress = "x.x.x.x",

  # The username to use for logging into vcenter
  VCenterUser = "Cisco_User",

  # The vcenter password for the user
  VCenterPassword = "*********",

  # The name of the Data Centre resource to use
  DCname = "Cisco-Crosswork",

  # The name of the vcenter network to attach to the Cw VM Management interface

  # NOTE: Escape any special characters using their URL escape codes, eg use
"%2f" instead of "/"
  MgmtNetworkName = "VM Network",

  # The name of the vcenter network to attach to the Cw VM Data interface.
  # Leave empty if not used.
  # NOTE: Escape any special characters using their URL escape codes, eg use
```

```
"%2f" instead of "/"
  DataNetworkName = "Crosswork-Internal",

  # The resource folder name on vcenter. Leave empty if not used.
  DCfolder = "",

  # List of the vcenter host resources along with the VMs names
  # that each that each resource will host. Add additional stanzas, separated
 by a ','
  # for each additional ESXi host or resource
  VMs = [{

    # The ESXi host, or ONLY the vcenter cluster/resource group name.
    Host = "x.x.x.x",

    # The datastore name available to be used by this host or resource group.

    Datastore = "Datastore-1",

    # The high speed datastore available for this host or resource group.
    # Set to same value as Datastore if unsure.
    HSDatastore = "Datastore-1"

   # The ids of the VMs to be hosted by the above ESXi host or resource. These
 have to match to the Cw VM
    # ids specified in the Cw VM map. Separate multiple VMs the given
    # host with a ',', eg ["0","1"].
    HostedCwVMs = ["0","1","2"]

    },
    {
    Host = "x.x.x.x"
    Datastore = "Datastore-2"
    HSDatastore = "Datastore-2"
    HostedCwVMs =["3","4"]
    }
  ]
}
```

Example 2: to deploy a cluster (3 hybrid, 2 workers) in a resource group

This example deploys a Crosswork cluster with resource groups.

```
/********* Crosswork Cluster Data  *********/


  ClusterName      = "CW-cluster-01"
  # When empty the image name will be populated from the uploaded image
  Cw_VM_Image = ""    # Line added automatically by installer.
  ClusterIPStack       = "IPv4"
  ManagementVIP     = "x.x.x.x"
  ManagementVIPName = ""
  ManagementIPNetmask = "x.x.x.x"
  ManagementIPGateway = "x.x.x.x"
  DataVIP           = "x.x.x.x"
  DataVIPName = ""
  DataIPNetmask       = "x.x.x.x"
  DataIPGateway       = "x.x.x.x"
```

```
  DNS                 = "x.x.x.x"
  DomainName            = "cisco.com"

  # Kubernetes Service Network Customization - The default network '10.96.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sServiceNetwork = "10.96.0.0"

  # Kubernetes Service Network Customization - The default network '10.244.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sPodNetwork = "10.244.0.0"


  CWPassword            = "*********"
  VMSize                = "Large"
  NTP                   = "x.x.x.x"
  SchemaVersion         = "7.2.1"

  # Data disk size for Manager/Hybrid nodes in GB. Min 485 Max 8000
  ManagerDataFsSize = 485
  # Data disk size for Worker nodes in GB. Min 485 Max 8000
  WorkerDataFsSize = 485

  // Thin or thick provisioning for all disks. Set to true for thin
provisioning, false for thick
  ThinProvisioned = false

  # Log partition size in GB. Min 20 Max 1000
  LogFsSize = 20

  # Minimum percentage of the data disk space to be used for the size of the
backup partition
  # Note: The final backup partition size will be calculated dynamically. This
 parameter defines the minimum.
  # Valid range 1 - 80
  BackupMinPercent = 35

  # Enforces VM profile reservations as "hard"
  EnableHardReservations = "False"

  # FOR DEMO USE ONLY - NOT TO BE USED IN PRODUCTION DEPLOYMENTS
  # Ram disk size in GB
  RamDiskSize           = 0

  # Pods that are marked as skip auto install will not be brought up until a
dependent application/pod explicitly asks for it
  EnableSkipAutoInstallFeature = "False"

  # DEMO/DEV USE ONLY - Enforce pod minimum resource reservations. Default and
 for production use is True
  EnforcePodReservations = "True"

  # Optional: Provide a standard IANA time zone. Default value is Etc/UTC if
not specified
  Timezone = ""

  # Vcenter or Non Vcenter deployment. false for Vcenter
  NonVcenter = false
```

```
/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 3; maximum: 3
# Number of Worker VMs minimum: 0; maximum: 3

CwVMs = {
  "0" = {
    VMName                = "cw-vm-0",
    ManagementIPAddress = "x.x.x.x",
    DataIPAddress       = "x.x.x.x",
    NodeType               = "Hybrid",
    #Op_Status = 2
  },
  "1" = {
    VMName                = "cw-vm-1",
    ManagementIPAddress = "x.x.x.x",
    DataIPAddress       = "x.x.x.x",
    NodeType               = "Hybrid",
    #Op_Status = 2
  },
  "2" = {
    VMName                = "cw-vm-2",
    ManagementIPAddress = "x.x.x.x",
    DataIPAddress       = "x.x.x.x",
    NodeType               = "Hybrid",
    #Op_Status = 2
  },
  "3" = {
    # This VM's name
    VMName                = "cw-worker-3",
    ManagementIPAddress = "x.x.x.x",
    DataIPAddress       = "x.x.x.x",
    NodeType               = "Worker",


    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    #Op_Status = 2
  },
  "4" = {
    # This VM's name
    VMName                = "cw-worker-4",
    ManagementIPAddress = "x.x.x.x",
    DataIPAddress       = "x.x.x.x",
    NodeType               = "Worker",
    #Op_Status = 2
  }
}



/********* vcenter Resource Data with Cw VM assignment *********/

VCenterDC = {
  VCenterAddress = "x.x.x.x",
  VCenterUser = "Cisco_User",
  VCenterPassword = "***********",
  DCname = "Cisco-Crosswork",
  MgmtNetworkName = "Management Network",
  DataNetworkName = "Data Network",
  DCfolder = ""
  VMs = [{
```

```
    Host = "{path to resource Group}",

    Datastore = "iSCSI-DataStore",

    HSDatastore = "iSCSI-DataStore",

    HostedCwVMs = ["0","1","2","3","4"],

    }
  ]
}
```

Example 3: to deploy a cluster (3 hybrid nodes, 2 worker nodes) on dual stack configuration

This example deploys a Crosswork cluster containing 3 Hybrid nodes (IDs 0,1, 2) and 2 worker nodes (IDs 3, 4) on a dual stack configuration.

```
/********* Crosswork Cluster Data  *********/

  ClusterName      = "geo-dualstack-setup"
  Cw_VM_Image = "cw-na-platform-release-7.2.1"    # Line added automatically
by installer.
  ClusterIPStack        = "DUALSTACK"
  ManagementVIPv4     = "x.x.x.x"
  ManagementVIPv6     = "2001:420:54ff:13::485:10"
  ManagementVIPv6     = "x:x:x:x::x:x"
  ManagementIPv4Netmask = "x.x.x.x"
  ManagementIPv4Gateway = "x.x.x.x"
  ManagementIPv6Netmask = "119"
  ManagementIPv6Gateway = "x:x:x:x::x:x"
  DataVIPv6           = "fd00::101:1a"
  DataVIPv4           = "x.x.x.x"
  DataIPv6Netmask       = "121"
  DataIPv6Gateway       = "fd00::101:1"
  DataIPv4Netmask       = "x.x.x.x"
  DataIPv4Gateway       = "x.x.x.x"
  DNSv6                 = "2001:x:284:2004:4:x:122:5058 x::x:x:x:x:x"
  DNSv4                 = "x.x.x.x"
  DomainName            = "cw.cisco"
  CWPassword            = "**************"
  VMSize                = "Large"
  NTP                   = "ntp.esl.cisco.com"
  SchemaVersion                    = "5.0.0"
  ThinProvisioned = false
  EnableHardReservations = "False"
  Timezone = "Asia/Kolkata"
  NonVcenter = false
/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 3; maximum: 3
# Number of Worker VMs minimum: 0; maximum: 3

CwVMs = {
    # Seed VMs' data.
    # NOTE: A VM with id "0" MUST be present in the initial day0 install
manifest and it's role MUST be
    # set to either MASTER or HYBRID.
  "0" = {
    VMName               = "cw-hybrid1-geo2-az2",
```

```
    ManagementIPv6Address = "x:x:x:x::x:x",
    DataIPv6Address       = "x::x:x",
    ManagementIPv4Address = "x.x.x.x",
    DataIPv4Address       = "x.x.x.x",
    NodeType               = "Hybrid"
  },

   # Second VMs' data
  "1" = {
    VMName                = "cw-hybrid2-geo2-az2",
    ManagementIPv6Address = "x:x:x:x::x:x",
    DataIPv6Address       = "x::x:x",
    ManagementIPv4Address = "x.x.x.x",
    DataIPv4Address       = "x.x.x.x",
    NodeType               = "Hybrid"
  },

  "2" = {
    VMName                = "cw-hybrid3-geo2-az2",
    ManagementIPv6Address = "x:x:x:x::x:x",
    DataIPv6Address       = "x::x:x",
    ManagementIPv4Address = "x.x.x.x",
    DataIPv4Address       = "x.x.x.x",
    NodeType               = "Hybrid"
  },

  "3" = {
    VMName                = "cw-worker1-geo2-az2",
    ManagementIPv6Address = "x:x:x:x::x:x",
    DataIPv6Address       = "x::x:x",
    ManagementIPv4Address = "x.x.x.x",
    DataIPv4Address       = "x.x.x.x",
    NodeType               = "Worker"
  },

  "4" = {
    VMName                = "cw-worker2-geo2-az2",
    ManagementIPv6Address = "x:x:x:x::x:x",
    DataIPv6Address       = "x::x:x",
    ManagementIPv4Address = "x.x.x.x",
    DataIPv4Address       = "x.x.x.x",
    NodeType               = "Worker"
  }
 }

/********* vcenter Resource Data with Cw VM assignment *********/

VCenterDC = {

  VCenterAddress = "x.x.x.x",
  VCenterUser = "Cisco_User",
  VCenterPassword = "*********",
  DCname = "Crosswork-LAB",
  MgmtNetworkName = "x.x.x.x",
  DataNetworkName = "DATA-VLAN-101",
  DCfolder = "",
  VMs = [{
    HostedCwVMs = ["0"],
    Host = "x.x.x.x",
    Datastore = "3.3TB-SSD-165-52",
    HSDatastore = "3.3TB-SSD-165-52"
  },
```

