



Cisco SVO Setup and Installation

Table 1: Feature History

| Feature Name | Release Information | Feature Description |
|------------------------|-----------------------------|---|
| Cisco SVO Installation | Cisco NCS 2000 Release 12.3 | <p>The following installation methods are introduced to ease SVO installation:</p> <ul style="list-style-type: none">• The <code>svoTools-12.3.sh</code> script is introduced to simplify system deployment in the server. You can use the script to install, uninstall, extract, and load the SVO images.• IPv4 Port Forwarding allows to save one IPv4 address of the management network for each running Admin Plane. When Port Forwarding is enabled, Admin Plane can share the same IPv4 address assigned to the host NIC.• Standalone configuration enables you to install the SVO instances in standalone mode for lab and development usage. |

| Feature Name | Release Information | Feature Description |
|---|-------------------------------|---|
| Layer 3 Management Network Connectivity through BGP | Cisco NCS 2000 Release 12.3.1 | Management interconnection for servers or VMs in different locations is now supported at Layer 3 through the Border Gateway Protocol (BGP). This simplifies the management of the Admin Plane servers by using core routers, which use the BGP applications in the VMs to route to the correct SVO instances. This approach allows you to configure the same management subnet for the SVO instances and different management subnets for distributed servers or VMs. Unlike in L2, which uses multiple protocols, the L3 management network needs only the BGP protocol to improve the performance of the network. |

- [Overview, on page 2](#)
- [Recommended Hardware, on page 6](#)
- [Recommended Software, on page 6](#)
- [Recommended Resource for Virtual Machines, on page 7](#)
- [Required Network Resources, on page 8](#)
- [Bandwidth and Latency Requirements, on page 10](#)
- [Install Docker Engine, on page 10](#)
- [Network Configuration File, on page 13](#)
- [Prepare the Network Configuration, on page 21](#)
- [Standalone SVO Configuration, on page 22](#)
- [Installation of SVO, on page 22](#)
- [Bring Up Admin Plane with the SVO Installation Tool, on page 26](#)
- [Deployment, on page 27](#)
- [Disaster Recovery, on page 30](#)
- [Use Cases, on page 32](#)

Overview

Cisco NCS 2000 Shelf Virtualization Orchestrator (SVO) supports different modes of deployment in the optical network. This chapter provides information related to the installation and configuration of the SVO software application on external servers or virtual machines (VMs).

SVO is a Docker-based application that leverages a complex networking configuration. The SVO application supports both IPv4 and IPv6 scenarios through different Docker networks created during the installation. The extreme flexibility of the SVO software solution and its network architecture allows satisfying any requirement with proper configuration.

IPv4 Port Forwarding

IPv4 port forwarding allows saving one IPv4 address of the management network for each running admin plane (two IPv4 addresses in a high availability setup). The admin plane application shares the same IPv4 address assigned to the host NIC.



Note IPv4 port forwarding can be used in the server only.

Use the `svoTools-12.3.1.sh` script described in [Install SVO in the External Server Using the Installation Script, on page 23](#) to enable IPv4 Port Forwarding.



Note When preparing the network configuration YAML file, the same IPv4 address must be assigned to the `adminplane` and `host-nic` fields in the management IPv4 section.

For information on how to create the network YAML configuration file, see [Network Configuration File, on page 13](#).

Layer 3 Management Interconnection for Geo Redundancy

Layer 3 management interconnection between servers for Geo redundancy allows you to avoid stretching the Layer 2 subnet between the two locations where the servers are installed.

Layer 3 management interconnection solution is based on the Border Gateway Protocol (BGP). Servers or VMs in different locations establish a BGP neighbor relationship with a core router. The SVO Admin Plane in the servers or VMs advertises the routes that are related to itself and each SVO instance running in Active mode to the core router.

Layer 3 management interconnection is implemented with goBGP 3.0.0. goBGP is an Open Source BGP implementation that can be downloaded from <https://github.com/osrg/gobgp>. We recommend goBGP 3.0.0 for better compatibility.

goBGP is available as an archive file. It includes a daemon and a client. After you download and extract the files in the local device, it is necessary to create a configuration file. The following examples show the configuration files for IPv4 setup and IPv4/IPv6 setup.

`gobgp_ipv4.yml`

```
global:
  config:
    as: 65001
    router-id: "192.168.85.2"
neighbors:
- config:
  peer-as: 65001
  neighbor-address: "192.168.85.1"
  auth-password: "cisco"
```

`gobgp_ipv4_ipv6.yml`

```
global:
  config:
```

```

    as: 65001
    router-id: "192.168.85.2"
  neighbors:
  - config:
    peer-as: 65001
    neighbor-address: "fd00::192:168:85:1"
    auth-password: "cisco"
  - config:
    peer-as: 65001
    neighbor-address: "192.168.85.1"
    auth-password: "cisco"

```



Note The field **auth-password** is optional, remove it if authentication is not a requirement.

When the configuration file is ready, you can execute the goBGP daemon with administrative privileges via command line (or configuring it as a service).

Use the following command for executing the **gobgp_ipv4.yml** file:

```
[gacrux@arturo-vm3 gobgp]$ sudo ./gobgpd -t yaml -f ./gobgp_ipv4.yml
```

Use the following command for executing the **gobgp_ipv4_ipv6.yml** file:

```
[gacrux@arturo-vm3 gobgp]$ sudo ./gobgpd -t yaml -f ./gobgp_ipv4_ipv6.yml
```

To connect to the SVO Installation Tool remotely, its route must be manually advertised with the goBGP client. This operation must be done on the local and remote servers using the following commands. The example commands contain the details of the local machine that must be substituted for the following custom configuration and the next-hop addresses:

- IP address of Admin Plane / SVO Installation Tool
- IP address of the server management network

Example

```
[gacrux@VM1]$ sudo ./gobgp global rib -a ipv4 add 10.58.253.2/32 nexthop 192.168.85.2
[gacrux@VM2]$ sudo ./gobgp global rib -a ipv4 add 10.58.253.3/32 nexthop 172.16.16.2
```

For information on how to create the network YAML configuration file, see [Network Configuration File](#), on page 13. Refer to [Deployment of Servers in Different Locations \(L3 Interconnection\)](#), on page 29 for details about SVO software application running with Layer 3 management interconnection and [Use Case 3 - Dislocated Servers \(L3 Interconnection\)](#), on page 34 for use case example.

Using **svoTools.sh** Script

Use the **svoTools-12.3.1.sh** script described in [Install SVO in the External Server Using the Installation Script](#) to enable Layer 3 Management Interconnection.

Running goBGP Daemon as a Service

After you download, extract the files, and create the configuration file as described in the previous section, you can configure the goBGP Daemon as a service.

Create the service file with the following content (substitute the highlighted parts with the correct path where you extracted gobgpd and with the path of the configuration file):

```
[gacrux@arturo-vm3 ~]$ cat /usr/lib/systemd/system/gobgpd.service
[Unit]
```

```

Description=goBGP 3.0 server daemon
Documentation=www.gobgp.com
After=network.target

[Service]
Type=exec
ExecStart=/home/gacruX/gobgp-3.0/gobgpd -t yaml -f /home/gacruX/gobgp-3.0/gobgp_ipv4.yml
ExecReload=/bin/kill -HUP $MAINPID
KillMode=process
Restart=on-failure
RestartSec=2s

[Install]
WantedBy=multi-user.target

```

Enable the Service

To enable the service, use the following command:

```

[gacruX@arturo-vm3 ~]$ sudo systemctl enable gobgpd.service
Created symlink from /etc/systemd/system/multi-user.target.wants/gobgpd.service to
/usr/lib/systemd/system/gobgpd.service.
[gacruX@arturo-vm3 ~]$

```

Start the Service

To start the service, use the following command:

```

[gacruX@arturo-vm3 ~]$ sudo systemctl start gobgpd.service

```

Check the Status of the Service

To check the status of the service, use the following command:

```

[gacruX@arturo-vm3 ~]$ sudo systemctl status gobgpd.service
● gobgpd.service - goBGP 3.0 server daemon
   Loaded: loaded (/usr/lib/systemd/system/gobgpd.service; enabled; vendor preset: disabled)
   Active: active (running) since Fri 2022-03-18 16:32:01 CET; 8s ago
     Main PID: 28139 (gobgpd)
        Tasks: 10
       Memory: 8.5M
      CGroup: /system.slice/gobgpd.service
              └─28139 /home/gacruX/LG/gobgp-3.0/gobgpd -t yaml -f
/home/gacruX/LG/gobgp-3.0/gobgp_ipv4.yml

Mar 18 16:32:01 arturo-vm3 systemd[1]: Started goBGP 3.0 server daemon.
Mar 18 16:32:01 arturo-vm3 gobgpd[28139]: {"level":"info","msg":"gobgpd
started","time":"2022-03-18T16:32:01+01:00"}
Mar 18 16:32:01 arturo-vm3 gobgpd[28139]: {"Topic":"Config","level":"info","msg":"Finished
reading the config file","time":"2022-03-18T16:32:01+01:00"}
Mar 18 16:32:01 arturo-vm3 gobgpd[28139]: {"Key":"192.168.85.1","Topic":"config","level":"info","msg":"Add
Peer","time":"2022-03-18T16:32:01+01:00"}
Mar 18 16:32:01 arturo-vm3 gobgpd[28139]: {"Key":"192.168.85.1","Topic":"Peer","level":"info","msg":"Add a peer
configuration","time":"2022-03-18T16:32:01+01:00"}
Mar 18 16:32:06 arturo-vm3 gobgpd[28139]: {"Key":"192.168.85.1","State":"BGP_FSM_OPENCONFIRM","Topic":"Peer","level":"info","msg":"Peer
Up","time":"2022-03-18T16:32:06+01:00"}
[gacruX@arturo-vm3 ~]$

```

Stop the Service

To stop the service, use the following command:

```
[gacrux@arturo-vm3 ~]$ sudo systemctl stop gobgpd.service
```

Disable the Service

To disable the service, use the following command:

```
[gacrux@arturo-vm3 ~]$ sudo systemctl disable gobgpd.service
```

Recommended Hardware

The hardware recommendation is based on the Cisco UCS configuration with the VMWare ESXi-7.0U1c.



