



Installing Cisco VTS on VMware

- [Installing Cisco VTS on a VMware Environment, on page 1](#)
- [Installing VTSR, on page 4](#)
- [Installing VTF on vCenter, on page 11](#)
- [Verifying VTS Installation, on page 13](#)
- [Changing Password for Cisco VTS from VTS GUI, on page 15](#)
- [Troubleshooting VTS Login Issues, on page 17](#)

Installing Cisco VTS on a VMware Environment

Installing Cisco VTS on a VMware environment involves:

- [Installing VTC VM on ESXi, on page 1](#)
- Installing vCenter Plugin

Installing VTC VM on ESXi

To install VTC VM on an ESXi host:

-
- Step 1** Connect to the ESXi host using the VMware vSphere Client.
- Step 2** In the vSphere Client, select **File > Deploy OVF Template**. The Deploy OVF Template wizard appears.
- Step 3** Specify the name and source location, and click Next.
- Note** You may place vtc.ovf and vtc.vmdk in different directories.
- Step 4** Select the appropriate host to spawn the VTC VM.
- Step 5** For VM disk format, use the default disk format settings (that is Thick Provision Lazy Zeroed).
- Step 6** Map VTC network connectivity to appropriate port-groups on vSwitch/DVS.
- vNIC1—Used for VTC network management
 - vNIC2—Used for VTC connectivity to VTF, VTSR
- Step 7** Enter the following properties:

- Hostname—VTS Hostname.
- Management IPv4 Address—Management IP address for VTC. This IP address is used for VTC network management.
- Management IPv4 Gateway—Management Gateway address
- Management IPv4 Netmask—Management IP Netmask
- Management IPv4 Method—DHCP / Static IP configuration for static IP .
- Management IPv6 Address—Management IP address for VTC. This IP address is used for VTC network management.
- Management IPv6 Gateway—Management Gateway address
- Management IPv6 Method—DHCP / Static IP configuration for static IP, or none.
- Management IPv6 Netmask—Management IP Netmask
- Underlay IPv4 Address—Underlay IP address. This is the IP address for internal network.
- Underlay IPv4 Gateway—Underlay Gateway IP
- Underlay IPv4 Method—DHCP / Static IP configuration for static IP.
- Underlay IPv4 Netmask—Underlay IP Netmask.
- Underlay IPv6 Address—Underlay IP address. This is the IP address for internal network.
- Underlay IPv6 Gateway—Underlay Gateway IP
- Underlay IPv6 Method—DHCP / Static IP configuration for static IP.

Note Cisco VTS does not support IPv6 Underlay configuration. You must specify the Underlay IP6 Method value as **None** to avoid errors.

- Underlay IPv6 Netmask—Underlay IP Netmask.
- DNSv4—IP address of the DNS server.
- Domain—The DNS Search domain.
- NTPv4—NTP address. Can be same as gateway IP address.
- vts-adminPassword—Password for the vts-admin user. Password used to access VTC via SSH for vts-admin account.
- AdministrativeUser—The Administrator User. Enter administrative username.
- AdministrativePassword—Password for administrator user.

Note admin/admin is used to log into GUI for 1st time. The password will be changed during first time login into GUI

While creating admin domain with large number of devices per L2GWgroup or L3GWGroup, you must add devices in smaller batches.

Note that locally we have tried batch of 40 devices and the admin domain creation gets completed within 11 minutes.

Installing vCenter Plugin

The vCenter plugin gets installed when you register the VMM from the Cisco VTS GUI.

Step 1 Go to **Administration > Virtual Machine Manager**.

Step 2 Click the Add (+) button.

The Register VMM page is displayed.

Step 3 Enter the VMM Details:

- Name—Name of the VMM.
- Version —Specify the version from the drop-down.
- Mode—Whether the VMM has been registered as Trusted or Untrusted.
- API Endpoint Details. This is optional.
 - API Endpoint Details:
 - API Protocol:IP Address:Port—VMM service endpoint's IPv4/IP6 address and port.
 - Datacenter—The name of the datacenter for which Cisco VTS acts as the controller.
 - Admin User Name—Username of the vCenter VMM.
 - Admin Passphrase —Password of the vCenter VMM.

Step 4 Click **Register**.

After the VMM is registered successfully, the Plugin sections opens up.