```
  {
    HostedCwVMs = ["1"],
    Host = "x.x.x.x",
    Datastore = "3.3TB-SSD-165-52",
    HSDatastore = "3.3TB-SSD-165-52"
  },
  {
    HostedCwVMs = ["2"],
    Host = "x.x.x.x",
    Datastore = "3.3TB-SSD-165-52",
    HSDatastore = "3.3TB-SSD-165-52"
  },
  {
    HostedCwVMs = ["3"],
    Host = "x.x.x.x",
    Datastore = "3.3TB-SSD-165-52",
    HSDatastore = "3.3TB-SSD-165-52"
  },
  {
    HostedCwVMs = ["4"],
    Host = "x.x.x.x",
    Datastore = "3.3TB-SSD-165-52",
    HSDatastore = "3.3TB-SSD-165-52"
  }
 ]
}

IgnoreDiagnosticsCheckFailure = "True"    # Line added automatically by
installer.
```

Crosswork Network Controller deployment templates for KVM

Provides KVM deployment templates and sample inventory files for Crosswork Network Controller cluster and single VM deployments.

This topic provides Crosswork Network Controller deployment and inventory templates for KVM.

Sample script

- Example 1: OVF template (.xml) to deploy Crosswork Network Controller on KVM (cluster deployment)
- Example 2: OVF template (.xml) to deploy Crosswork Network Controller VM on KVM (single VM deployment)
- Example 3: sample .tfvars file for dual stack inventory import (single VM deployment)
- Example 4: sample .tfvars file for IPv6 inventory import (single VM deployment)
- Example 5: sample .tfvars file for IPv4 inventory import (single VM deployment)

Note

When manually importing the inventory for KVM deployments, ensure that the **NonVcenter** flag is set to **true** in the `.tfvars` file.

Example 1: OVF template (.xml) to deploy Crosswork Network Controller on KVM (cluster deployment)

```
<Environment
     xmlns="http://schemas.dmtf.org/ovf/environment/1"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xmlns:oe="http://schemas.dmtf.org/ovf/environment/1"
     xmlns:ve="http://www.vmware.com/schema/ovfenv"
     oe:id="">
   <PlatformSection>
      <Kind>KVM</Kind>
      <Version>7.2.1</Version>
      <Vendor>KVM</Vendor>
      <Locale>en</Locale>
   </PlatformSection>
    <PropertySection>
         <Property oe:key="CWPassword" oe:value="**********"/>
         <Property oe:key="CWUsername" oe:value="cw-admin"/>
         <Property oe:key="CwInstaller" oe:value="False"/>
         <Property oe:key="DNSv4" oe:value="171.70.168.183"/>
         <Property oe:key="DNSv6" oe:value="::0"/>
         <Property oe:key="DataIPv4Address" oe:value="192.168.5.43"/>
         <Property oe:key="DataIPv4Gateway" oe:value="192.168.5.1"/>
         <Property oe:key="DataIPv4Netmask" oe:value="255.255.255.0"/>
         <Property oe:key="DataIPv6Address" oe:value="::0"/>
         <Property oe:key="DataIPv6Gateway" oe:value="::1"/>
         <Property oe:key="DataIPv6Netmask" oe:value="64"/>
         <Property oe:key="DataVIP" oe:value="192.168.5.42"/>
         <Property oe:key="DataVIPName" oe:value=""/>
         <Property oe:key="Deployment" oe:value="cw_ipv4"/>
         <Property oe:key="Disclaimer" oe:value="Cisco Crosswork"/>
         <Property oe:key="Domain" oe:value="cisco.com"/>
         <Property oe:key="EnableSkipAutoInstallFeature" oe:value="False"/>
         <Property oe:key="EnforcePodReservations" oe:value="True"/>
         <Property oe:key="IgnoreDiagnosticsCheckFailure" oe:value="True"/>
         <Property oe:key="InitMasterCount" oe:value="3"/>
         <Property oe:key="InitNodeCount" oe:value="5"/>
         <Property oe:key="IsSeed" oe:value="True"/>
         <Property oe:key="K8Orch" oe:value=""/>
         <Property oe:key="K8sPodNetworkV4" oe:value="10.244.0.0"/>
         <Property oe:key="K8sServiceNetworkV4" oe:value="10.96.0.0"/>
         <Property oe:key="ManagementIPv4Address" oe:value="10.19.70.143"/>
         <Property oe:key="ManagementIPv4Gateway" oe:value="10.19.70.1"/>
         <Property oe:key="ManagementIPv4Netmask" oe:value="255.255.255.0"/>
         <Property oe:key="ManagementIPv6Address" oe:value="::0"/>
         <Property oe:key="ManagementIPv6Gateway" oe:value="::1"/>
         <Property oe:key="ManagementIPv6Netmask" oe:value="112"/>
         <Property oe:key="ManagementVIP" oe:value="10.19.70.142"/>
         <Property oe:key="ManagementVIPName" oe:value=""/>
         <Property oe:key="NTP" oe:value="ntp.esl.cisco.com"/>
         <Property oe:key="Timezone" oe:value="US/Pacific"/>
         <Property oe:key="VMLocation" oe:value="default"/>
         <Property oe:key="VMType" oe:value="Hybrid"/>
```

```
         <Property oe:key="bckup_min_percent" oe:value="35"/>
         <Property oe:key="corefs" oe:value="18"/>
         <Property oe:key="ddatafs" oe:value="485"/>
         <Property oe:key="logfs" oe:value="20"/>
         <Property oe:key="ramdisk" oe:value="0"/>
         <Property oe:key="ssd" oe:value="15"/>
         <Property oe:key="VMSize" oe:value="Large"/>
         <Property oe:key="ThinProvisioned" oe:value="False"/>
         <Property oe:key="UseNonDefaultCalicoBgpPort" oe:value="False"/>
        <Property oe:key="bootOptions.efiSecureBootEnabled" oe:value="True"/>

   </PropertySection>
</Environment>
```

Example 2: OVF template (.xml) to deploy Crosswork Network Controller VM on KVM (single VM deployment)

```
<Environment
     xmlns="http://schemas.dmtf.org/ovf/environment/1"
     xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
     xmlns:oe="http://schemas.dmtf.org/ovf/environment/1"
     xmlns:ve="http://www.vmware.com/schema/ovfenv"
     oe:id="">
   <PlatformSection>
      <Kind>KVM</Kind>
      <Version>7.2.1</Version>
      <Vendor>KVM</Vendor>
      <Locale>en</Locale>
   </PlatformSection>
    <PropertySection>
         <Property oe:key="CWPassword" oe:value="**********"/>
         <Property oe:key="CWUsername" oe:value="cw-admin"/>
         <Property oe:key="CwInstaller" oe:value="False"/>
         <Property oe:key="DNSv4" oe:value="171.70.168.183"/>
         <Property oe:key="DNSv6" oe:value="::0"/>
         <Property oe:key="DataIPv4Address" oe:value="192.168.5.48"/>
         <Property oe:key="DataIPv4Gateway" oe:value="192.168.5.1"/>
         <Property oe:key="DataIPv4Netmask" oe:value="255.255.255.0"/>
         <Property oe:key="DataIPv6Address" oe:value="::0"/>
         <Property oe:key="DataIPv6Gateway" oe:value="::1"/>
         <Property oe:key="DataIPv6Netmask" oe:value="64"/>
         <Property oe:key="DataVIP" oe:value="192.168.5.51"/>
         <Property oe:key="DataVIPName" oe:value=""/>
         <Property oe:key="Deployment" oe:value="cw_ipv4"/>
         <Property oe:key="Disclaimer" oe:value="Cisco Crosswork"/>
         <Property oe:key="Domain" oe:value="cisco.com"/>
         <Property oe:key="EnableSkipAutoInstallFeature" oe:value="True"/>
         <Property oe:key="EnforcePodReservations" oe:value="True"/>
         <Property oe:key="IgnoreDiagnosticsCheckFailure" oe:value="True"/>
         <Property oe:key="InitMasterCount" oe:value="1"/>
         <Property oe:key="InitNodeCount" oe:value="1"/>
         <Property oe:key="IsSeed" oe:value="True"/>
         <Property oe:key="K8Orch" oe:value=""/>
         <Property oe:key="K8sPodNetworkV4" oe:value="10.244.0.0"/>
         <Property oe:key="K8sServiceNetworkV4" oe:value="10.96.0.0"/>
         <Property oe:key="ManagementIPv4Address" oe:value="10.19.70.148"/>
         <Property oe:key="ManagementIPv4Gateway" oe:value="10.19.70.1"/>
         <Property oe:key="ManagementIPv4Netmask" oe:value="255.255.255.0"/>
         <Property oe:key="ManagementIPv6Address" oe:value="::0"/>
         <Property oe:key="ManagementIPv6Gateway" oe:value="::1"/>
         <Property oe:key="ManagementIPv6Netmask" oe:value="112"/>
         <Property oe:key="ManagementVIP" oe:value="10.19.70.151"/>
```

```
         <Property oe:key="ManagementVIPName" oe:value=""/>
         <Property oe:key="NTP" oe:value="ntp.esl.cisco.com"/>
         <Property oe:key="Timezone" oe:value="US/Pacific"/>
         <Property oe:key="VMLocation" oe:value="default"/>
         <Property oe:key="VMType" oe:value="Hybrid"/>
         <Property oe:key="bckup_min_percent" oe:value="35"/>
         <Property oe:key="corefs" oe:value="18"/>
         <Property oe:key="ddatafs" oe:value="485"/>
         <Property oe:key="logfs" oe:value="20"/>
         <Property oe:key="ramdisk" oe:value="0"/>
         <Property oe:key="ssd" oe:value="15"/>
         <Property oe:key="VMSize" oe:value="XLarge"/>
         <Property oe:key="ThinProvisioned" oe:value="False"/>
         <Property oe:key="UseNonDefaultCalicoBgpPort" oe:value="False"/>
        <Property oe:key="bootOptions.efiSecureBootEnabled" oe:value="True"/>

   </PropertySection>
</Environment>
```

Example 3: sample .tfvars file for dual stack inventory import (single VM deployment)

```
/*****************************************
* Cw SVM deployment input data TEMPLATE *
*              EDIT BEFORE USE           *
*****************************************
#See at the end of the file for a configured sample

/********* Crosswork SVM Data  *********/

  # The name of the Crosswork SVM.
  # ClusterName      = "day0-cluster"

  # Provide  name of Cw VM image in vcenter or leave empty
  # When empty the image name will be populated from the uploaded image
  Cw_VM_Image           = "Give_the_Image_Name"

  # The IP stack protocol: IPv4 or IPv6 or DUALSTACK
  ClusterIPStack        = "DUALSTACK"

  ManagementVIPv4     = "172.29.77.174"
  ManagementVIPv6     = "2001:420:284:2004:4:164:77:1574"

  # Optional: The Management Virtual IP host-name
  ManagementVIPName = ""

  ManagementIPv4Netmask = ""
  ManagementIPv6Netmask = ""

  ManagementIPv4Gateway = ""
  ManagementIPv6Gateway = ""

  DataVIPv4     = ""
  DataVIPv6     = ""


  # Optional: The Data Virtual IP host-name
  DataVIPName = ""

  DataIPv4Netmask       = ""
  DataIPv6Netmask       = ""
```

```
  DataIPv4Gateway       = ""
  DataIPv6Gateway       = ""

  DNSv4                 = ""
  DNSv6                 = ""

  # The domain name to use for the cluster
  DomainName            = "cisco.com"

  # Kubernetes Service Network Customization - The default network '10.96.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sServiceNetwork = "10.96.0.0"

  # Kubernetes Service Network Customization - The default network '10.244.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sPodNetwork = "10.244.0.0"

  # Sets the cw-admin user ssh login password for all VMs in the cluster
  # The password MUST be of min length 8 and strong
  CWPassword            = "Password123!"