Note In VMWare ESXi, configure the Security Policy of virtual switches as follow:

```
Allow promiscuous mode: yes
Allow forged transmits: yes
Allow MAC changes: yes
```

- Cisco UCS-C220-M5SX
- 32 CPUs (2x Intel Xeon Gold 5218 CPU @ 2.30GHz)
- 256GB RAM
- 2x 480GB SSD (Raid)
- MLOM NIC for additional network interfaces

Recommended Software

Table 2: Feature History

| Feature Name | Release Information | Feature Description |
|--|-----------------------------|---|
| Support for Hosting SVO Server on Red Hat Enterprise Linux | Cisco NCS 2000 Release 12.3 | This feature allows you to host SVO on an external server running Red Hat Enterprise Linux 7.9. |

The following are the recommended OS versions.

- CentOS Linux release 7.9.2009 (Docker Engine 20.10.9)
- Red Hat Enterprise Linux 7.9 (Docker Engine 20.10.9)

Recommended Resource for Virtual Machines

The following sections provide information on the resources that are required to define SVO instances.

In case of high availability, it is recommended to have the same resources on both the VMs.

Release 12.3.1

Release 12.3.1 introduces effective resource management, which has a significant impact on resource optimization. As a reference of load, a VM created on a UCS server with 64 virtual CPUs, 240 GB RAM and 400 GB of available disk space, can accommodate admin plane and 90 SVO instances (the verification is done with 60 percent of ROADM instances and 40 percent of OLA instances).

- **CPU**

Seven virtual CPUs for 10 SVO instances (including admin plane)

- **Memory**

- 2 GB for admin plane
- 2.1 GB for each node, irrespective of the SVO type such as ROADM and OLA



Note We recommend you to use 80 percent of the total available memory of the server or VM to create SVO instances, and use the remaining 20 percent as shared resources to accommodate memory peaks during heavy operations.

- **Disk Space**

2.5 GB for each SVO instance



Note SVO application stores data in the root (/) filesystem (/var and /misc folders). It is in charge of the administration of the server/VM to apply the proper partition (advanced disk options during installation).

Release 12.3 and Earlier

- **CPU**

- One virtual CPU for each ROADM node
- One virtual CPU for every three OLA, DGE, or TXP nodes

- **Memory**

- 2 GB for admin plane
- 3 GB for each OLA, DGE, TXP, or ROADM with 2, 3 or 4 degrees
- 4 GB for each ROADM with 5, 6, 7 or 8 degrees

- 8 GB for each ROADM with more than 8 degrees

- **Disk Space**

- 5 GB for each SVO instance

Required Network Resources

Network resources that must be planned in the design phase are related to three different subnets.

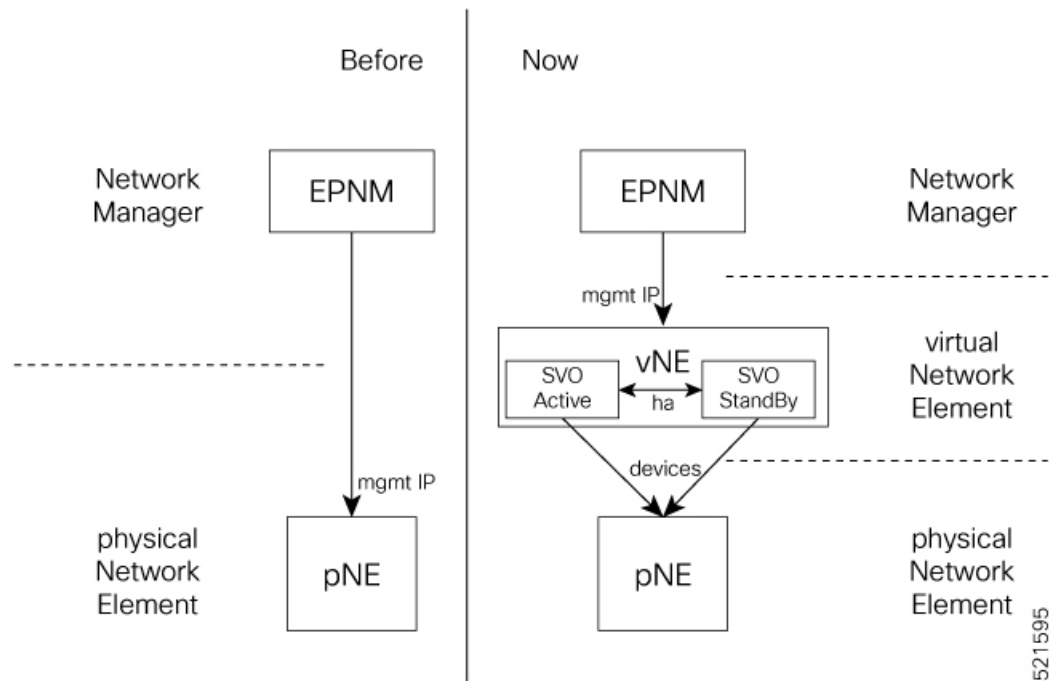
The following table describes the different networks.

Table 3: Network Type

| Network | Network Type | Description |
|---------|-----------------------------------|--|
| mgmt | Management network | Offers Netconf NBI and Web UIs for admin plane and SVO instances (used by users) |
| ha | High availability network | For data replication Note Used only between two servers or VMs |
| devices | Network to reach NCS 2000 devices | Communicate with NCS 2000 devices Note Used only between the servers or VMs and the NCS 2000 devices |

The following figure illustrates the management of an NE before and after the introduction of SVO.

Figure 1: Network Element Before and After SVO



Note For each virtual NE, represented by a SVO instance, there is a single management IP address. The status of the management interface, Up or Down, is aligned with the status of the SVO, active or standby. If SVO is active, the management interface is Up. If SVO is standby, the management interface is Down. The status, active or standby, is orchestrated by admin plane.

In the following table, the data to define the size of the subnets to set up HA network and the description of each network and its usage.

| Network | No. of IP Addresses | Description |
|------------|--|--|
| management | One for each gateway Two for each host NIC Two for each admin plane One for each SVO instance | <p>Must be the same subnet on both the servers or VMs.</p> <p>On this subnet, SVO software application offers Netconf NBI and Web UIs.</p> <p>When used from the customer DCN, this network becomes visible and routable in the customer DCN.</p> <p>Admin planes work in active/active status, that is, on each server or VM, when an admin plane is running, it is possible to connect to the Web UI through the management IP address (https://mgmt-ip).</p> <p>Each SVO instance works in active/standby status. The instances share the same management IP address.</p> <p>Note The network interface is Down on the standby side.</p> |

| Network | No. of IP Addresses | Description |
|---------|--|---|
| HA | One for each gateway Two for each host NIC Two for each admin plane Two for each SVO instance | In release 12.1, HA network must have the same subnet on both the servers or VMs. From release 12.2, HA network can have different and routable subnets. Allows data replication between active and standby SVO instances. It is the primary communication channel between admin planes. |
| Devices | One for each gateway Two for each host NIC Two for each admin plane Two for each SVO instance | In release 12.1, Devices network must have the same subnet on both the servers and the VMs. From release 12.2, the subnets can be different. Devices shall be routable towards NCS 2000 network. Used by SVO instances to communicate with NCS 2000 devices. It is the secondary communication channel between admin planes. |

Bandwidth and Latency Requirements

The bandwidth required for the high availability (HA) networks ranges from 255 to 977 Mbps, depending on the workload of the server.

The latency requirements are:

- EPNM from and to SVO is < 80 ms
- SVO servers from and to NCS 2000 devices is < 80 ms
- SVO servers high availability is < 100 ms

Install Docker Engine

Use this task to install the docker engine.



Note During the installation on RHEL, the following error message could be generated during the execution of the command `sudo yum install docker-ce-<version> ...`:

```
*****
yum can be configured to try to resolve such errors by temporarily enabling
disabled repos and searching for missing dependencies.
To enable this functionality please set 'notify_only=0' in
/etc/yum/pluginconf.d/search-disabled-repos.conf
*****
```

Edit the file `/etc/yum/pluginconf.d/search-disabled-repos.conf` and set `notify_only=0`, then execute again the command.