Step 5 Enter the following in the Plugin details section:

- IP Address : Port
- Admin User Name
- Admin Passphrase

Note If you had entered the API endpoint details, the Plugin details will get populated automatically.

Notes Regarding VMware vSphere Distributed Switch

The following points need to be taken care of while you create a vDS.



Note

- All the ToRs in the inventory should be part of the vDS.
 - One vDS can represent one or more ToRs.
 - All the hosts that are connected to a particular ToR should be part of the same vDS.
-

For Non-vPC Specific Configuration

If you are not using vPC on the leaves:

- Associate one or more leafs per vDS.
- Attach the hosts data interface to the vDS uplinks.



Note See VMware documentation for the detailed procedure.

For vPC Specific Configuration

If you are using vPC on the leaves:

-
- Step 1** Create one vDS switch for one or more vPC pairs.
- Step 2** Enable enhanced LACP.
See VMware documentation for the detailed procedure.
- Step 3** Create a Link Aggregation Group for each vDS.
See VMware documentation for the detailed procedure.
- Step 4** You may remove the default port group that gets created as it will not be used.
-

Installing VTSR

The VTSR VM acts as the control plane for the VTF. You need to install VTSR only if you plan to have a VTF in your set up.

Installing VTSR involves:

- Generating an ISO file. See [Generating an ISO for VTSR](#), for details.
To generate VTSR day0 config, we need to create the site on VTC GUI first and use the generated site-id in vtsr day0 config file to generate the vtsr day0 iso file.
- Deploying the VTSR on the VMM. See [Deploying VTSR on OpenStack](#) or [Deploying VTSR on VMware, on page 7](#), for details.

Generating an ISO for VTSR

To create an ISO for VTSR:



Note For an HA installation, you need to create two ISOs and deploy them separately.

If you are upgrading from 2.6, you need to generate the VTSR ISO again with Monit details in the `vtsr_template.cfg` file. See also, [Upgrading VTSR](#).

Step 1 Go to `/opt/cisco/package/vts/share`.