  # Sets the VM size for the cluster.
  # Options are: Small | Large | XLarge.
  VMSize                = "Large"

  # NTP server address or name
  NTP                   = ""

  # Vcenter or Non Vcenter deployment. false for Vcenter
  NonVcenter = true

  # Configuration Manifest schema version
  SchemaVersion                   = "7.2.1"

  # Data disk size for Manager/Hybrid nodes in GB. Min 450 Max 8000
  ManagerDataFsSize = 485
  # Data disk size for Worker nodes in GB. Min 450 Max 8000
  WorkerDataFsSize = 485

  // Thin or thick provisioning for all disks. Set to true for thin
provisioning, false for thick
  ThinProvisioned = false

  # Log partition size in GB. Min 10 Max 1000
  LogFsSize = 20

  # Minimum percentage of the data disk space to be used for the size of the
backup partition
  # Note: The final backup partition size will be calculated dynamically. This
 parameter defines the minimum.
  # Valid range 1 - 80
  BackupMinPercent = 35

  # Enforces VM profile reservations as "hard"
  EnableHardReservations = true

  # FOR DEMO USE ONLY - NOT TO BE USED IN PRODUCTION DEPLOYMENTS
  # Ram disk size in GB
```

```
  RamDiskSize           = 0

  # Pods that are marked as skip auto install will not be brought up until a
dependent application/pod explicitly asks for it
  #Refer product documentation for picking the correct value
  EnableSkipAutoInstallFeature = "False"

  # DEMO/DEV USE ONLY - Enforce pod minimum resource reservations. Default and
 for production use is True
  EnforcePodReservations = "True"

  # Optional: Provide a standard IANA time zone. Default value is Etc/UTC if
not specified
  Timezone = ""

/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 1; maximum: 3
# Number of Worker VMs minimum: 0; maximum: 3

CwVMs = {
    # Seed VMs' data.
    # IMPORTANT: A VM with id "0" MUST be present in the initial day0 install
 manifest and its role MUST be
    # set to either MASTER or HYBRID.
  "0" = {

    # This VM's name
    VMName                = "App-Hybrid-Node-1",

    # This VMs' management IP address,To be filled only incase of
SingleStack(IPv4/IPv6)
    ManagementIPv4Address = "172.29.77.163",
    ManagementIPv6Address = "2001:420:284:2004:4:164:77:1563",

    DataIPv4Address       = "192.168.5.163",
    DataIPv6Address       = "fded:1bc1:fc3e:96d0:192:168:5:1563",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType               = "Hybrid",

    # The state for this VM; 2 = running. Only uncomment when doing a manual
inventory import
    Op_Status = 2
  },
}
```

Example 4: sample .tfvars file for IPv6 inventory import (single VM deployment)

```
/*******************************************
* Cw SVM deployment input data TEMPLATE *
*              EDIT BEFORE USE              *
*******************************************
#See at the end of the file for a configured sample

/********* Crosswork SVM Data  *********/

  # The name of the Crosswork cluster.
  ClusterName      = "day0-cluster"

  # Provide  name of Cw VM image in vcenter or leave empty
```

```
 # When empty the image name will be populated from the uploaded image
 Cw_VM_Image           = "Give_the_Image_Name"

 ClusterIPStack        = "IPv6"

 # The Management Virtual IP for the cluster
 ManagementVIP     = ""

 # Optional: The Management Virtual IP host-name
 ManagementVIPName = ""

 # The Management IP subnet in dotted decimal format for ipv4 or prefix length
for ipv6
 ManagementIPNetmask = ""

 # The Gateway IP on the Management Network
 ManagementIPGateway = ""

 # The Data Virtual IP. Use 0.0.0.0 or ::0 to disable data interface.
 ## To be filled only incase of SingleStack(IPv4/IPv6)
 DataVIP           = ""

 # Optional: The Data Virtual IP host-name
 DataVIPName = ""

 # The Data IP subnet in dotted decimal format for ipv4 or prefix length for
ipv6
 # Provide any regular mask when not in use
 DataIPNetmask       = ""

 # The Gateway IP on the Management Network
 DataIPGateway       = ""

 #  The IP address of the DNS server
 DNS                 = ""

 # The domain name to use for the cluster
 DomainName            = ""

 # Kubernetes Service Network Customization - The default network '10.96.0.0'.

 # NOTE: The CIDR range is fixed '/16', no need to enter.
 #       Only IPv4 is supported, IPv6 customization is NOT supported.
 K8sServiceNetwork = "10.96.0.0"

 # Kubernetes Service Network Customization - The default network '10.244.0.0'.

 # NOTE: The CIDR range is fixed '/16', no need to enter.
 #       Only IPv4 is supported, IPv6 customization is NOT supported.
 K8sPodNetwork = "10.244.0.0"

 # Sets the cw-admin user ssh login password for all VMs in the cluster
 # The password MUST be of min length 8 and strong
 CWPassword            = "Password123!"

 # Sets the VM size
 # Options are: Small | Large | XLarge.
 VMSize                = ""

 # NTP server address or name
 NTP                   = ""
```

```
  # Vcenter or Non Vcenter deployment. false for Vcenter
  NonVcenter = true

  # Configuration Manifest schema version
  SchemaVersion                   = "7.2.1"

  # Data disk size for Manager/Hybrid nodes in GB. Min 450 Max 8000. The
specified value must be within the storage capacity defined for the deployment.

  ManagerDataFsSize = 485

  // Thin or thick provisioning for all disks. Set to true for thin
provisioning, false for thick
  ThinProvisioned = false

  # Log partition size in GB. Min 10 Max 1000. The specified value must be
within the storage capacity defined for the deployment.
  LogFsSize = 20

  # Minimum percentage of the data disk space to be used for the size of the
backup partition
  # Note: The final backup partition size will be calculated dynamically. This
 parameter defines the minimum.
  # Valid range 1 - 80
  BackupMinPercent = 35

  # Enforces VM profile reservations as "hard"
  EnableHardReservations = true

  # FOR DEMO USE ONLY - NOT TO BE USED IN PRODUCTION DEPLOYMENTS
  # Ram disk size in GB
  RamDiskSize           = 0

  # Pods that are marked as skip auto install will not be brought up until a
dependent application/pod explicitly asks for it
  #Refer product documentation for picking the correct value
  EnableSkipAutoInstallFeature = "True"

  # DEMO/DEV USE ONLY - Enforce pod minimum resource reservations. Default and
 for production use is True
  EnforcePodReservations = "True"

  # Optional: Provide a standard IANA time zone. Default value is Etc/UTC if
not specified
  Timezone = ""

/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 1; maximum: 3
# Number of Worker VMs minimum: 0; maximum: 3

CwVMs = {
    # Seed VMs' data.
    # IMPORTANT: A VM with id "0" MUST be present in the initial day0 install
 manifest and its role MUST be
    # set to either MASTER or HYBRID.
  "<0>" = {

    # This VM's name
    VMName                = "Give_A_Name",

    # This VMs' management IP address,To be filled only incase of
```

```
SingleStack(IPv4/IPv6)
    ManagementIPAddress = "",

    # This VMs' data IP address. Use 0.0.0.0 or ::0 to disable data interface.

    ## To be filled only incase of SingleStack(IPv4/IPv6)
    DataIPAddress       = "",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType              = "Hybrid",

    # The state for this VM; 0 => 'Initializing', 1 => 'Down', 2 => 'Running',
 3 => 'Degraded', 4 => 'Erased', 5 => 'Failed', 6 => 'Suspended', 7 =>
'Configuring', 8 => 'Erasing'.
    # Uncomment 'Op_Status = <value>' below only when performing a manual
inventory import.
    # Example: setting to Running
    # Op_Status = 2
  }
}
```

Example 5: sample .tfvars file for IPv4 inventory import (single VM deployment)

```
/*******************************************
* Cw SVM deployment input data TEMPLATE *
*              EDIT BEFORE USE              *
*******************************************
#See at the end of the file for a configured sample

/********* Crosswork SVM Data  *********/

  # The name of the Crosswork cluster.
  ClusterName      = "day0-cluster"

  # Provide  name of Cw VM image in vcenter or leave empty
  # When empty the image name will be populated from the uploaded image
  Cw_VM_Image           = "Give_the_Image_Name"

  # The IP stack protocol: IPv4 or IPv6 or DUALSTACK
  ClusterIPStack        = "IPv4"

  # The Management Virtual IP for the cluster
  ManagementVIP     = ""

  # Optional: The Management Virtual IP host-name
  ManagementVIPName = ""

  # The Management IP subnet in dotted decimal format for ipv4 or prefix length
 for ipv6
  ManagementIPNetmask = ""

  # The Gateway IP on the Management Network
  ManagementIPGateway = ""

  # The Data Virtual IP. Use 0.0.0.0 or ::0 to disable data interface.
  ## To be filled only incase of SingleStack(IPv4/IPv6)
  DataVIP           = ""

  # Optional: The Data Virtual IP host-name
  DataVIPName = ""
```

```
  # The Data IP subnet in dotted decimal format for ipv4 or prefix length for
 ipv6
  # Provide any regular mask when not in use
  DataIPNetmask       = ""

  # The Gateway IP on the Management Network
  DataIPGateway       = ""

  #  The IP address of the DNS server
  DNS                 = ""

  # The domain name to use for the cluster
  DomainName            = ""

  # Kubernetes Service Network Customization - The default network '10.96.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sServiceNetwork = "10.96.0.0"

  # Kubernetes Service Network Customization - The default network '10.244.0.0'.

  # NOTE: The CIDR range is fixed '/16', no need to enter.
  #       Only IPv4 is supported, IPv6 customization is NOT supported.
  K8sPodNetwork = "10.244.0.0"

  # Sets the cw-admin user ssh login password for all VMs in the cluster
  # The password MUST be of min length 8 and strong
  CWPassword            = "Password123!"