Procedure

Step 1 Check repositories for available package updates.

```
[gacrux@arturo-vm3 ~]$ sudo yum check-update
Loaded plugins: fastestmirror
Determining fastest mirrors
 * base: ams.edge.kernel.org
 * extras: ams.edge.kernel.org
 * updates: ams.edge.kernel.org
...
```

Step 2 Install the required dependencies from the Docker repositories.

```
[gacrux@arturo-vm3 ~]$ sudo yum install -y yum-utils device-mapper-persistent-data lvm2
Loaded plugins: fastestmirror
Loading mirror speeds from cached hostfile
 * base: ams.edge.kernel.org
 * extras: ams.edge.kernel.org
 * updates: ams.edge.kernel.org
Resolving Dependencies
--> Running transaction check
---> Package device-mapper-persistent-data.x86_64 0:0.8.5-3.e17 will be updated
---> Package device-mapper-persistent-data.x86_64 0:0.8.5-3.e17_9.2 will be an update
---> Package lvm2.x86_64 7:2.02.187-6.e17 will be updated
...
Installed:
  yum-utils.noarch 0:1.1.31-54.e17_8

Dependency Installed:
  libxml2-python.x86_64 0:2.9.1-6.e17.5          python-chardet.noarch 0:2.2.1-3.e17
  python-kitchen.noarch 0:1.1.1-5.e17

Updated:
  device-mapper-persistent-data.x86_64 0:0.8.5-3.e17_9.2          lvm2.x86_64
  7:2.02.187-6.e17_9.3

Dependency Updated:
  device-mapper.x86_64 7:1.02.170-6.e17_9.3          device-mapper-event.x86_64
  7:1.02.170-6.e17_9.3
  device-mapper-event-libs.x86_64 7:1.02.170-6.e17_9.3          device-mapper-libs.x86_64
  7:1.02.170-6.e17_9.3
  lvm2-libs.x86_64 7:2.02.187-6.e17_9.3
```

Complete!

Step 3 Set up the Docker repository.

```
[gacrux@arturo-vm3 ~]$ sudo yum-config-manager --add-repo
https://download.docker.com/linux/centos/docker-ce.repo
Loaded plugins: fastestmirror
adding repo from: https://download.docker.com/linux/centos/docker-ce.repo
grabbing file https://download.docker.com/linux/centos/docker-ce.repo to
/etc/yum.repos.d/docker-ce.repo
repo saved to /etc/yum.repos.d/docker-ce.repo
```

Step 4 Install the 20.10.9 version of Docker engine and containerd.

```
[gacrux@arturo-vm3 ~]$ sudo yum install docker-ce-20.10.9 docker-ce-cli-20.10.9 containerd.io
Loaded plugins: langpacks, product-id, search-disabled-repos, subscription-manager
Resolving Dependencies
--> Running transaction check
---> Package containerd.io.x86_64 0:1.4.10-3.1.el7 will be installed
--> Processing Dependency: container-selinux >= 2:2.74 for package:
containerd.io-1.4.10-3.1.el7.x86_64
---> Package docker-ce.x86_64 3:20.10.9-3.el7 will be installed
...
Installed:
  containerd.io.x86_64 0:1.4.10-3.1.el7                docker-ce.x86_64 3:20.10.9-3.el7
  docker-ce-cli.x86_64 1:20.10.9-3.el7

Dependency Installed:
  container-selinux.noarch 2:2.119.2-1.911c772.el7_8      docker-ce-rootless-extras.x86_64
  0:20.10.9-3.el7      docker-scan-plugin.x86_64 0:0.8.0-3.el7
  fuse-overlayfs.x86_64 0:0.7.2-6.el7_8                      fuse3-libs.x86_64 0:3.6.1-4.el7
  slirp4netns.x86_64 0:0.4.3-4.el7_8
```

Complete!

Step 5 Start the Docker engine.

```
[gacrux@arturo-vm3 ~]$ systemctl start docker
==== AUTHENTICATING FOR org.freedesktop.systemd1.manage-units ====
Authentication is required to manage system services or units.
Authenticating as: gacrux
Password:
==== AUTHENTICATION COMPLETE ====
```

Step 6 Enable the Docker engine.

```
[gacrux@arturo-vm3 ~]$ systemctl enable docker
==== AUTHENTICATING FOR org.freedesktop.systemd1.manage-unit-files ====
Authentication is required to manage system service or unit files.
Authenticating as: gacrux
Password:
==== AUTHENTICATION COMPLETE ====
Created symlink from /etc/systemd/system/multi-user.target.wants/docker.service to
/usr/lib/systemd/system/docker.service.
==== AUTHENTICATING FOR org.freedesktop.systemd1.reload-daemon ====
Authentication is required to reload the systemd state.
Authenticating as: gacrux
Password:
==== AUTHENTICATION COMPLETE ====
```

Step 7 Add the user to the Docker group.

```
[gacrux@arturo-vm3 ~]$ sudo usermod -aG docker $USER
```

Step 8 Exit the command line.

```
[gacrux@arturo-vm3 ~]$ exit
```

Step 9 Disconnect and reconnect to refresh user permissions.

Network Configuration File

Network configuration file defines all the information that is related to the servers and the networking infrastructure that is required to set up the SVO functionalities. See [Required Network Resources](#) section to define the networks.

The *network.yml* configuration file starts with the **server-name** field. The **server-name** field refers to the **name** field of the server in which the installation executes.

The configuration files that are used for the installation in high availability networks on two servers or VMs contain different values in the **server-name** field for the respective servers.

The *network.yml* configuration file contains two main sections, one for each server or VM of the cluster.

For each server section, there are three fields:

- **name**—Server name (editable)
- **mgmt-port**—Not editable (value is 443)
- **ha-agent-port**—Not editable (value is 5480)

SVO network architecture is based on four networks. Each section in a *network.yml* file defines the following networks:

- **Management**—The management ports of the admin plane and SVO instances belong to this network.
- **High Availability**—This network is used for SVO database replication and communication between the admin planes.
- **Private HA Network**—This network is used only for communicating between the admin plane and the local SVO instances in certain processes such as node activation. It doesn't use external networking resources.
- **Devices**—This network is used for connecting the NCS 2000 devices. The devices must be reachable through this subnet. If the admin plane is unable to communicate with its peer server on the primary high availability network, it uses a link on this network as an alternative way to communicate.

The **Private HA Network** is local to the server or VM. The other three networks are exposed outwards and they must have three different subnets.



Note Ensure that IP addresses from the *172.16.xx.xx/16* subnet are not specified in *.yml* files as this subnet is reserved for Docker's internal use.

The following table describes the fields in the *network.yml* configuration files.

| Field | Value | Editable | Description |
|-------------------------|---|----------|---|
| server-name | — | Yes | Name of the server to which the file is applied. |
| mgmt-address-family | IPv4 IPv4_IPv6 IPv6 | Yes | IP address family of the management network. |
| name | management hanetwork haprivate devices | No | Name of the network |
| mgmt-port | 443 | No | Port that is used for admin plane web UI |
| ha-agent-port | 5480 | No | Port that is used for internal communication between the admin plane and SVO |
| host-nic-name | — | Yes | (Optional) Interface that is used in docker network Note Only in management , hanetwork , and devices |
| ha-port | 10001 10002 | Yes | Port used for admin plane high-availability on the specific network (primary path) Note Only in HA and Devices . |
| IPv4 | | | |
| ip | — | Yes | Subnet |
| prefix | — | Yes | Mask |
| gateway | — | Yes | Gateway |
| host-nic | — | Yes | (Optional) IPv4 address that is assigned to the host interface Note Only in management , hanetwork , and devices |
| adminplane | — | Yes | IPv4 address that is assigned to the admin plane |
| IPv6: (Optional) | | | |
| ip | — | Yes | Subnet |
| prefix | — | Yes | Mask |
| gateway | — | Yes | Gateway |
| host-nic | — | Yes | (Optional) IPv6 address that is assigned to the host interface Note Only in management , hanetwork , and devices |

| Field | Value | Editable | Description |
|------------|-------|----------|--|
| adminplane | — | Yes | IPv6 address that is assigned to the admin plane |

The following sections contain one example of *network.yml* file for IPv4 configuration and one for IPv6 configuration.

To create the configuration file, use one of the following templates.

IPv4 Network YAML Configuration File Sample

```
server-name: VM1

servers:
- name: VM1
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens192
    ipv4:
      ip: 10.58.233.0
      prefix: 24
      gateway: 10.58.233.1
      host-nic: 10.58.233.75
      adminplane: 10.58.233.76
  ha:
    name: hanetwork
    host-nic-name: ens224
    ha-port: 10001
    ipv4:
      ip: 192.168.1.0
      prefix: 24
      gateway: 192.168.1.1
      adminplane: 192.168.1.4
      host-nic: 192.168.1.2
  haprivate:
    name: haprivate
    ipv4:
      ip: 192.168.3.0
      prefix: 24
      gateway: 192.168.3.1
      adminplane: 192.168.3.2
  devices:
    name: devices
    host-nic-name: ens256
    ha-port: 10002
    ipv4:
      ip: 192.168.2.0
      prefix: 24
      gateway: 192.168.2.1
      adminplane: 192.168.2.4
      host-nic: 192.168.2.2

- name: VM2
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens192
    ipv4:
```

```

    ip: 10.58.233.0
    prefix: 24
    gateway: 10.58.233.1
    host-nic: 10.58.233.77
    adminplane: 10.58.233.78
  ha:
    name: hanetwork
    host-nic-name: ens224
    ha-port: 10001
    ipv4:
      ip: 192.168.1.0
      prefix: 24
      gateway: 192.168.1.1
      adminplane: 192.168.1.5
      host-nic: 192.168.1.3
  haprivate:
    name: haprivate
    ipv4:
      ip: 192.168.3.0
      prefix: 24
      gateway: 192.168.3.1
      adminplane: 192.168.3.2
  devices:
    name: devices
    host-nic-name: ens256
    ha-port: 10002
    ipv4:
      ip: 192.168.2.0
      prefix: 24
      gateway: 192.168.2.1
      adminplane: 192.168.2.5
      host-nic: 192.168.2.3

```

IPv6 Network YAML Configuration File Sample

Networking infrastructure requires all the *network.yml* files with IPv6 section to have an IPv4 section.