Step 2 Make a copy of the new `vtsr_template.cfg` template and edit for your VTSR instance. A sample `vtsr_template.cfg` file is given below:

```
# This is a sample VTSR configuration file
# Copyright (c) 2015 cisco Systems

# Please protect the generated ISO, as it contains authentication data
# in plain text.

# VTS Registration Information:
# VTS_ADDRESS should be the IP for VTS. The value must be either an ip or a mask.
# VTS_ADDRESS is mandatory. If only the V4 version is specified,
# The V4 management interface for the VTSR (NODE1_MGMT_NETWORK_IP_ADDRESS)
# will be used. If the V6 version is specified, the V6 management interface
# for the VTSR (NODE1_MGMT_NETWORK_IPV6_ADDRESS) must be specified and will be used.
VTS_ADDRESS="10.85.88.152"
#VTS_IPV6_ADDRESS="a1::10"
# VTS_REGISTRATION_USERNAME used to login to VTS.
VTS_REGISTRATION_USERNAME="admin"
# VTS_REGISTRATION_PASSWORD is in plaintext.
VTS_REGISTRATION_PASSWORD="Cisco123!"
# VTSR VM Admin user/password
USERNAME="cisco"
PASSWORD="cisco123"

# Mandatory Management-VRF name for VTSR.
VTS_MANAGEMENT_VRF="vtsr-mgmt-vrf"

# VTSR VM Network Configuration for Node 1:
# NETWORK_IP_ADDRESS, NETWORK_IP_NETMASK, and NETWORK_IP_GATEWAY
# are required to complete the setup. Netmask can be in the form of
# "24" or "255.255.255.0"
# The first network interface configured with the VTC VM will be used for
# underlay connectivity; the second will be used for the management network.
# For both the MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory; they are used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# currently not supported, so if both are specified V6 will take priority (and
# requires VTS_IPV6_ADDRESS to be set).
# The *V6* parameters for the mgmt network are optional. Note that if V6 is used for mgmt
# it must be V6 on both nodes. Netmask must be the prefix length for V6.
NODE1_MGMT_NETWORK_IP_ADDRESS="19.1.0.20"
NODE1_MGMT_NETWORK_IP_NETMASK="255.255.255.0"
NODE1_MGMT_NETWORK_IP_GATEWAY="19.1.0.1"
#NODE1_MGMT_NETWORK_IPV6_ADDRESS="a1::20"
#NODE1_MGMT_NETWORK_IPV6_NETMASK="64"
#NODE1_MGMT_NETWORK_IPV6_GATEWAY="a1::1"
NODE1_UNDERLAY_NETWORK_IP_ADDRESS="19.0.128.20"
NODE1_UNDERLAY_NETWORK_IP_NETMASK="255.255.255.0"
NODE1_UNDERLAY_NETWORK_IP_GATEWAY="19.0.128.1"
# AUX network is optional
```

```

#NODE1_AUX_NETWORK_IP_ADDRESS="169.254.20.100"
#NODE1_AUX_NETWORK_IP_NETMASK="255.255.255.0"
#NODE1_AUX_NETWORK_IP_GATEWAY="169.254.20.1"
# XR Hostname
NODE1_XR_HOSTNAME="vtsr01"
# Loopback IP and netmask
NODE1_LOOPBACK_IP_ADDRESS="128.0.0.10"
NODE1_LOOPBACK_IP_NETMASK="255.255.255.255"

# Operational username and password - optional
# These need to be configured to start monit on VTSR

#VTSR_OPER_USERNAME="monit-ro-oper"
# Password needs an encrypted value
# Example : "openssl passwd -1 -salt <salt-string> <password>"
#VTSR_OPER_PASSWORD="$1$cisco$b88M8bkCN2ZpXgEEc2sG9/"

# VTSR monit interval - optional - default is 30 seconds
#VTSR_MONIT_INTERVAL="30"

# VTSR VM Network Configuration for Node 2:
# If there is no HA then the following Node 2 configurations will remain commented and
# will not be used and Node 1 configurations alone will be applied
# For HA , the following Node 2 configurations has to be uncommented
# VTSR VM Network Configuration for Node 2
# NETWORK_IP_ADDRESS, NETWORK_IP_NETMASK, and NETWORK_IP_GATEWAY
# are required to complete the setup. Netmask can be in the form of
# "24" or "255.255.255.0"
# The first network interface configured with the VTC VM will be used for
# underlay connectivity; the second will be used for the management network.
# For both the MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory; they are used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# currently not supported, so if both are specified V6 will take priority (and
# requires VTS_IPV6_ADDRESS to be set).
# The *V6* parameters for the mgmt network are optional. Note that if V6 is used for mgmt
# it must be V6 on both nodes. Netmask must be the prefix length for V6.
#NODE2_MGMT_NETWORK_IP_ADDRESS="19.1.0.21"
#NODE2_MGMT_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_MGMT_NETWORK_IP_GATEWAY="19.1.0.1"
##NODE2_MGMT_NETWORK_IPV6_ADDRESS="a1::21"
##NODE2_MGMT_NETWORK_IPV6_NETMASK="64"
##NODE2_MGMT_NETWORK_IPV6_GATEWAY="a1::1"
#NODE2_UNDERLAY_NETWORK_IP_ADDRESS="19.0.128.21"
#NODE2_UNDERLAY_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_UNDERLAY_NETWORK_IP_GATEWAY="19.0.128.1"
# AUX network is optional
# Although Aux network is optional it should be either present in both nodes
# or not present in both nodes.
# It cannot be present on Node1 and not present on Node2 and vice versa
#NODE2_AUX_NETWORK_IP_ADDRESS="179.254.20.200"
#NODE2_AUX_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_AUX_NETWORK_IP_GATEWAY="179.254.20.1"
# XR Hostname
#NODE2_XR_HOSTNAME="vtsr02"
# Loopback IP and netmask
#NODE2_LOOPBACK_IP_ADDRESS="130.0.0.1"
#NODE2_LOOPBACK_IP_NETMASK="255.255.255.255"

# VTS site uuid
VTS_SITE_UUID="abcdefab-abcd-abcd-abcd-abcdefabcdef"

```

Step 3 Update the following on *vtsr_template.cfg* for your deployment.

Note To deploy VTSR in HA mode, you need to create two ISOs. To create two ISOs, comment out the parameters starting `NODE2_` in the sample file, and provide the appropriate values.