  # Sets the VM size
  # Options are: Small | Large | XLarge.
  VMSize                = ""

  # NTP server address or name
  NTP                   = ""

  # Vcenter or Non Vcenter deployment. false for Vcenter
  NonVcenter = true

  # Configuration Manifest schema version
  SchemaVersion                   = "7.2.1"

  # Data disk size for Manager/Hybrid nodes in GB. Min 450 Max 8000. The
specified value must be within the storage capacity defined for the deployment.

  ManagerDataFsSize = 485

  // Thin or thick provisioning for all disks. Set to true for thin
provisioning, false for thick
  ThinProvisioned = false

  # Log partition size in GB. Min 10 Max 1000. The specified value must be
within the storage capacity defined for the deployment.
  LogFsSize = 20

  # Minimum percentage of the data disk space to be used for the size of the
backup partition
  # Note: The final backup partition size will be calculated dynamically. This
 parameter defines the minimum.
  # Valid range 1 - 80
  BackupMinPercent = 35
```

```
  # Enforces VM profile reservations as "hard"
  EnableHardReservations = true

  # FOR DEMO USE ONLY - NOT TO BE USED IN PRODUCTION DEPLOYMENTS
  # Ram disk size in GB
  RamDiskSize           = 0

  # Pods that are marked as skip auto install will not be brought up until a
dependent application/pod explicitly asks for it
  #Refer product documentation for picking the correct value
  EnableSkipAutoInstallFeature = "True"

  # DEMO/DEV USE ONLY - Enforce pod minimum resource reservations. Default and
 for production use is True
  EnforcePodReservations = "True"

  # Optional: Provide a standard IANA time zone. Default value is Etc/UTC if
not specified
  Timezone = ""

/********* Crosswork VM Data Map *********/
# Configure named entries for each Cw VM.
# Number of Hybrid VMs minimum: 1; maximum: 3
# Number of Worker VMs minimum: 0; maximum: 3

CwVMs = {
    # Seed VMs' data.
    # IMPORTANT: A VM with id "0" MUST be present in the initial day0 install
 manifest and its role MUST be
    # set to either MASTER or HYBRID.
  "<0>" = {

    # This VM's name
    VMName                = "Give_A_Name",

    # This VMs' management IP address,To be filled only incase of
SingleStack(IPv4/IPv6)
    ManagementIPAddress = "",

    # This VMs' data IP address. Use 0.0.0.0 or ::0 to disable data interface.

    ## To be filled only incase of SingleStack(IPv4/IPv6)
    DataIPAddress       = "",

    # This Cw VM's type - use "Hybrid" for initial install
    NodeType               = "Hybrid",

    # The state for this VM; 0 => 'Initializing', 1 => 'Down', 2 => 'Running',
 3 => 'Degraded', 4 => 'Erased', 5 => 'Failed', 6 => 'Suspended', 7 =>
'Configuring', 8 => 'Erasing'.
    # Uncomment 'Op_Status = <value>' below only when performing a manual
inventory import.
    # Example: setting to Running
    # Op_Status = 2
  }
}
```

Crosswork Network Controller VM templates for single VM deployments

Provides dual-stack templates for deploying a single Crosswork Network Controller VM with the OVF tool or Docker installer.

This topic provides templates for deploying Crosswork Network Controller on a single VM.

Sample script

- Example 1: dual stack template to deploy using OVF tool (.sh file)
- Example 2: dual stack template to deploy using docker installer (.tfvars file)

Example 1: dual stack template to deploy using OVF tool (.sh file)

```
#!/usr/bin/env bash
SVM_OVA_PATH=$1
Host="10.104.52.105"
VCENTER_LOGIN="admin@vsphere.local@x.x.x.x"
VCENTER_PATH="SVM/host"

/root/user/cwinstall/ovftool/ovftool --acceptAllEulas --skipManifestCheck
--X:injectOvfEnv -ds="10.104.52.75-ssd" \
--numberOfCpus:"*"=24 --viCpuResource=:50000: \
--memorySize:"*"=131072 --viMemoryResource=:131072: \
--diskMode="thin" --overwrite --powerOffTarget --powerOn --noSSLVerify \
--allowExtraConfig \
--deploymentOption="cw_dual" \
--prop:"NTP=ntp.esl.cisco.com" \
--prop:"Timezone=Asia/Kolkata" \
--prop:"EnforcePodReservations=True" \
--prop:"EnableSkipAutoInstallFeature=True" \
--prop:"Domain=cisco.com" \
--prop:"Disclaimer=Managed by SVM Infra Test" \
--name="cwi-svm-cnc-adv-test" \
--net:"Data Network=DPortGroup10" \
--net:"Management Network=VM Network" \
--prop:"ManagementVIPName=cwi-svm-cnc-adv-1" \
--prop:"ManagementVIPv4=x.x.x.x" \
--prop:"ManagementVIPv6=x:420:x:24::x:87" \
--prop:"ManagementIPv4Address=x.x.x.x" \
--prop:"ManagementIPv6Address=x:420:x:24::x:88" \
--prop:"ManagementIPv4Netmask=255.255.255.128" \
--prop:"ManagementIPv6Netmask=112" \
--prop:"ManagementIPv4Gateway=x.x.x.x" \
--prop:"ManagementIPv6Gateway=x:420:x:24::656:1" \
--prop:"DataVIPv4=x.1.1.x" \
--prop:"DataVIPv6=fded:x:fc3e:x:20:1:1:87" \
--prop:"DataIPv4Address=x.1.1.x" \
--prop:"DataIPv6Address=fded:x:fc3e:x:20:1:1:88" \
--prop:"DataIPv4Netmask=255.0.0.0" \
--prop:"DataIPv6Netmask=112" \
--prop:"DataIPv4Gateway=x.1.x.1" \
--prop:"DataIPv6Gateway=fded:x:fc3e:x:20:1:1:1" \
--prop:"DNSv4=x.x.x.183" \
--prop:"DNSv6=2001:420:200:1::a" \
--prop:"K8sServiceNetworkV4=10.96.0.0" \
--prop:"K8sPodNetworkV4=10.244.0.0" \
--prop:"CWPassword=*********" \
--prop:"VMType=Hybrid" \
```

```
--prop:"IsSeed=True" \
--prop:"InitNodeCount=1" \
--prop:"InitMasterCount=1" \
$SVM_OVA_PATH \
vi://$VCENTER_LOGIN/$VCENTER_PATH/$Host
```

Example 2: dual stack template to deploy using docker installer (.tfvars file)

```
Cw_VM_Image = "cw-na-cnc-advantage-svm-7.2.0_dual"    # Line added automatically
 by installer.
ClusterIPStack        = "DUALSTACK"
ManagementVIPName     = "cwi-svm-cnc-adv-1"
ManagementVIPv4       = "x.x.x.x"
ManagementVIPv6       = "x:x:x:x::x:x"
ManagementIPv4Netmask = "255.255.255.128"
ManagementIPv6Netmask = "112"
ManagementIPv4Gateway = "x.x.x.x"
ManagementIPv6Gateway = "x:x:x:x::x:x"
DataVIPv4             = "x.x.x.x"
DataVIPv6             = "x:x:x:x:x:x:x:x"
DataIPv4Netmask       = "255.0.0.0"
DataIPv6Netmask       = "112"
DataIPv4Gateway       = "x.x.x.x"
DataIPv6Gateway       = "x:x:x:x:x:1:1:1"
DNSv4                 = "x.x.x.x"
DNSv6                 = "x:x:x:x::x"
DomainName            = "cisco.com"
K8sServiceNetwork     = "10.96.0.0"
K8sPodNetwork         = "10.244.0.0"
CWPassword            = "*******"
VMSize                = "XLarge"
NTP                   = "ntp.esl.cisco.com"
Timezone              = "Asia/Kolkata"
NonVcenter            = "false"
EnableSkipAutoInstallFeature = "True"
EnableHardReservations = true
CwVMs = {
    "0" = {
      VMName              = "cwi-svm-cnc-adv-test",
      ManagementIPv4Address = "x.x.x.x",
      ManagementIPv6Address = "2001:420:54ff:24::656:88",
      DataIPv4Address     = "x.x.x.x",
      DataIPv6Address     = "fded:1bc1:fc3e:96d0:20:1:1:88",
      NodeType            = "Hybrid"
    }
}
VCenterDC = {
  VCenterAddress = "x.x.x.x",
  VCenterUser = "abc@vsphere.local",
  VCenterPassword = "********",
  DCname = "SVM",
  MgmtNetworkName = "VM Network",
  DataNetworkName = "DPortGroup10",
  VMs = [
      {
        HostedCwVMs = ["0"],
        Host = "10.104.52.105",
        Datastore = "10.104.52.75-ssd",
        HSDatastore="10.104.52.75-ssd"
      }
  ]
```

```
}
SchemaVersion = "7.2.0"
```

Crosswork Data Gateway deployment templates and scripts

Groups sample Crosswork Data Gateway deployment scripts for geo-redundant, dual-stack, IPv4, and IPv6 deployments.

Refer to this topic for sample Crosswork Data Gateway deployment scripts.

Sample script for Crosswork Data Gateway dual-stack deployment with geo-redundancy

Provides a sample script for deploying Crosswork Data Gateway with dual-stack addressing in a geo-redundant deployment.

This topic provides the sample Crosswork Data Gateway deployment script identified in the title.

Sample script

Before running the script, ensure that you are using OVFtool version 4.4.x.