In the following example, the IPv6 section has been configured for the three networks: management, high availability, and devices.

It is not mandatory that all networks must be configured in IPv6. Each network is independent of others.

Table 4: Network Types

| Network Types | Description |
|---------------|--|
| Management | <p>Dual stack (IPv4+IPv6)—User wants to use both IPv4 and IPv6 addresses. It is required that the same IPv4 subnet is used between the two servers</p> <p>IPv6 only—User wants to use only IPv6 addresses. Use different private IPv4 subnets between the two servers, to avoid any conflict</p> <p>In both scenarios, the IPv4 information are required by the networking infrastructure.</p> |

| Network Types | Description |
|-------------------------------|---|
| High availability and devices | When IPv6 information are available, IPv4 information are not used directly by the SVO application but are anyway required by the networking infrastructure. The suggestion is to use different private IPv4 subnets between the two servers, to avoid any conflict. |
| Private HA Network | It is a local network. You must use the IPv4 information as required by the networking infrastructure. There is no reason to add IPv6 section. |

```
server-name: VM1
```

```
servers:
```

```
- name: VM1
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens192
    ipv4:
      ip: 10.58.233.0
      prefix: 24
      gateway: 10.58.233.1
      adminplane: 10.58.233.72
    ipv6:
      ip: 2001:420:4491:2004::233:0
      prefix: 112
      gateway: 2001:420:4491:2004::233:2
      host-nic: 2001:420:4491:2004::233:70
      adminplane: 2001:420:4491:2004::233:72
```

```
ha:
```

```
  name: hanetwork
  host-nic-name: ens224
  ha-port: 10001
  ipv4:
    ip: 192.168.80.0
    prefix: 24
    gateway: 192.168.80.1
    adminplane: 192.168.80.2
  ipv6:
    ip: 2001:db8:abc:111::0
    prefix: 64
    gateway: 2001:db8:abc:111::1
    adminplane: 2001:db8:abc:111::4
    host-nic: 2001:db8:abc:111::2
```

```
haprivate:
```

```
  name: haprivate
  ipv4:
    ip: 192.168.3.0
    prefix: 24
    gateway: 192.168.3.1
    adminplane: 192.168.3.2
```

```
devices:
```

```
  name: devices
  host-nic-name: ens256
  ha-port: 10002
  ipv4:
```

```

        ip: 192.168.81.0
        prefix: 24
        gateway: 192.168.81.1
        adminplane: 192.168.81.2
    ipv6:
        ip: 2001:db8:abc:222::0
        prefix: 64
        gateway: 2001:db8:abc:222::1
        adminplane: 2001:db8:abc:222::4
        host-nic: 2001:db8:abc:222::2
- name: VM2
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens192
    ipv4:
      ip: 10.58.233.0
      prefix: 24
      gateway: 10.58.233.1
      adminplane: 10.58.233.73
    ipv6:
      ip: 2001:420:4491:2004::233:0
      prefix: 112
      gateway: 2001:420:4491:2004::233:2
      host-nic: 2001:420:4491:2004::233:71
      adminplane: 2001:420:4491:2004::233:73
  ha:
    name: hanetwork
    host-nic-name: ens224
    ha-port: 10001
    ipv4:
      ip: 192.168.83.0
      prefix: 24
      gateway: 192.168.83.1
      adminplane: 192.168.83.3
    ipv6:
      ip: 2001:db8:abc:111::0
      prefix: 64
      gateway: 2001:db8:abc:111::1
      adminplane: 2001:db8:abc:111::5
      host-nic: 2001:db8:abc:111::3
  haprivate:
    name: haprivate
    ipv4:
      ip: 192.168.3.0
      prefix: 24
      gateway: 192.168.3.1
      adminplane: 192.168.3.3
  devices:
    name: devices
    host-nic-name: ens256
    ha-port: 10002
    ipv4:
      ip: 192.168.84.0
      prefix: 24
      gateway: 192.168.84.1
      adminplane: 192.168.84.3
    ipv6:
      ip: 2001:db8:abc:222::0
      prefix: 64
      gateway: 2001:db8:abc:222::1

```

```
adminplane: 2001:db8:abc:222::5
host-nic: 2001:db8:abc:222::3
```

IPv4 Port Forwarding Network YAML Configuration

```
...

servers:
- name: VM1
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens192
    ipv4:
      ip: 10.58.233.0
      prefix: 24
      gateway: 10.58.233.1
      host-nic: 10.58.233.75
      adminplane: 10.58.233.75
  ha:
    name: hanetwork
    host-nic-name: ens224
    ha-port: 10001
    ipv4:
      ip: 192.168.1.0
      prefix: 24
      gateway: 192.168.1.1
      host-nic: 192.168.1.2
      adminplane: 192.168.1.4

...
```

L3 Management Interconnection Network YAML Configuration

Layer 3 management interconnection between servers simplifies Geo redundancy deployment.



Note A new field is added in the *network.yml* file to configure the management interconnection at Layer 3. The following field is an optional field with the default value set as **LAYER2**.

```
mgmt-interconnection: LAYER3
```

The following *network.yml* file contains the IP addresses used in the [Use Case 3 - Dislocated Servers \(L3 Interconnection\)](#), on page 34 section.

```
server-name: VM1
mgmt-address-family: IPv4_IPv6
mgmt-interconnection: LAYER3
servers:
- name: VM1
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens161
    ipv4:
      ip: 10.58.253.0
      prefix: 27
      gateway: 10.58.253.1
      host-nic: 192.168.85.2
```

```

    adminplane: 10.58.253.2
  ipv6:
    ip: 2000::10:58:253:0
    prefix: 123
    gateway: 2000::10:58:253:1
    host-nic: fd00::192:168:85:2
    adminplane: 2000::10:58:253:2
ha:
  name: hanetwork
  host-nic-name: ens224
  ha-port: 10001
  ipv4:
    ip: 192.168.85.32
    prefix: 27
    gateway: 192.168.85.33
    host-nic: 192.168.85.34
    adminplane: 192.168.85.35
  ipv6:
    ip: fd00::192:168:85:20
    prefix: 123
    gateway: fd00::192:168:85:21
    host-nic: fd00::192:168:85:22
    adminplane: fd00::192:168:85:23
haprivate:
  name: haprivate
  ipv4:
    ip: 192.168.85.96
    prefix: 27
    gateway: 192.168.85.97
    adminplane: 192.168.85.98
devices:
  name: devices
  host-nic-name: ens256
  ha-port: 10002
  ipv4:
    ip: 192.168.85.64
    prefix: 27
    gateway: 192.168.85.65
    host-nic: 192.168.85.66
    adminplane: 192.168.85.67
  ipv6:
    ip: fd00::192:168:85:40
    prefix: 123
    gateway: fd00::192:168:85:41
    host-nic: fd00::192:168:85:42
    adminplane: fd00::192:168:85:43
- name: VM2
  mgmt-port: 443
  ha-agent-port: 5480
  mgmt:
    name: management
    host-nic-name: ens161
    ipv4:
      ip: 10.58.253.0
      prefix: 27
      gateway: 10.58.253.1
      host-nic: 172.16.16.2
      adminplane: 10.58.253.3
    ipv6:
      ip: 2000::10:58:253:0
      prefix: 123
      gateway: 2000::10:58:253:1
      host-nic: fd00::172:16:16:2

```

```
    adminplane: 2000::10:58:253:3
ha:
  name: hanetwork
  host-nic-name: ens224
  ha-port: 10001
  ipv4:
    ip: 172.16.16.32
    prefix: 27
    gateway: 172.16.16.33
    host-nic: 172.16.16.34
    adminplane: 172.16.16.35
  ipv6:
    ip: fd00::172:16:16:20
    prefix: 123
    gateway: fd00::172:16:16:21
    host-nic: fd00::172:16:16:22
    adminplane: fd00::172:16:16:23
haprivate:
  name: haprivate
  ipv4:
    ip: 172.16.16.96
    prefix: 27
    gateway: 172.16.16.97
    adminplane: 172.16.16.98
devices:
  name: devices
  host-nic-name: ens256
  ha-port: 10002
  ipv4:
    ip: 172.16.16.64
    prefix: 27
    gateway: 172.16.16.65
    host-nic: 172.16.16.66
    adminplane: 172.16.16.67
  ipv6:
    ip: fd00::172:16:16:40
    prefix: 123
    gateway: fd00::172:16:16:41
    host-nic: fd00::172:16:16:42
    adminplane: fd00::172:16:16:43
```

Prepare the Network Configuration

Use this task to prepare the network configuration for the external server model or the SVO card model.

Before you begin

See [Network Configuration File](#), on page 13.