- `VTS_ADDRESS` - VTS IP address
- `NODE1_MGMT_NETWORK_IP_ADDRESS` - VTSR IP address
- `NODE1_MGMT_NETWORK_IP_GATEWAY` - VTSR gateway address
- `NODE1_UNDERLAY_NETWORK_IP_ADDRESS` - This is the place where TOR is connected directly
- `NODE1_UNDERLAY_NETWORK_IP_GATEWAY` - Underlay network IP address and Underlay network IP gateway should be brought where the VTS underlay network is configured.

Note `VTSR_OPER_USERNAME` and `VTSR_OPER_PASSWORD` are mandatory to start Monit on VTSR. `VTSR_MONIT_INTERVAL` is optional. It is 30 seconds, by default. See *Monitoring Cisco VTS* chapter in the *Cisco VTS User Guide* for details about Monit.

Step 4 Run the `build_vts_config_iso.sh` vtsr script: This will generate the ISO file that you need to attach to the VM before booting it.

Note Ensure that you log in as a root user.

For example:

```
admin@dev: #/opt/cisco/package/vts/bin/build_vts_config_iso.sh vtsr
/opt/cisco/package/vts/share/vtsr_template.cfg
Validating input.
validating
Generating ISO File.
Done!
admin@dev:~$ ls -l
-rw-r--r-- 1 admin vts-admin 360448 Jan 4 18:16 vtsr_nodel_cfg.iso
```

Note In case you had entered the parameters for the second ISO, for HA deployment, running the script generates two ISOs.

Deploying VTSR on VMware

Deploying the VTSR.ova is similar to XRNC.

Step 1 Generate an ISO file for the VTSR VM. See [Generating an ISO for VTSR](#) .

Step 2 In the vSphere Client, select **File > Deploy OVF Template**. The Deploy OVF Template wizard appears.

Step 3 Select VTSR.ova from the source location, and click **Next**.The OVF template details are displayed.

Step 4 Click **Next** to specify the destination. Enter the following details:

- Name for the VM
- Folder or datacenter where the VM will reside

Step 5 Click **Next** to select the storage location to store the files for the template. The default values for virtual disk format and VM Storage Policy need not be changed.

- Step 6** Click **Next** to set up the networks. Specify the first network as the Underlay Network and the second network as the Management Network.
- Step 7** Click **Next**. Review the settings selections.
- Step 8** Click **Finish** to start the deployment.
- Step 9** After the deployment is complete, edit the VM settings. Add a CD/DVD Drive selecting Datastore ISO file and point to the `vtsr.iso` file which was generated and uploaded to the host.
- Step 10** Power on the VM.
- Step 11** To ensure VTSR is configured with the proper Day Zero configuration, SSH to VTSR and then run:

```
RP/0/RP0/CPU0:vtsr01-vcenter#bash
[xr-vm_node0_RP0_CPU0:~]$docker ps
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
31f6cbe6a048 vtsr:dev "/usr/bin/supervisord" 3 weeks ago Up 7 days vtsr
```

- Step 12** Run either of the following commands:

- `[xr-vm_node0_RP0_CPU0:~]$docker exec -it vtsr bash`

Or,

- `[xr-vm_node0_RP0_CPU0:~]$docker exec -it 31 bash`

In the second option, 31 is the process ID, which you can get from Step 11.

An output similar to the below example is displayed:

```
connecting to confd_cli
root@host:/opt/cisco/package# confd_cli -u admin -C
Welcome to the ConfD CLI
admin connected from 127.0.0.1 using console on host
host> en
host# show running-config vtsr-?
Possible completions:
vtsr-config vtsr-day0-config
host(config)# vtsr-config ?
Possible completions:
dhcp-relays global-config interfaces ip-routes l2-networks vm-macs vrf vtsr
host(config)# vtsr-config
```

Do not press or Enter key when the VTSR is loading or getting registered with VTC. For vCenter, VTSR may take approximately 30-45 minutes to come up.

Applying VTSR Device Templates Using `vts-cli.sh` Script

The Day Zero configuration (OSPF, loopback0) has to be configured on VTSR using the `vts-cli.sh` script. You can apply the following templates:



Note This procedure is not required in case you have VTF in L2 switch mode.

Run `vts-cli.sh`, after you run `sudo su -`.