```
#!/usr/bin/env bash
VM_NAME="<VM name on vcenter>"
DM="<thin/thick>"
DS="<Datastore>"
Vcenter="<vCenter IP address>"
Host="<CDG hostname>"
DC="<Data Center>"
CDGIpv4MgmtIpv4="<CDG management IPv4>"
ManagementIPv4Netmask="<Management network IPv4 mask>"
ManagementIPv4Gateway="<Management IPv4 gateway>"
NorthDataIPv4Address="<Northbound data IPv4>"
NorthDataIPv4Netmask="<Northbound data IPv4 mask>"
NorthDataIPv4Gateway="<Northbound data IPv4 gateway>"
SouthDataIPv4Address="<Southbound data IPv4>"
SouthDataIPv4Netmask="<Southbound data IPv4 mask>"
SouthDataIPv4Gateway="<Southbound data IPv4 gateway>"
CDGIpv6MgmtIpv6="<CDG management IPv6>"
ManagementIPv6Netmask="<Management network IPv6 mask>"
ManagementIPv6Gateway="<Management network IPv6 gateway>"
NorthDataIPv6Address="<Northbound data network IPv6 address>"
NorthDataIPv6Netmask="<Northbound data network IPv6 mask>"
NorthDataIPv6Gateway="<Northbound data network IPv6 gateway>"
SouthDataIPv6Address="<Southbound data network IPv6 address>"
SouthDataIPv6Netmask="<Southbound data network IPv6 mask>"
SouthDataIPv6Gateway="<Southbound data network IPv6 gateway>"
PrivateKey="<Private key>"
Disclaimer="Cisco CDG VM for Crosswork Deployment"
DNSv6="<DNS address>"
NTP="<NTP address>"
Domain="<Domain>"
CtrlerCertChainPwd="<Controller certificate password>"
ClientCertChainPwd="<Client certificate password>"
DgAdminPwd="<Data gateway admin password>"
DgOperPwd="<Data gateway operator password>"
ControllerIP="<Unified FQDN>"
ControllerPassword="<CNC Password>"
ControllerPort="30607"
cdgDomain="<Data gateway domain>"
SouthDataNetwork="<South data network>"
NorthDataNetwork="<North data network>"
```

```
ovftool --version
ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv --overwrite
--powerOffTarget --powerOn --noSSLVerify  --allowExtraConfig \
      -ds=$DS \
--deploymentOption="onpremise-standard" \
      --diskMode=$DM \
--prop:"ControllerIP=${ControllerIP}" \
--prop:"ControllerPort=${ControllerPort}" \

--prop:"ControllerSignCertChain=cw-admin@${CwIpv6Mgmt}:/home/cw-admin/controller.pem"
 \
      --prop:"ControllerCertChainPwd=${ControllerPassword}" \
      --name="${VM_NAME}" \
      --prop:"Hostname=${cdgDomain}" \
      --prop:"Description=CDG Base VM for Automation" \
      --prop:"Vnic0IPv6Method=Static" \
      --prop:"Vnic0IPv6Address=${CDGIpv6MgmtIpv6}" \
      --prop:"Vnic0IPv6Netmask=${ManagementIPv6Netmask}" \
      --prop:"Vnic0IPv6Gateway=${ManagementIPv6Gateway}" \
      --prop:"Vnic1IPv6Method=Static" \
      --prop:"Vnic1IPv6Address=${NorthDataIPv6Address}" \
      --prop:"Vnic1IPv6Netmask=${NorthDataIPv6Netmask}" \
      --prop:"Vnic1IPv6Gateway=${NorthDataIPv6Gateway}" \
      --prop:"Vnic0IPv4Method=Static" \
      --prop:"Vnic0IPv4Address=${CDGIpv4MgmtIpv4}" \
      --prop:"Vnic0IPv4Netmask=${ManagementIPv4Netmask}" \
      --prop:"Vnic0IPv4Gateway=${ManagementIPv4Gateway}" \
      --prop:"Vnic1IPv4Method=Static" \
      --prop:"Vnic1IPv4Address=${NorthDataIPv4Address}" \
      --prop:"Vnic1IPv4Netmask=${NorthDataIPv4Netmask}" \
      --prop:"Vnic1IPv4Gateway=${NorthDataIPv4Gateway}" \
      --prop:"dg-adminPassword=${DgAdminPwd}" \
      --prop:"dg-operPassword=${DgOperPwd}" \
      --prop:"DNS=${DNSv6}" \
      --net:"vNIC0=DPortGMgmt12" \
      --net:"vNIC1=${NorthDataNetwork}" \
      --net:"vNIC2=${SouthDataNetwork}" \
      --prop:"NTP=${NTP}" \
      --prop:"Domain=${Domain}" \
      $CDG_OVA_PATH \

vi://Administrator%40vsphere%2Elocal:Crosswork123%21@$Vcenter/$DC/host/$Host
```

Sample script for Crosswork Data Gateway dual-stack deployment

Provides a sample script for deploying Crosswork Data Gateway with dual-stack addressing.

This topic provides the sample Crosswork Data Gateway deployment script identified in the title.

Sample script

Before running the script, ensure that you are using OVFtool version 4.4.x.

Note

The controller IP address parameter can be a controller VIP or FQDN of controller VIP.

```
#!/usr/bin/env bash
VM_NAME="<VM name on vcenter>"
DM="<thin/thick>"
DS="<Datastore>"
Vcenter="<vCenter IP address>"
Host="<CDG hostname>"
DC="<Data Center>"
CDGIpv4MgmtIpv4="<CDG management IPv4>"
ManagementIPv4Netmask="<Management network IPv4 mask>"
ManagementIPv4Gateway="<Management IPv4 gateway>"
NorthDataIPv4Address="<Northbound data IPv4>"
NorthDataIPv4Netmask="<Northbound data IPv4 mask>"
NorthDataIPv4Gateway="<Northbound data IPv4 gateway>"
SouthDataIPv4Address="<Southbound data IPv4>"
SouthDataIPv4Netmask="<Southbound data IPv4 mask>"
SouthDataIPv4Gateway="<Southbound data IPv4 gateway>"
CDGIpv6MgmtIpv6="<CDG management IPv6>"
ManagementIPv6Netmask="<Management network IPv6 mask>"
ManagementIPv6Gateway="<Management network IPv6 gateway>"
NorthDataIPv6Address="<Northbound data network IPv6 address>"
NorthDataIPv6Netmask="<Northbound data network IPv6 mask>"
NorthDataIPv6Gateway="<Northbound data network IPv6 gateway>"
SouthDataIPv6Address="<Southbound data network IPv6 address>"
SouthDataIPv6Netmask="<Southbound data network IPv6 mask>"
SouthDataIPv6Gateway="<Southbound data network IPv6 gateway>"
PrivateKey="<Private key>"
Disclaimer="Cisco CDG VM for Crosswork Deployment"
DNSv6="<DNS address>"
NTP="<NTP address>"
Domain="<Domain>"
CtrlerCertChainPwd="<Controller certificate password>"
ClientCertChainPwd="<Client certificate password>"
DgAdminPwd="<Data gateway admin password>"
DgOperPwd="<Data gateway operator password>"
ControllerIP="<Controller VIP or FQDN of the controller VIP>"
ControllerPassword="<CNC Password>"
ControllerPort="30607"
cdgDomain="<Data gateway domain>"
SouthDataNetwork="<South data network>"
NorthDataNetwork="<North data network>"

ovftool --version
ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv --overwrite
--powerOffTarget --powerOn --noSSLVerify  --allowExtraConfig \
      -ds=$DS \
--deploymentOption="onpremise-standard" \
      --diskMode=$DM \
--prop:"ControllerIP=${ControllerIP}" \
--prop:"ControllerPort=${ControllerPort}" \

--prop:"ControllerSignCertChain=cw-admin@${CwIpv6Mgmt}:/home/cw-admin/controller.pem"
 \
      --prop:"ControllerCertChainPwd=${ControllerPassword}" \
      --name="${VM_NAME}" \
```

```
      --prop:"Hostname=${cdgDomain}" \
      --prop:"Description=CDG Base VM for Automation" \
      --prop:"Vnic0IPv6Method=Static" \
      --prop:"Vnic0IPv6Address=${CDGIpv6MgmtIpv6}" \
      --prop:"Vnic0IPv6Netmask=${ManagementIPv6Netmask}" \
      --prop:"Vnic0IPv6Gateway=${ManagementIPv6Gateway}" \
      --prop:"Vnic1IPv6Method=Static" \
      --prop:"Vnic1IPv6Address=${NorthDataIPv6Address}" \
      --prop:"Vnic1IPv6Netmask=${NorthDataIPv6Netmask}" \
      --prop:"Vnic1IPv6Gateway=${NorthDataIPv6Gateway}" \
      --prop:"Vnic0IPv4Method=Static" \
      --prop:"Vnic0IPv4Address=${CDGIpv4MgmtIpv4}" \
      --prop:"Vnic0IPv4Netmask=${ManagementIPv4Netmask}" \
      --prop:"Vnic0IPv4Gateway=${ManagementIPv4Gateway}" \
      --prop:"Vnic1IPv4Method=Static" \
      --prop:"Vnic1IPv4Address=${NorthDataIPv4Address}" \
      --prop:"Vnic1IPv4Netmask=${NorthDataIPv4Netmask}" \
      --prop:"Vnic1IPv4Gateway=${NorthDataIPv4Gateway}" \
      --prop:"dg-adminPassword=${DgAdminPwd}" \
      --prop:"dg-operPassword=${DgOperPwd}" \
      --prop:"DNS=${DNSv6}" \
      --net:"vNIC0=DPortGMgmt12" \
      --net:"vNIC1=${NorthDataNetwork}" \
      --net:"vNIC2=${SouthDataNetwork}" \
      --prop:"NTP=${NTP}" \
      --prop:"Domain=${Domain}" \
      $CDG_OVA_PATH \

vi://Administrator%40vsphere%2Elocal:Crosswork123%21@$Vcenter/$DC/host/$Host
```

Sample script for deploying Crosswork Data Gateway with IPv6 address

Provides prerequisites, considerations, and a sample script for deploying Crosswork Data Gateway with an IPv6 address.

This topic provides the sample Crosswork Data Gateway deployment script identified in the title.

Prerequisites and considerations before running the script

Ensure these prerequisites are met:

- You are using OVFtool version 4.4.x.
- You are aware that the values for these rules may vary based on the number of vNICs deployed:
  - NicDefaultGateway
  - NicAdministration
  - NicExternalLogging
  - NicManagement
  - NicControl
  - NicNBExternalData
  - NicSBData

Sample script

```
#!/usr/bin/env bash
DM="<thin/thick>"
Disclaimer="<Disclaimer>"
DNSv4="<DNS Server>"
NTP="<NTP Server>"
Domain="<Domain>"
Hostname="<CDG hostname>"

VM_NAME="<VM name on vcenter>"
DeploymentOption="<onpremise-standard/onpremise-extended>"
DS="<Datastore>"
Host="<ESXi host>"
ManagementNetwork="<vSwitch/dvSwitch>"
DataNetwork="<vSwitch/dvSwitch>"
DeviceNetwork="<vSwitch/dvSwitch>"
ManagementIPv6Address="<CDG managment IP>"
ManagementIPv6Netmask="<CDG managment mask>"
ManagementIPv6Gateway="<CDG managment gateway>"
DataIPv6Address="<CDG Data network IP>"
DataIPv6Netmask="<CDG Data network mask>"
DataIPv6Gateway="<CDG Data network gateway>"
dgadminpwd="<CDG password for dg-admin user>"
dgoperpwd="<CDG password for dg-admin user>"
ControllerIP="<CNC Managment VIP>"
ControllerPassword="<CNC Password>"
ControllerPort="30607"