Procedure

-
- Step 1** Create a configuration file (*network.yml*) for each server with the networking configuration data. This file is uploaded during the installation.
 - Step 2** Cable and configure the related network interfaces (physical or virtual).

Both the configuration files are identical and contain the data for both the servers in HA. The only difference is the *server-name* attribute that contains the name of the server to which the file is applied.

Standalone SVO Configuration

From Release 12.3, you can configure and install SVO in Standalone mode. Independent of IPv4 or IPv6 configuration, the network YAML file contains information of a single server, instead of the local and remote instances.

Limitation of Standalone Mode

An SVO installed in Standalone mode cannot be upgraded to high availability. The SVO application must be re-installed with a standard network YAML file containing the information of both servers.

Installation of SVO

The SVO application can be installed in the following setups:

- [Install SVO Card, on page 22](#)
- [Install the External Server, on page 23](#)

The SVO software installation for the external server can be separated in different steps:

- Creation of the VM and resource allocation
- Installation of OS and Docker packages
- Installation of SVO software

VM Creation and Resource Allocation

Create a VM allocating the proper resources, in terms of CPU, memory and disk space.

Installation of OS and Docker packages

Installation steps described in the following sections refer to CentOS 7.9.2009 and RHEL 7.9.

Install SVO Card



Note The SVO card model comes pre-installed with the required software packages. You need to power up the card and start the configuration.

Use this task to install the SVO application using the SVO card.

Before you begin

Ensure following recommendations are met:

- [Recommended Hardware, on page 6](#)
- [Recommended Software, on page 6](#)
- [Recommended Resource for Virtual Machines, on page 7](#)
- [Required Network Resources, on page 8](#)

Procedure

- Step 1** Power up the SVO card.
- Step 2** Run the SVO installation tool. See [Bring Up Admin Plane with the SVO Installation Tool, on page 26](#).
-

Install the External Server

Before you begin installing the SVO application using the installation script or manual installation, ensure the following recommendations are met.

- [Recommended Hardware, on page 6](#)
- [Recommended Software, on page 6](#)
- [Recommended Resource for Virtual Machines, on page 7](#)
- [Required Network Resources, on page 8](#)
- [Install Docker Engine, on page 10](#)

You can install the SVO application in the following methods:

- [Install SVO in the External Server Using the Installation Script, on page 23](#)

This method enables you to install the SVO application using the **svoTools-12.3.1.sh** script. When you install the SVO application using the installation script, you can skip the manual installation procedure.

- [Install SVO in the External Server Manually, on page 25](#)

This method enables you to install the SVO application using the CLI commands and docker images.

Install SVO in the External Server Using the Installation Script

The **svoTools-12.3.1.sh** script enables automatic installation of the SVO application in the server for Release 12.3.1. You need not use the Docker engine to install the SVO application. The script is available for download within **svo-utilities-12.3.0.tar** file that is located in the Utilities folder of SVO Release 12.3.1 at the [Cisco Software Downloads page](#). The tar archive file contains the related signed RPM with instructions to extract the actual script file. This utility script simplifies the following operations:

1. Extracts and loads SVO images
2. Brings up and starts SVO installation tool

3. Uninstalls SVO



Note `svoTools-12.3.1.sh` script can be used to install the SVO application in the server only.



Note If any compatibility issue arises between the selected operating system and the script, manually install the SVO application.



Note If the management network configuration is **IPv4/IPv6** or **IPv6-only**, the `svoTools-12.3.1.sh` script requires subnet, mask, and gateway for both the stacks.

The difference between the **IPv4/IPv6** and **IPv6-only** configurations is the IPv4 information. For the IPv4/IPv6 dual stacks management network, the same public IPv4 subnet network information must be used between the two servers. For an IPv6-only management network, two different private IPv4 subnet network information must be used between the two servers.

Procedure

Run the `sudo ./svoTools-<version>.sh` command as an admin.

```
[gacrux@arturo-vm2 LG]$ sudo ./svoTools-12.3.1.sh
=====
=====

          SVO Tools 1.3 (compatible with SVO 12.3.1)

Cisco Systems, Inc.                                     Copyright 2022
=====

Welcome to the SVO Tools

This script simplifies the following operations:
* Load SVO images from SW release (ncs2k-server-12.3.1_REL.tar)
* Start SVO Installation Tool
* Uninstall SVO

In any moment you can abort the operation pressing CTRL-C
-----

Enter your choice:

1) Load SVO images
2) Start SVO
3) Uninstall SVO
4) Exit
#?

Type the number of your choice.
```


Load SVO images require downloading the SVO software release from Cisco.com (as a TAR archive file) in the same folder where the script is executing.

Start SVO option allows starting the SVO Installation Tool. After selecting the **Start SVO Tool** option, see [Bring Up Admin Plane with the SVO Installation Tool, on page 26](#) about how to connect via browser to the tool.

Uninstall SVO allows uninstalling the SVO files.

Note From Release 12.3.1, the `svoTools-12.3.1.sh` allows you to choose between **Layer2** and **Layer3** management interconnection. Select **Layer3** option to enable Layer 3 Management Interconnection solution when the following prompt appears:

```
Select management interconnection between servers at Layer2 or Layer3:
1) Layer2
2) Layer3
#? 2
```

Install SVO in the External Server Manually



Note If you successfully completed the SVO application installation using the [Install SVO in the External Server Using the Installation Script, on page 23](#) procedure, you can skip the following procedure.

Use this task to install the external server.

Procedure

- Step 1** Create the network configuration. See [Prepare the Network Configuration, on page 21](#).
- Step 2** Obtain the admin plane (es-admin-plane) and SVO (svo-dos) docker images. Perform these steps:
- Get the `ncs2k-server-12.1.0_REL.tar` file and extract its contents.


```
> tar xvf ncs2k-server-12.1.0_REL.tar
> cd SIGNED_RPM; ls
NCS2K-S-1210.020K.2311.x86_64.rpm
es-admin-plane-12.1.0.B0582.x86_64.rpm
svo-dos-12.1.0.R0582.x86_64.rpm
> rpm2cpio es-admin-plane-12.1.0.B0582.x86_64.rpm | cpio -D <OUTDIR> -idmv
> rpm2cpio svo-dos-12.1.0.R0582.x86_64.rpm | cpio -D <OUTDIR> -idmv
```
 - Load the admin plane and SVO images.


```
docker load -i es-admin-plane.tgz
docker load -i svo-dos.tgz
```
 - Verify the images using the **docker images** command.
- Step 3** Create the configuration folder `/misc/disk1/data/adminplane`. The folder path need to used as a parameter the next step.
- Step 4** Create a text file called `installer.properties` inside the configuration folder. Add the the following line to the file.

```
server.address=<adminplane-mgmt-ip>
installer.remote-connection=true
```

The `adminplane-mgmt-ip` is the IP address that is assigned to the admin plane in the management network.

Step 5 Create a docker network using any one of the following commands:

- **IPv4**

```
docker network create -d macvlan --attachable --subnet <subnet>/<mask> --gateway <gw>
-o parent=<mgmt-interface> management
```

```
docker create --name adminplane --network management --ip <adminplane-mgmt-ip> -m 2g
--memory-swap 2g --cap-add NET_ADMIN --restart
always -v /misc/disk1/data:/misc/disk1/data -v /var/run/docker.sock:/var/run/docker.sock
-v /misc/disk1/data/adminplane:/opt/config es-admin-plane:<version>
```

- **IPv4 and IPv6**

```
docker network create -d macvlan --attachable --subnet <subnet>/<mask> --gateway <gw>
--ipv6 --subnet <ipv6-subnet>/<prefix> --gateway <ipv6-gw> -o parent=<mgmt-interface>
management
```

```
docker create --name adminplane --network management --ip <adminplane-mgmt-ip> --ip6
<adminplane-mgmt-ipv6> -m 2g --memory-swap 2g --cap-add NET_ADMIN --restart
always -v /misc/disk1/data:/misc/disk1/data -v /var/run/docker.sock:/var/run/docker.sock
-v /misc/disk1/data/adminplane:/opt/config es-admin-plane:<version>
```

Step 6 Start the application using the following command:

```
docker start adminplane
```

Step 7 Run the SVO installation tool. See [Bring Up Admin Plane with the SVO Installation Tool, on page 26](#).

Bring Up Admin Plane with the SVO Installation Tool

Use this task to bring up the Admin Plane with the SVO installation tool for the external server model or the SVO card model.



Note In the SVO card model, the user created is the superuser in all the SVO instances created later.

Procedure

Step 1 Start the SVO installation tool using the IP address as follows:

- **External server model**—Use the management IP address, `http://adminplane-mgmt-ip`.
- **SVO card model**—Use the pre-defined IP address, `http://192.168.0.66`.

Note The client must be connected to the SVO craft port using an ethernet cable.

Step 2 In the **Credentials** area, perform these steps:

- a. Enter a username in the **Username** field.
- b. Enter a password in the **Password** field.
The password must be a minimum of eight characters, and it can be a maximum of 127 characters. The password must have at least one uppercase letter, one lowercase character, one number, and one special character.
- c. Retype the password in the **Retype Password** field.

Step 3 In the **Networks** area, perform these steps:

- a. Click **Browse** to select the *network.yml* configuration file **Configuration File** field.
For more information about the configuration file (*network.yml*), see [Prepare the Network Configuration, on page 21](#).

Step 4 Click **Submit**.