CDG_OVA_PATH=$1

VCENTER_LOGIN="Administrator%40vsphere.local@<vCenter-IP>"
VCENTER_PATH="<vCenter-DC-NAME>/host"

ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv -ds=$DS
--diskMode=$DM --overwrite --powerOffTarget --powerOn --noSSLVerify \
--allowExtraConfig \
--name=$VM_NAME \
--deploymentOption=${DeploymentOption} \
--net:"vNIC0=${ManagementNetwork}" \
--prop:"ControllerIP=${ControllerIP}" \
--prop:"ControllerPort=${ControllerPort}" \
--prop:"ControllerSignCertChain=cw-admin@${ControllerIP}:/home/cw-admin/controller.pem"
 \
--prop:"ControllerCertChainPwd=${ControllerPassword}" \
--prop:"Hostname=${Hostname}" \
--prop:"Description=${Disclaimer}" \
--prop:"DNS=${DNSv4}" \
--prop:"NTP=${NTP}" \
--prop:"Domain=${Domain}" \
--prop:"Vnic0IPv6Method=Static" \
--prop:"Vnic0IPv6Address=${ManagementIPv6Address}" \
--prop:"Vnic0IPv6Gateway=${ManagementIPv6Gateway}" \
--prop:"Vnic0IPv6Netmask=${ManagementIPv6Netmask}" \
--prop:"NicDefaultGateway=eth0" \
--prop:"NicAdministration=eth0" \
--prop:"NicExternalLogging=eth0" \
--prop:"NicManagement=eth0" \
--prop:"NicControl=eth0" \
--prop:"NicNBExternalData=eth0" \
--prop:"NicSBData=eth0" \
--prop:"dg-adminPassword=${dgadminpwd}" \
--prop:"dg-operPassword=${dgoperpwd}" \
```

```
$CDG_OVA_PATH \
vi://$VCENTER_LOGIN/$VCENTER_PATH/$Host

############################################################
Append section below for two NIC deployment
############################################################
#--net:"vNIC1=${DataNetwork}" \
#--prop:"Vnic1IPv6Method=Static" \
#--prop:"Vnic1IPv6Address=${DataIPv6Address}" \
#--prop:"Vnic1IPv6Gateway=${DataIPv6Gateway}" \
#--prop:"Vnic1IPv6Netmask=${DataIPv6Netmask}" \
#--prop:"NicDefaultGateway=eth0" \
#--prop:"NicAdministration=eth0" \
#--prop:"NicExternalLogging=eth0" \
#--prop:"NicManagement=eth0" \
#--prop:"NicControl=eth1" \
#--prop:"NicNBExternalData=eth1" \
#--prop:"NicSBData=eth1" \

############################################################
Append section below for three NIC deployment
############################################################
#--net:"vNIC1=${DataNetwork}" \
--net:"vNIC2=${DeviceNetwork}" \
#--prop:"Vnic1IPv6Method=Static" \
#--prop:"Vnic1IPv6Address=${DataIPv6Address}" \
#--prop:"Vnic1IPv6Gateway=${DataIPv6Gateway}" \
#--prop:"Vnic1IPv6Netmask=${DataIPv6Netmask}" \
#--prop:"NicDefaultGateway=eth0" \
#--prop:"NicAdministration=eth0" \
#--prop:"NicExternalLogging=eth0" \
#--prop:"NicManagement=eth0" \
#--prop:"NicControl=eth1" \
#--prop:"NicNBExternalData=eth1" \
#--prop:"NicSBData=eth2" \
```

Sample script for deploying Crosswork Data Gateway with IPv4 address

Provides prerequisites, considerations, and a sample script for deploying Crosswork Data Gateway with an IPv4 address.

This topic provides the sample Crosswork Data Gateway deployment script identified in the title.

Prerequisites and considerations before running the script

Ensure these prerequisites are met:

- You are using OVFtool version 4.4.x.
- You are aware that the values for these rules may vary based on the number of vNICs deployed:
  - NicDefaultGateway
  - NicAdministration
  - NicExternalLogging
  - NicManagement
  - NicControl
  - NicNBExternalData
  - NicSBData

Sample script

```
#!/usr/bin/env bash
DM="<thin/thick>"
Disclaimer="<Disclaimer>"
DNSv4="<DNS Server>"
NTP="<NTP Server>"
Domain="<Domain>"
Hostname="<CDG hostname>"

VM_NAME="<VM name on vcenter>"
DeploymentOption="<onpremise-standard/onpremise-extended>"
DS="<Datastore>"
Host="<ESXi host>"
ManagementNetwork="<vSwitch/dvSwitch>"
DataNetwork="<vSwitch/dvSwitch>"
DeviceNetwork="<vSwitch/dvSwitch>"
ManagementIPv4Address="<CDG managment IP>"
ManagementIPv4Netmask="<CDG managment mask>"
ManagementIPv4Gateway="<CDG managment gateway>"
DataIPv4Address="<CDG Data network IP>"
DataIPv4Netmask="<CDG Data network mask>"
DataIPv4Gateway="<CDG Data network gateway>"
dgadminpwd="<CDG password for dg-admin user>"
dgoperpwd="<CDG password for dg-admin user>"
ControllerIP="<CNC Managment VIP>"
ControllerPassword="<CNC Password>"
ControllerPort="30607"


CDG_OVA_PATH=$1

VCENTER_LOGIN="Administrator%40vsphere.local@<vCenter-IP>"
VCENTER_PATH="<vCenter-DC-NAME>/host"

ovftool --acceptAllEulas --skipManifestCheck --X:injectOvfEnv -ds=$DS
--diskMode=$DM --overwrite --powerOffTarget --powerOn --noSSLVerify \
--allowExtraConfig \
--name=$VM_NAME \
--deploymentOption=${DeploymentOption} \
--net:"vNIC0=${ManagementNetwork}" \
--prop:"ControllerIP=${ControllerIP}" \
--prop:"ControllerPort=${ControllerPort}" \
--prop:"ControllerSignCertChain=cw-admin@${ControllerIP}:/home/cw-admin/controller.pem"
 \
--prop:"ControllerCertChainPwd=${ControllerPassword}" \
--prop:"Hostname=${Hostname}" \
--prop:"Description=${Disclaimer}" \
--prop:"DNS=${DNSv4}" \
--prop:"NTP=${NTP}" \
--prop:"Domain=${Domain}" \
--prop:"Vnic0IPv4Method=Static" \
--prop:"Vnic0IPv4Address=${ManagementIPv4Address}" \
--prop:"Vnic0IPv4Gateway=${ManagementIPv4Gateway}" \
--prop:"Vnic0IPv4Netmask=${ManagementIPv4Netmask}" \
--prop:"NicDefaultGateway=eth0" \
--prop:"NicAdministration=eth0" \
--prop:"NicExternalLogging=eth0" \
--prop:"NicManagement=eth0" \
--prop:"NicControl=eth0" \
--prop:"NicNBExternalData=eth0" \
--prop:"NicSBData=eth0" \
--prop:"dg-adminPassword=${dgadminpwd}" \
```

```
--prop:"dg-operPassword=${dgoperpwd}" \
$CDG_OVA_PATH \
vi://$VCENTER_LOGIN/$VCENTER_PATH/$Host

#############################################################
Append section below for two NIC deployment
#############################################################
#--net:"vNIC1=${DataNetwork}" \
#--prop:"Vnic1IPv4Method=Static" \
#--prop:"Vnic1IPv4Address=${DataIPv4Address}" \
#--prop:"Vnic1IPv4Gateway=${DataIPv4Gateway}" \
#--prop:"Vnic1IPv4Netmask=${DataIPv4Netmask}" \
#--prop:"NicDefaultGateway=eth0" \
#--prop:"NicAdministration=eth0" \
#--prop:"NicExternalLogging=eth0" \
#--prop:"NicManagement=eth0" \
#--prop:"NicControl=eth1" \
#--prop:"NicNBExternalData=eth1" \
#--prop:"NicSBData=eth1" \

#############################################################
Append section below for three NIC deployment
#############################################################
#--net:"vNIC1=${DataNetwork}" \
--net:"vNIC2=${DeviceNetwork}" \
#--prop:"Vnic1IPv4Method=Static" \
#--prop:"Vnic1IPv4Address=${DataIPv4Address}" \
#--prop:"Vnic1IPv4Gateway=${DataIPv4Gateway}" \
#--prop:"Vnic1IPv4Netmask=${DataIPv4Netmask}" \
#--prop:"NicDefaultGateway=eth0" \
#--prop:"NicAdministration=eth0" \
#--prop:"NicExternalLogging=eth0" \
#--prop:"NicManagement=eth0" \
#--prop:"NicControl=eth1" \
#--prop:"NicNBExternalData=eth1" \
#--prop:"NicSBData=eth2" \
```

Sample auto action templates

Provides sample auto action definition and manifest templates in JSON and YAML formats.

This topic includes examples of auto action templates.

- Sample auto-action definition (JSON) file
- Sample auto-action manifest (YAML) file

> **Note**
>
> Make sure to replace the placeholder values with actual values relevant to your environment.

Sample auto-action definition (JSON) file

```
{
    "auto_action": {
```

```
        "add_to_repository_requests": [
            {
                "file_location": {
                    "uri_location": {
                        "uri":
"https://example.com/path/to/cw-na-cncessential-7.2.0-260130.tar.gz"
                    }
                }
            },
            {
                "file_location": {
                    "uri_location": {
                        "uri":
"https://example.com/path/to/cw-na-cncadvantage-7.2.0-260130.tar.gz"
                    }
                }
            },
            {
                "file_location": {
                    "uri_location": {
                        "uri":
"https://example.com/path/to/cw-na-cncaddon-7.2.0-260130.tar.gz"
                    }
                }
            }
        ],
        "install_activate_requests": [
            {
                "package_identifier": {
                    "_comment": "Part of essentials capp",
                    "version": "7.2.0",
                    "id": "capp-common-ems-services"
                }
            },
            {
                "package_identifier": {
                    "_comment": "Part of advantage capp",
                    "version": "7.2.0",
                    "id": "capp-cat"
                }
            },
            {
                "package_identifier": {
                    "_comment": "Part of advantage capp",
                    "version": "7.2.0",
                    "id": "capp-coe"
                }
            },
            {
                "package_identifier": {
                    "_comment": "Part of advantage capp",
                    "version": "7.2.0",
                    "id": "capp-aa"
                }
            },
            {
                "package_identifier": {
                    "_comment": "Part of add on capp",
                    "version": "7.2.0",
                    "id": "capp-ca"
                }
            },
```

```
            {
                "package_identifier": {
                    "_comment": "Part of add on capp",
                    "version": "7.2.0"",
                    "id": "capp-hi"
                }
            }
        ]
    }
}
```

Sample auto-action manifest (YAML) file

```
# auto_action.yaml
# This YAML file configures automated actions to be performed by the
orchestrator.
# It includes actions to add files to a repository and install/activate
packages.
# Modify the sample file based on the need. Add or remove
add_to_repository_requests accordingly.