The system creates the credentials, verifies the network configuration file and brings up the system. You are now able to connect to Cisco SVO admin plane login page at <https://adminplane-mgmt-ip>. See [Log into the Cisco SVO Admin Plane](#).

Deployment

The SVO networking architecture offers a degree of flexibility that meets all requirements.

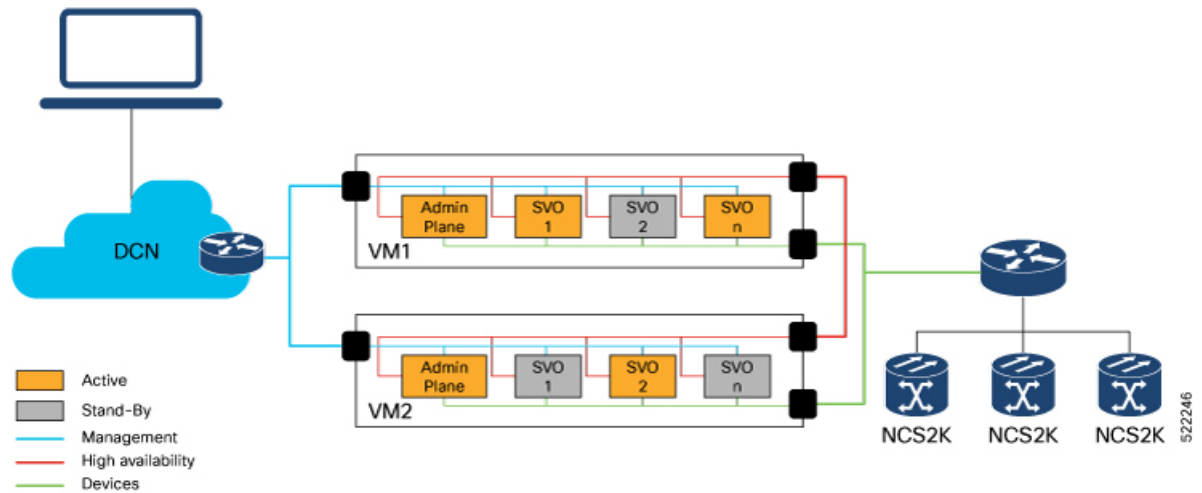
The following list describes the types of deployment:

- [Deployment of Colocated Servers](#)
- [Deployment of Servers in Different Locations \(L2 Interconnection\)](#)

Deployment of Colocated Servers

Deployment of the colocated servers is simple. In this deployment, both the servers or VMs share the same subnets for all three networks. The following image displays the schema of the networks and the connections.

Figure 2: Deployment of Colocated Servers



Deployment of Servers in Different Locations (L2 Interconnection)

Deployment of servers in different locations is complicated. This deployment needs more attention in the design and more checks in the existing subnet.

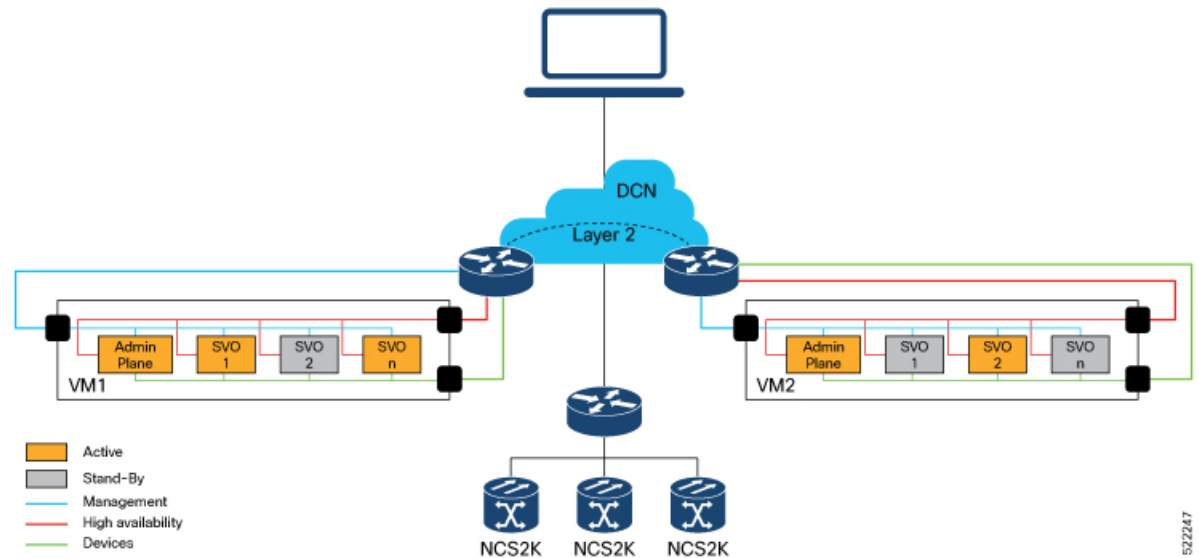
Management subnet reaches the servers in different locations. A Layer 2 connection is required for both the servers and VMs in different locations. The valid Layer 2 connections are L2VPN, L2TPv3, VXLAN, or a similar solution.

In Release 12.1, some limitations can generate more complexities for this kind of deployment. The **high availability** networks on the two servers must be in the same subnets. The same limitation impacts the **devices** networks on the two servers. Cisco recommends following the same approach that is selected for the **management** network. It is valid for both **high availability** and **devices** networks.

From Release 12.2, the high availability and devices networks can be configured in different routable subnets.

The following figure shows the schema of the networks and the connections.

Figure 3: Deployment of Servers in Different Locations



532247

Deployment of Servers in Different Locations (L3 Interconnection)

From Release 12.3.1, deployment of servers in different locations with L3 management interconnection is supported. It enables a Layer 3 management interconnection between servers or VMs that are deployed in different locations.

A single and common **management** subnet is used only by the SVO Admin Plane. The management network interfaces of the servers or VMs have IP addresses on different subnets.

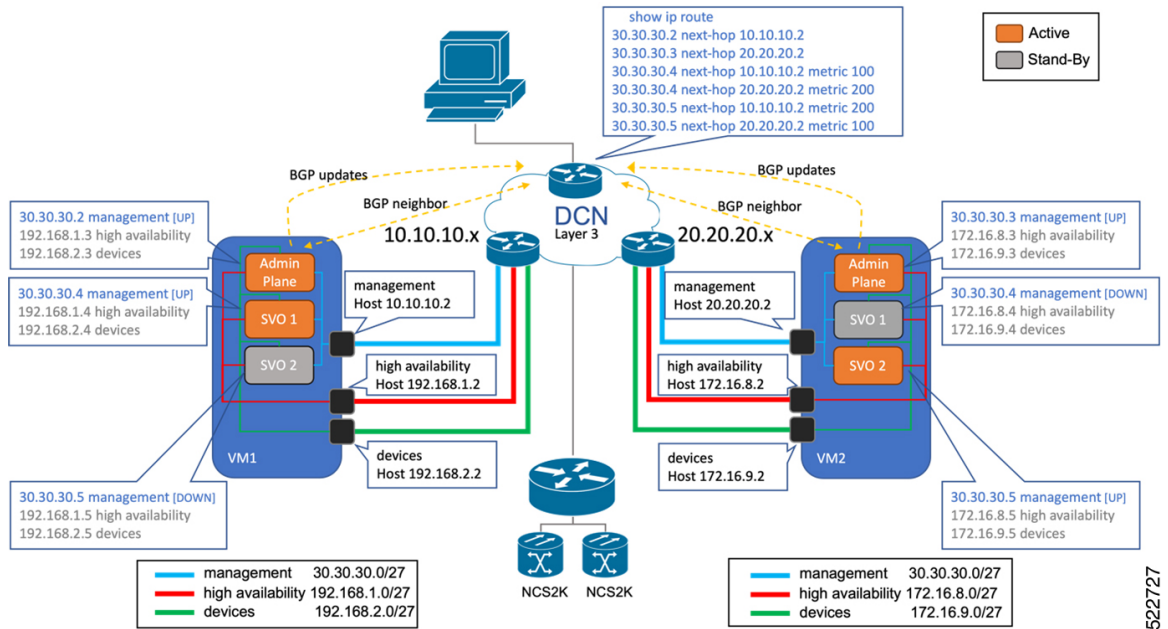
The two servers or VMs establish a BGP neighbor relationship with a core router.

SVO Admin Plane advertises the route that is related to itself and each SVO instance running in Active mode on the local server, with next-hop address pointing to the local server host management IP address.

The IP addresses of the servers are not visible to you, only the SVO management subnet is visible. In the following example, the management subnet is **30.30.30.0/27**.

The following image displays the schema of connections, BGP relationships, and networks advertisements.

Figure 4: Deployment of Servers in Different Locations (L3 Interconnection)



Admin Planes constantly monitor the BGP daemon running on each server and exchange information between themselves about the BGP neighbor relationship status. The status is displayed on the SVO Admin Plane Web UI (near the high availability bell icon) with two colored dots over a router icon.

In case of failure in the BGP neighbor relationship on one server, a switchover of all Active SVO instances is triggered toward the server where the BGP neighbor relationship is still established.

You can configure BGP communities per server through the Web UI or through the configuring file. For configuration in Web UI, refer to [Edit Admin Plane Properties](#).



Note You can configure the BGP communities in the `adminplane.properties` file from the `/misc/disk1/data/adminplane/adminplane.properties` path. Enter a list of elements that are separated by commas as shown in the following example:

```
adminplane.bgp.advertiser.communities=100,102
```

Disaster Recovery

This section explains the analysis of what happens during a disaster, where a server or VM failed, disconnected, became unreachable or unavailable, and the details of the behavior of SVO instances connected to NCS 2000 devices.

The following lists the implementation of high availability:

- SVO instances use a **high availability** network for database synchronization.

- Admin Planes use a **high availability** network as the primary path and **devices** network as the secondary path. Admin Plane continuously monitors the connectivity between an SVO instance and the associated NCS 2000 device.

Consider there are multiple SVO instances created. Each SVO instance is associated with an NCS 2000 device. Some SVO instances are Active on one server or VM, and other SVO instances are Active on the peer server or VM.

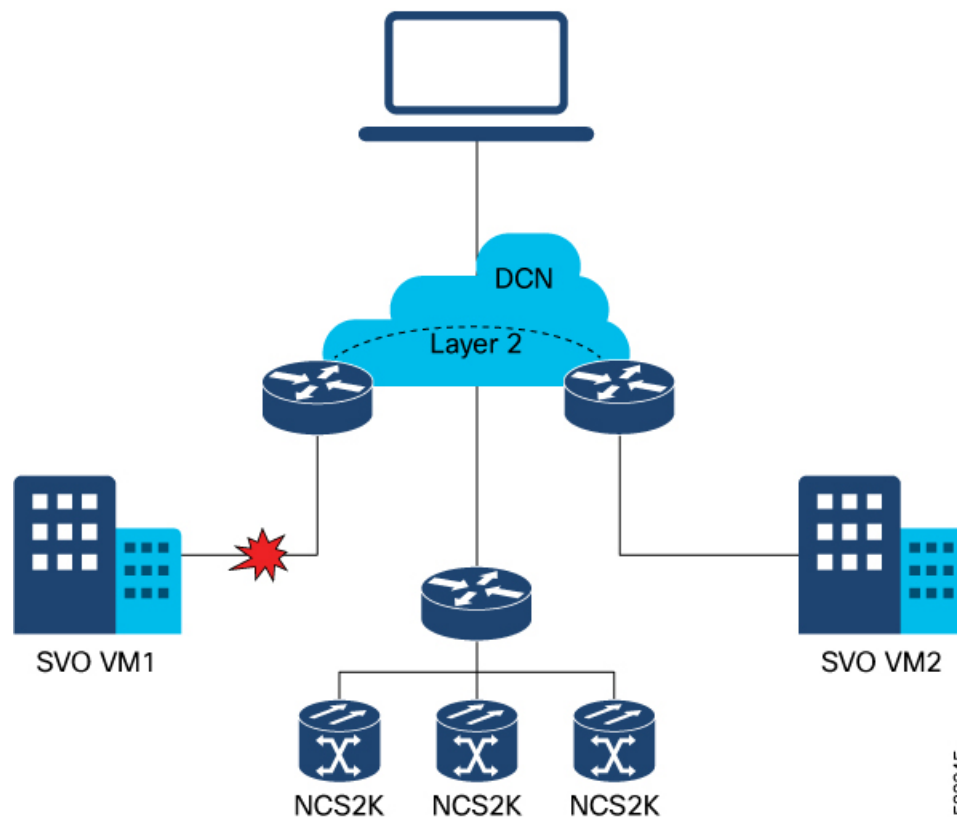
Data Center Restoration

Admin Plane and SVO instances periodically verify the communication channels for high availability.

When network connectivity restores, the SVO recognizes the networks in any order. The SVO can recognize the **devices** network first and then the **high availability** network, or vice versa.

Based on the order of recognition, there can be a transient (less than a minute) delay where an SVO instance is in an Active/Active state. Irrespective of the order, the Admin Planes manage all the situations, assigning the roles.

Figure 5: Data Center Restoration



Sudden Data Center Disconnection

Admin Plane and SVO Instances running on one server are not more able to communicate with the peers.

SVO Instances in Stand-By move to None because they are not able to communicate with the Active.

Admin Plane running in the disconnected data center, detects SVO Instances are not able to communicate with NCS 2000 devices, and aware of the high availability issue, move all the SVO Instances from Active to None.

Admin Plane running in the working data center is not able to communicate with the peer Admin Plane on both primary and secondary paths, it verifies SVO Instances are able to communicate with NCS 2000 devices, then moves SVO Instances from Stand-By to Active.

Use Cases

The following is the list of use cases for the SVO application setup.

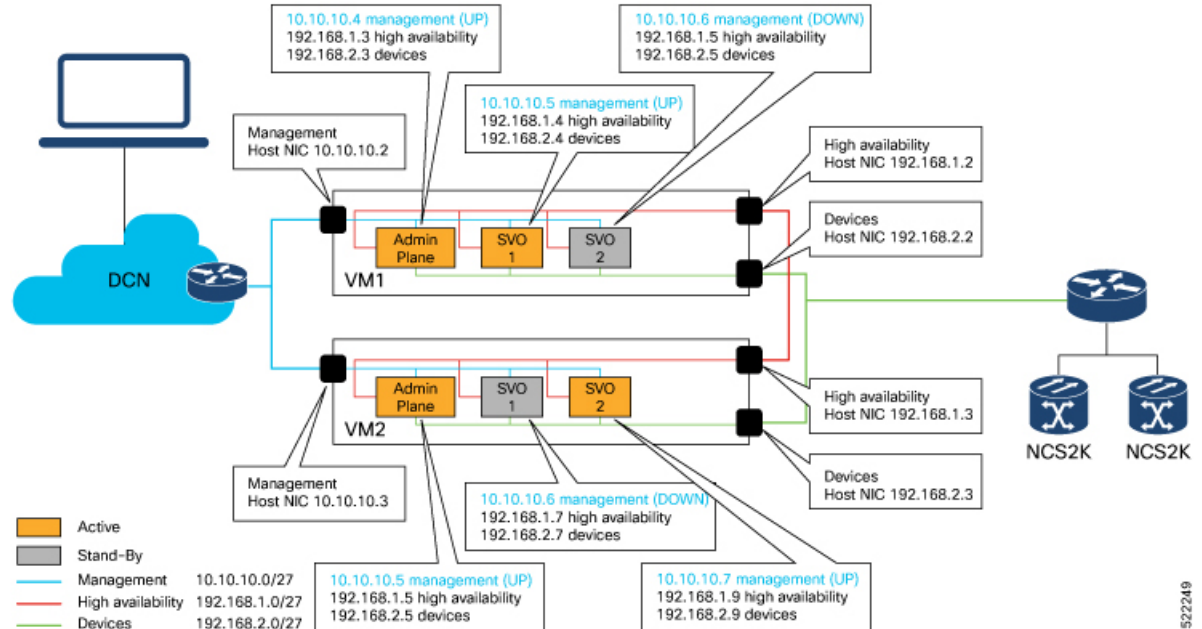
- [Use Case 1 - Colocated Servers, on page 32](#)
- [Use Case 2 - Dislocated Servers \(L2 Interconnection\), on page 33](#)

Use Case 1 - Colocated Servers

The following image displays typical use case of servers or VMs colocated in the same location, it shows an example of IP addresses assignment for host NICs, Admin Planes, and SVO instances.

From Release 12.1, the management, high availability, and devices networks are the same on both the servers or VMs.

Figure 6: Colocated Servers



522249



Note Host NIC IP addresses (text in black) can be configured on the server or VM, before continuing with the SVO installation.

IP addresses in gray are automatically assigned by the Admin Plane during SVO creation.

Use Case 2 - Dislocated Servers (L2 Interconnection)

The following image displays the typical use case of servers or VMs distributed in different locations. It describes the IP addresses assigned for the host NICs, Admin Planes, and SVO instances.

From Release 12.2, the management network is common for both the servers and VMs in different locations. However, the high availability and devices networks are different for both the servers and VMs in different locations.