# Usage Instructions:
# 1. Ensure the file paths and URIs are accessible from your orchestrator
environment.
# 2. Supported types for adding files to the repository include:
#    - Local file URIs (e.g.,
file:///mnt/app_capp/cw-na-collectors-7.xx.xx.tar.gz)
#    - Remote file URIs over HTTP/HTTPS (e.g.,
https://example.com/sample.tar.gz)
#    - Remote files accessed via SCP (e.g., /path/to/remote/file.tar.gz)
# 3. Pick the add_to_repository_requests type based on the need.
# 4. Either `add_to_repository_requests` or `install_activate_requests` should
 be present, or both can be included.
# 5. The orchestrator will add the specified files to the repository and then
 install/activate the listed packages if both are present.
# 6. Double-check the package identifiers and versions for accuracy.
# 7. Ensure `install_activate_requests` is defined in the dependency order.
If CAPP2 is dependent on CAPP1, CAPP1 should be defined first.
# 8. Package ID references:
#    - E-CDG: capp-cdg
#    - EMS: capp-common-ems-services
#    - COE: capp-coe
#    - CAT: capp-cat
#    - Service Health: capp-aa
#    - E-NSO: capp-enso
#    - CP-INFRA: capp-cpinfra
#    - CP-Design: capp-design
#    - CP-Collector: capp-collector

auto_action:
  # List of requests to add files to the repository
  add_to_repository_requests:
    - file_location:
        uri_location:
          # Remote file URI to be added to the repository.
          # Example: https://example.com/sample.tar.gz
          uri: <remote_file_uri>
          # Optional: Basic authentication credentials for HTTPS
          # basic_auth:
          #   username: <username>
```

```
          #   password: <password>
          # Optional: Skip TLS verification if https
          # skip_tls_verify: <true/false>
    - file_location:
        scp_location:
          # Remote file path to be added to the repository.
          # Example: /path/to/remote/file.tar.gz
          remote_file: <remote_file_path>
          # SSH configuration for remote file access
          ssh_config:
            remote_host: <ssh_host>
            port: <ssh_port>
            username: <ssh_username>
            password: <ssh_password>
    - file_location:
        uri_location:
          # Local file URI to be added to the repository.
          # Example: file:///mnt/app_capp/cw-na-collectors-7.xx.xx.tar.gz
          uri: <local_file_uri_1>
  # List of requests to install and activate packages. Ensure the dependency
order of CAPP activation is retained
  install_activate_requests:
    - package_identifier:
        # Identifier of the package to be installed/activated
        # Example: capp-cdg
        id: <package_id_1>
        version: <optional_package_version_1>
    - package_identifier:
        # Identifier of the package to be installed/activated
        # Example: capp-common-ems-services
        id: <package_id_2>
        version: <optional_package_version_2>

# Note:
# - If using SSH for remote file access, replace <ssh_host>, <ssh_port>,
<ssh_username>, and <ssh_password> with the appropriate credentials and
configuration.
# - If using HTTPS with basic authentication, replace <username> and <password>
 with the appropriate credentials.
```

Geo redundancy templates

Groups sample cross-cluster inventory and multi-cluster upgrade inventory templates for geo redundancy.

Refer to this topic for sample templates used with geo redundancy.

Sample cross cluster inventory templates

Provides sample cross-cluster inventory files for geo redundancy across Day 0 and Day N deployment scenarios.

Here are some examples of the cross cluster inventory file (.yaml) that you need to prepare to enable geo redundancy:

- Day 0: geo inventory for active cluster deployed with peer cluster (standby)
- Day 0: geo inventory for active cluster deployed with peer clusters (standby and arbiter)
- Day N: geo inventory for active cluster deployed without peer cluster (standby)
- Day N: geo inventory for active cluster deployed without arbiter VM

Note

- Cross Cluster inventory supports only IPv6 addresses. If both IPv4 and IPv6 addresses are present, only the IPv6 addresses will be used. If the Cross Cluster inventory is provided with only IPv4 addresses, the system will generate an error.
- In a geo redundant setup with dual stack deployment, the inventory file must reference the IPv6 VIP for both clusters instead of the IPv4 VIP, as the unified cross cluster uses IPv6. It is recommended to use IPv6 addresses for the unified cross cluster endpoint since all geo-redundancy traffic is routed via IPv6.
- For details on each parameter, see the **sample file** downloaded from the **Geo Redundancy Manager** window.

Day 0: geo inventory for active cluster deployed with peer cluster (standby)

```
##############################################################################################
#
# Use case:
#   - Geo-inventory yaml to be used when geo-activating the active cluster
after standby
#     cluster is fully deployed.
#
# Flag settings:
#   - Set 'is_skip_peer_check_enabled' to false since the peer cluster is
already deployed.
#   - Set 'is_post_migration_activation' to false since this is not a migration
 flow.
#
# Important Notes:
#   - The 'is_skip_peer_check_enabled: true' is required when:
#     - The standby cluster is not deployed before geo-activating the active
cluster.
#     - Enabling geo mode on a system post disaster restore after multi cluster
 failure.
#     - Additionally, if 'is_skip_peer_check_enabled: true', following details
 are required:
#       - Cluster connectivity
#       - Cluster node's connectivity details
#   - See geo-inventory for the case when standby cluster is not deployed
before geo-activating.
#
##############################################################################################
---
meta_version: 1.0.0
crosscluster_name: mycnc-geo-cluster
crosscluster_unified_connectivity:
  unified_end_point:
    unified_endpoint_type:
      fqdn_type: {}
    unified_endpoint_implementation: DNS
  data_fqdn:
    domain_name: cw.cisco
    host_name: geodata
  management_fqdn:
    domain_name: cw.cisco
    host_name: geomanagement
clusters:
- cluster_name: cluster-sjc
```

```
  cluster_id: fded:1bc1:fc3e:96d0:192:168:5:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: ACTIVE
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:10:500
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:5:500
    management_vip_mask: 112
  site_location:
    location: San Jose
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
- cluster_name: cluster-nyc
  cluster_id: fded:1bc1:fc3e:96d0:192:168:6:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: STANDBY
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:11:500
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:6:500
    management_vip_mask: 112
  site_location:
    location: New York
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
secret: Your-secret1
is_post_migration_activation: false
is_skip_peer_check_enabled: false
```

Day 0: geo inventory for active cluster deployed with peer clusters (standby and arbiter)

```
#############################################################################################
#
# Use case:
#   - Geo-inventory yaml to be used when geo-activating the active cluster
after standby
#     and arbiter clusters are fully deployed.
#
# Flag settings:
#   - Set 'is_skip_peer_check_enabled' to false since both peer clusters are
already deployed.
#   - Set 'is_post_migration_activation' to false since this is not a migration
 flow.
```

```
#
# Important Notes:
#   - The 'is_skip_peer_check_enabled: true' is required when:
#     - The standby cluster is not deployed before geo-activating the active
cluster.
#     - Enabling geo mode on a system post disaster restore after multi cluster
 failure.
#     - Additionally, if 'is_skip_peer_check_enabled: true', following details
 are required:
#       - Cluster connectivity
#       - Cluster node's connectivity details
#   - See geo-inventory for the case when standby cluster is not deployed
before geo-activating.
#
##############################################################################################
---
meta_version: 1.0.0
crosscluster_name: mycnc-geo-cluster
crosscluster_unified_connectivity:
  unified_end_point:
    unified_endpoint_type:
      fqdn_type: {}
    unified_endpoint_implementation: DNS
  data_fqdn:
    domain_name: cw.cisco
    host_name: geodata
  management_fqdn:
    domain_name: cw.cisco
    host_name: geomanagement
clusters:
- cluster_name: cluster-sjc
  cluster_id: fded:1bc1:fc3e:96d0:192:168:5:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: ACTIVE
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:10:500
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:5:500
    management_vip_mask: 112
  site_location:
    location: San Jose
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
- cluster_name: cluster-nyc
  cluster_id: fded:1bc1:fc3e:96d0:192:168:6:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: STANDBY
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:11:500
```

```
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:6:500
    management_vip_mask: 112
  site_location:
    location: New York
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
- cluster_name: cluster-aus
  cluster_id: fded:1bc1:fc3e:96d0:192:168:5:506
  cluster_type: ARBITER
  initial_preferred_leadership_state: STANDBY
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:10:506
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:5:506
    management_vip_mask: 112
  site_location:
    location: Austin
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
secret: Your-secret1
is_post_migration_activation: false
is_skip_peer_check_enabled: false
```

Day N: geo inventory for active cluster deployed without peer cluster (standby)

```
#############################################################################################
#
# Use case:
#   - Geo-inventory yaml to be used when geo-activating the active cluster
before standby
#     cluster is fully deployed.
#
# Flag settings:
#   - Set 'is_skip_peer_check_enabled' to true since the peer cluster is not
deployed.
#   - Set 'is_post_migration_activation' to true.
#
# Important Notes:
#   - The 'is_skip_peer_check_enabled: true' is required when:
#     - The standby cluster is not deployed before geo-activating the active
cluster.
#     - Enabling geo mode on a system post disaster restore after multi cluster
 failure.
#     - Additionally, if 'is_skip_peer_check_enabled: true', following details
 are required:
#       - Cluster connectivity
```

```
#        - See explicit IPv4/IPv6 VIP and netmasks under cluster connectivity.
#       - Cluster node's connectivity details
#         - See 'cluster_node_details' object under the cluster, which is not
 deployed.
#
##############################################################################################
---
meta_version: 1.0.0
crosscluster_name: mycnc-geo-cluster
crosscluster_unified_connectivity:
  unified_end_point:
    unified_endpoint_type:
      fqdn_type: {}
    unified_endpoint_implementation: DNS
  data_fqdn:
    domain_name: cw.cisco
    host_name: geodata
  management_fqdn:
    domain_name: cw.cisco
    host_name: geomanagement
clusters:
- cluster_name: cluster-sjc
  cluster_id: fded:1bc1:fc3e:96d0:192:168:5:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: ACTIVE
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:10:500
    data_vip_mask: 112
    management_vip: fded:1bc1:fc3e:96d0:192:168:5:500
    management_vip_mask: 112
    data_vip_ipv4: 10.10.10.50
    data_vip_ipv4_mask: 24
    management_vip_ipv4: 192.168.5.50
    management_vip_ipv4_mask: 24
    data_vip_ipv6: fded:1bc1:fc3e:96d0:10:10:10:500
    data_vip_ipv6_mask: 112
    management_vip_ipv6: fded:1bc1:fc3e:96d0:192:168:5:500
    management_vip_ipv6_mask: 112
  site_location:
    location: San Jose
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
- cluster_name: cluster-nyc
  cluster_id: fded:1bc1:fc3e:96d0:192:168:6:500
  cluster_type: CROSSWORK_CLUSTER
  initial_preferred_leadership_state: STANDBY
  connectivity:
    unified_end_point:
      unified_endpoint_type:
        ip_type: {}
      unified_endpoint_implementation: VRRP
    data_vip: fded:1bc1:fc3e:96d0:10:10:11:500
    data_vip_mask: 112
```