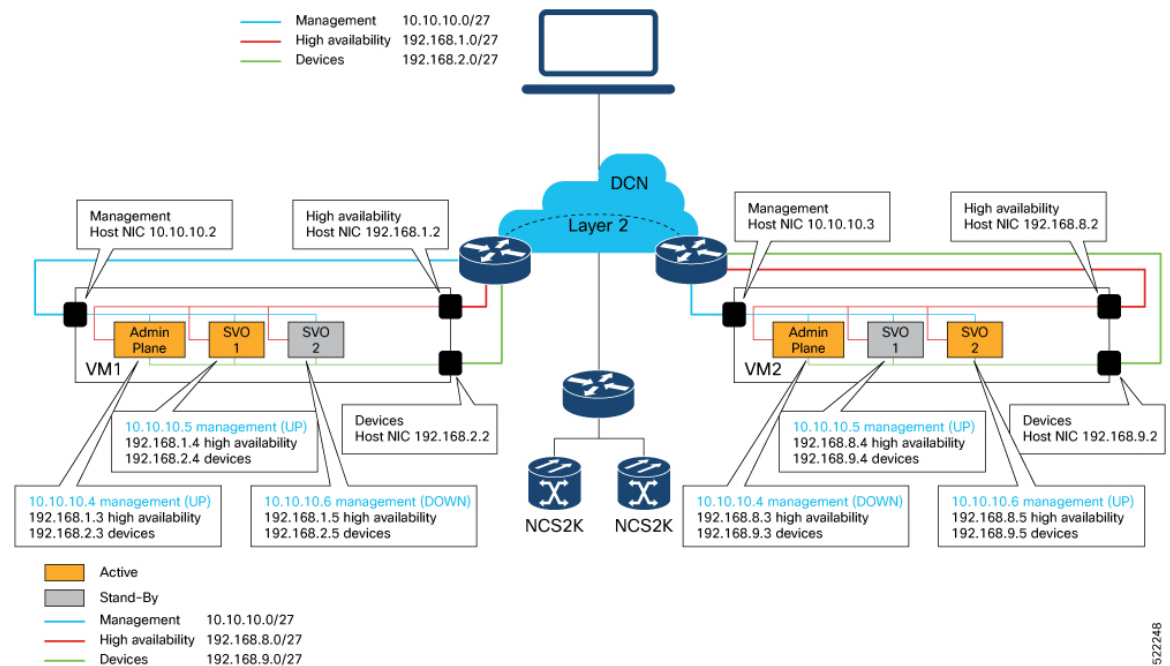


Note To implement the Layer 2 connection for the router interfaces, use the IP addresses of the **management** subnet from the end. The valid Layer 2 connections are L2VPN, L2TPv3, VXLAN, or a similar solution.

The following lists the IP addresses to use for router interfaces on a specific **management** subnet:

- For the 10.10.10.0/27 subnet, use 10.10.10.29 and 10.10.10.30 IP addresses.
- For the 10.10.10.0/24 subnet, use 10.10.10.253 and 10.10.10.254 IP addresses.

Figure 7: Dislocated Servers



522248



Note Host NIC IP addresses (text in black) can be configured on the server or VM before continuing with the SVO installation.

IP addresses in gray are automatically assigned by the Admin Plane during SVO creation.

Use Case 3 - Dislocated Servers (L3 Interconnection)

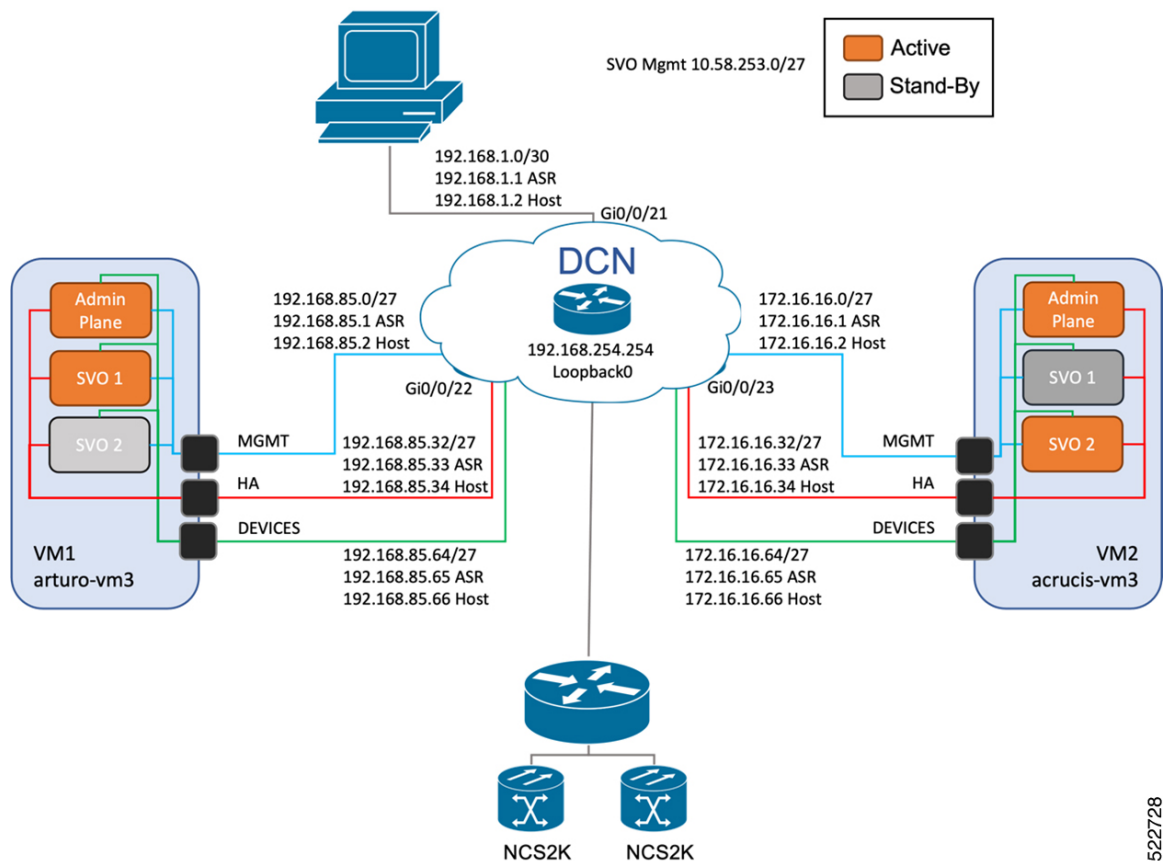
The following images display the typical use case of servers or VMs dislocated in different locations. It shows an example of IP addresses assignment for host NICs, Admin Planes, and SVO instances.

In this use case, the management interconnection at Layer 3 exists between the servers or VMs. The Layer 3 interconnection is the difference between this use case and the previous use cases.

From Release 12.3.1, the management, high availability, and devices networks are different for both the servers and VMs in different locations. However, the SVO applications share the same management subnet between the servers or VMs in different locations.

The following image displays the L3 interconnection with IPv4 addresses assigned to the server NICs. The SVO configured management subnet is **10.58.253.0/27**.

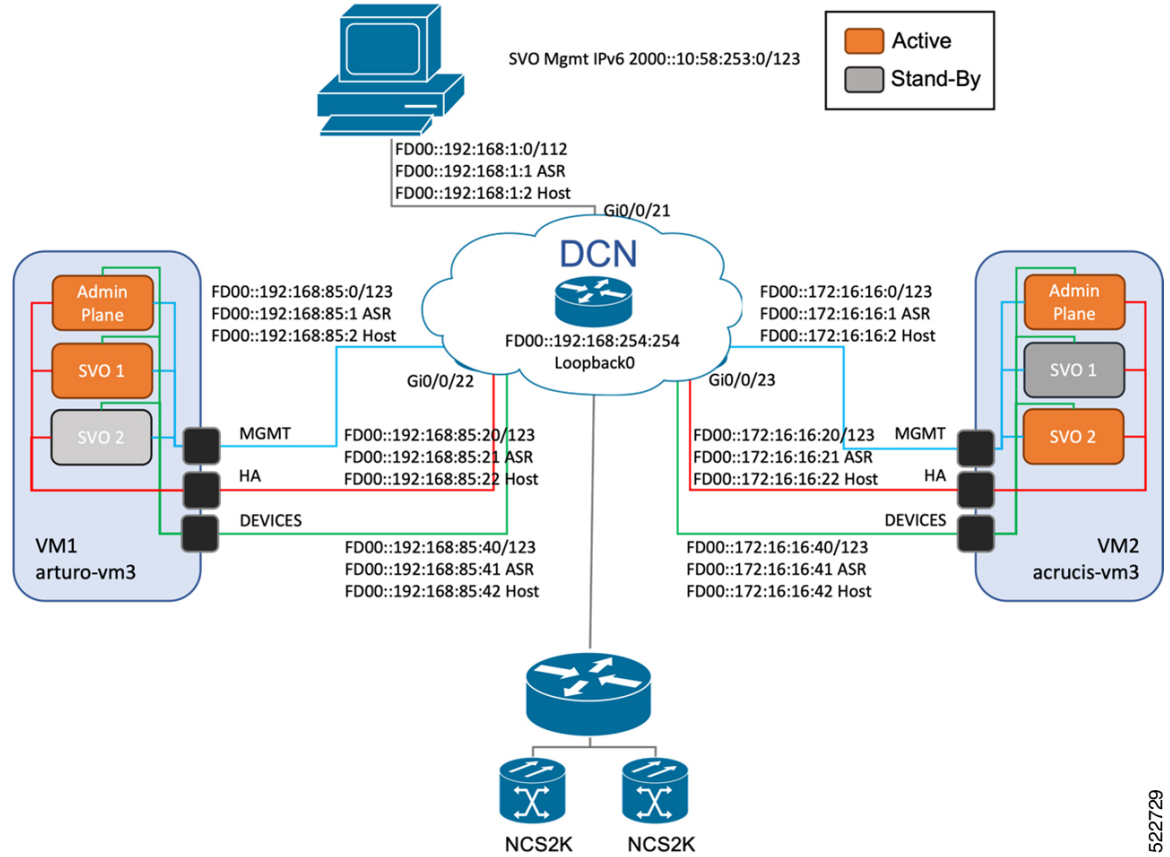
Figure 8: Dislocated Servers (L3 Interconnection) - IPv4 Configuration



522728

The following image displays L3 interconnection with the IPv6 addresses assigned to the server NICs. The SVO configured management subnet is **2000::10:58:253:0/123**.

Figure 9: Dislocated Servers (L3 Interconnection) - IPv6 Configuration



522729