```
    management_vip: fded:1bc1:fc3e:96d0:192:168:6:500
    management_vip_mask: 112
    data_vip_ipv4: 10.10.11.50
    data_vip_ipv4_mask: 24
    management_vip_ipv4: 192.168.6.50
    management_vip_ipv4_mask: 24
    data_vip_ipv6: fded:1bc1:fc3e:96d0:10:10:11:500
    data_vip_ipv6_mask: 112
    management_vip_ipv6: fded:1bc1:fc3e:96d0:192:168:6:500
    management_vip_ipv6_mask: 112
  site_location:
    location: New York
  cluster_credential:
    https_credential:
      username: admin
      password: ********
    ssh_credential:
      username: cw-admin
      password: ********
  cluster_node_details:
    - unified_end_point:
        unified_endpoint_type:
          ip_type: {}
        unified_endpoint_implementation: VRRP
      data_ipv4: 10.10.11.51
      data_ipv4_mask: 24
      management_ipv4: 192.168.6.51
      management_ipv4_mask: 24
      data_ipv6: fded:1bc1:fc3e:96d0:10:10:11:501
      data_ipv6_mask: 112
      management_ipv6: fded:1bc1:fc3e:96d0:192:168:6:501
      management_ipv6_mask: 112
    - unified_end_point:
        unified_endpoint_type:
          ip_type: {}
        unified_endpoint_implementation: VRRP
      data_ipv4: 10.10.11.52
      data_ipv4_mask: 24
      management_ipv4: 192.168.6.52
      management_ipv4_mask: 24
      data_ipv6: fded:1bc1:fc3e:96d0:10:10:11:502
      data_ipv6_mask: 112
      management_ipv6: fded:1bc1:fc3e:96d0:192:168:6:502
      management_ipv6_mask: 112
    - unified_end_point:
        unified_endpoint_type:
          ip_type: {}
        unified_endpoint_implementation: VRRP
      data_ipv4: 10.10.11.53
      data_ipv4_mask: 24
      management_ipv4: 192.168.6.53
      management_ipv4_mask: 24
      data_ipv6: fded:1bc1:fc3e:96d0:10:10:11:503
      data_ipv6_mask: 112
      management_ipv6: fded:1bc1:fc3e:96d0:192:168:6:503
      management_ipv6_mask: 112
secret: Your-secret1
is_post_migration_activation: true
is_skip_peer_check_enabled: true
```

Day N: geo inventory for active cluster deployed without arbiter VM

- This is the day-N arbiter deployment of geo HA setup, i.e., the arbiter cluster is deployed after deploying the active and standby clusters.
- The inventory is same as the one used in Day 0: geo inventory for active cluster deployed with peer clusters (standby and arbiter), but you must set the **is_post_migration_activation** flag to `true`..
- You must enable geo redundancy on the active and standby clusters with "re-import" flow, followed by normal geo activation on the arbiter VM.

Sample multi-cluster upgrade inventory template

Provides a sample inventory template for upgrading multiple clusters.

This topic provides a sample inventory template for a multi-cluster upgrade.

Sample script

Here is an example of the inventory file (.yaml) that you need to prepare for multi-cluster upgrade:

```
#############################################################################################
#
# Use case:
#   - Geo inventory yaml to be used when performing geo upgrade from site-2.
#
#############################################################################################


## meta version of yaml ###
meta_version: 1.0.0
## Crosscluster name , mutable
crosscluster_name: mycnc-geo-cluster
### Unified endpoint of multi cluster for management and data endpoint across
 clusters
crosscluster_unified_connectivity:
  unified_end_point:
    unified_endpoint_type:
      #### fqdn_type,ip_type are options , only fqdn_type is supported.
      fqdn_type: {}
    #### DNS,BGP,NLB are options, only DNS is supported for now.
    unified_endpoint_implementation: DNS
     ### The below is needed if fqdn_type is chosen,else data_vip,mgmt_vip
could be used for ip_type endpoint type
  management_fqdn:
    ## cnc domain zone, DNS server would be checked for resolution
    domain_name: your-name.domain
    host_name: your-unified-cnc-mgmt-hostname
  data_fqdn:
    ## cnc domain zone name, DNS server would be checked for resolution
    domain_name: your-name.domain
    host_name: your-unified-cnc-data-hostname

### Constituent clusters ####
clusters:
##### Mutable cluster name
  - cluster_name: cluster-sjc
    ## Fill cluster_id with cluster Management VIP
    cluster_id: fded:1bca:fc3f:96d0:192:168:5:101
    connectivity:
      ### Intra cluster (within a cluster) unified endpoint  ###
      ### Endpoint type is ip_type,fqdn_type , Implementation could be
```

```
VRRP,NLB,BGP.
      unified_end_point:
        unified_endpoint_type:
          ip_type: {}
        #### VRRP,BGP,NLB are options, only VRRP,ip_type is supported for now
 in on prem. For cloud NLB,
        ### fqdn_type could be used.
        unified_endpoint_implementation: VRRP
      ### The below is needed if ip_type is chosen,else data_fqdn,mgmt_fqdn
could be used for fqdn_type endpoint type
      ## Your intra cluster data vip
      data_vip: 10.10.10.11
      ## data vip subnet mask
      data_vip_mask: 0
      ## Your intra cluster management vip
      management_vip: 20.20.20.11
      ## management vip subnet mask
      management_vip_mask: 0
      ## management and data fqdn for crosscluster instance
      ## management and data fqdn is applicable for only for unified
crosscluster instance
    ## STANDBY or ACTIVE for leadership state
    initial_preferred_leadership_state: ACTIVE
   ### DC location , needs to be unique per cluster, For cloud region-az could
 be used
    site_location:
      location: San Jose
      #Mutable credentials
    cluster_credential:
      ## This is the https credential post first time cluster login
      https_credential:
        username: admin
        #### pwd/secrets are within single quotes,if special chars are used
        password: your-password
      ssh_credential:
        username: admin
        #### pwd/secrets are within single quotes,if special chars are used
        password: your-password
##### Mutable cluster name
  - cluster_name: cluster-nyc
    ## Fill cluster_id with cluster Management VIP
    cluster_id: fded:1bca:fc3f:96d0:192:168:6:101
    ## STANDBY or ACTIVE for leadership state
    initial_preferred_leadership_state: STANDBY
    connectivity:
      ### Intra cluster (within a cluster) unified endpoint  ###
      ### Endpoint type is ip_type,fqdn_type , Implementation could be
VRRP,NLB,BGP.
      unified_end_point:
        unified_endpoint_type:
          ip_type: {}
        #### VRRP,BGP,NLB are options, only VRRP,ip_type is supported for now
 in on prem. For cloud NLB,
        ### fqdn_type could be used.
        unified_endpoint_implementation: VRRP
      ### The below is needed if ip_type is chosen,else data_fqdn,mgmt_fqdn
could be used for fqdn_type endpoint type
      ## Your intra cluster data vip
      data_vip: 30.30.30.11
      ## data vip subnet mask
      data_vip_mask: 0
      ## Your intra cluster management vip
```

```
      management_vip: 40.40.40.11
      ## management vip subnet mask
      management_vip_mask: 0
      ## management and data fqdn for crosscluster instance
      ## management and data fqdn is applicable for only for unified
crosscluster instance
    ### DC location , needs to be unique per cluster, For cloud region-az could
 be used
    site_location:
      location: New York City
    #Mutable credentials
    cluster_credential:
      ## This is the https credential post first time cluster login
      https_credential:
        username: admin
        #### pwd/secrets are within single quotes if special chars are used
        password: your-password
      ssh_credential:
        username: admin
        #### pwd/secrets are within single quotes if special chars are used
        password: your-password
#### Mutable secret are within single quotes if special chars are used, used
to kick-start inter cluster mTLS
### needs to be >= 10 chars with at-least 1 special,upper,numerical characters
secret: Your-secret1
### Set this to true , if one is enabling geo mode on a system post migration
 setup, rather than a fresh first time
## install
is_post_migration_activation: false
### Set this to true, if one wants to skip peer check, if one is enabling geo
 mode on a system post disaster restore
## after multi cluster failure
is_skip_peer_check_enabled: false
# Disable nso backup from peer ,set true if nso not deployed ,else false.
disable_nso_backup_from_peer: true
# Backup job timeout  ,
# Valid values for eg. "3h" ,"4h", "5h30m40s" , "100m"
# By default in system is "3h",
# If set timeout less than "3h",it will set to default "3h" internally.
# If cluster deployed on test bed with all apps with not much data ,then default
 timeout is ok,
# If its scale setup/or setup with huge data , then calculate and set timeout
 like :
# for eg: if postgres/timescale data size is ~300 GB then it take approx 8
hours  for backup ,
# so calculate in way  : for 100GB approx ~3h ,
# for 300 GB approx 9 hours ,
# for 1 TB approx 30 hours ,and so on
backup_job_poll_timout: "3h"
# Migration job timeout,
# Valid values for eg. "3h" ,"4h", "5h30m40s" , "100m"
# by default in system is "3h",
# If set timeout less than "3h",it will set to default "3h" internally.
# If cluster deployed on test bed with all apps with not much data ,then default
 timeout is ok,
# If its scale setup/or setup with huge data , then calculate and set timeout
 like :
# for eg: if postgres/timescale data size is ~300 GB then it take approx 8
hours  for backup ,
# so calculate in way  : for 100GB approx ~3h ,
# for 300 GB approx 9 hours ,
# For 1 TB approx 30 hours ,and so on
```

```
migration_job_poll_timout: "3h"
# Storage setting for migration
storage_settings:
  ##scp host setting type
  scp_host:
    ## remote location
    remote_location: "/mnt/robot_datafs/public/l1"
    ## ssh configuration
    ssh_config:
      ## ssh configuration host ip
      remote_host: 172.20.80.84
      ## ssh configuration username
      username: cw-admin
      ## ssh configuration password
      password: your-password
       ## ssh  port
      port: 22
```