In VTC L3HA scenario, cluster installation will configure loop back and ospf/isis configs on VTSRs based on the information provided in the `cluster.conf` file. No need to run these templates again for VTEP mode.

- vtsr-underlay-loopback-template. See [Applying Loopback Template](#)
- vtsr-underlay-ospf-template. See [Applying OSPF Template](#)
- vtsr-underlay-isis-template. See [Applying IS-IS Template](#)

To determine the usage go to /opt/vts/bin and enter ./vts-cli.sh

```
# This is a sample VTSR configuration file
# Copyright (c) 2015 cisco Systems

# Please protect the generated ISO, as it contains authentication data
# in plain text.

# VTS Registration Information:
# VTS_ADDRESS should be the IP for VTS. The value must be either an ip or a mask.
# VTS_ADDRESS is mandatory. If only the V4 version is specified,
# The V4 management interface for the VTSR (NODE1_MGMT_NETWORK_IP_ADDRESS)
# will be used. If the V6 version is specified, the V6 management interface
# for the VTSR (NODE1_MGMT_NETWORK_IPV6_ADDRESS) must be specified and will be used.
VTS_ADDRESS="10.85.88.152"
#VTS_IPV6_ADDRESS="a1::10"
# VTS_REGISTRATION_USERNAME used to login to VTS.
VTS_REGISTRATION_USERNAME="admin"
# VTS_REGISTRATION_PASSWORD is in plaintext.
VTS_REGISTRATION_PASSWORD="Cisco123!"
# VTSR VM Admin user/password
USERNAME="cisco"
PASSWORD="cisco123"

# Mandatory Management-VRF name for VTSR.
VTS_MANAGEMENT_VRF="vtsr-mgmt-vrf"

# VTSR VM Network Configuration for Node 1:
# NETWORK_IP_ADDRESS, NETWORK_IP_NETMASK, and NETWORK_IP_GATEWAY
# are required to complete the setup. Netmask can be in the form of
# "24" or "255.255.255.0"
# The first network interface configured with the VTC VM will be used for
# underlay connectivity; the second will be used for the management network.
# For both the MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory; they are used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# currently not supported, so if both are specified V6 will take priority (and
# requires VTS_IPV6_ADDRESS to be set).
# The *V6* parameters for the mgmt network are optional. Note that if V6 is used for mgmt
# it must be V6 on both nodes. Netmask must be the prefix length for V6.
NODE1_MGMT_NETWORK_IP_ADDRESS="19.1.0.20"
NODE1_MGMT_NETWORK_IP_NETMASK="255.255.255.0"
NODE1_MGMT_NETWORK_IP_GATEWAY="19.1.0.1"
#NODE1_MGMT_NETWORK_IPV6_ADDRESS="a1::20"
#NODE1_MGMT_NETWORK_IPV6_NETMASK="64"
#NODE1_MGMT_NETWORK_IPV6_GATEWAY="a1::1"
NODE1_UNDERLAY_NETWORK_IP_ADDRESS="19.0.128.20"
NODE1_UNDERLAY_NETWORK_IP_NETMASK="255.255.255.0"
NODE1_UNDERLAY_NETWORK_IP_GATEWAY="19.0.128.1"
# AUX network is optional
#NODE1_AUX_NETWORK_IP_ADDRESS="169.254.20.100"
#NODE1_AUX_NETWORK_IP_NETMASK="255.255.255.0"
#NODE1_AUX_NETWORK_IP_GATEWAY="169.254.20.1"
# XR Hostname
NODE1_XR_HOSTNAME="vtsr01"
# Loopback IP and netmask
NODE1_LOOPBACK_IP_ADDRESS="128.0.0.10"
```

```

NODE1_LOOPBACK_IP_NETMASK="255.255.255.255"

# Operational username and password - optional
# These need to be configured to start monit on VTSR

#VTSR_OPER_USERNAME="monit-ro-oper"
# Password needs an encrypted value
# Example : "openssl passwd -1 -salt <salt-string> <password>"
#VTSR_OPER_PASSWORD="$!$cisco$b88M8bkCN2ZpXgEEc2sG9/"

# VTSR monit interval - optional - default is 30 seconds
#VTSR_MONIT_INTERVAL="30"

# VTSR VM Network Configuration for Node 2:
# If there is no HA then the following Node 2 configurations will remain commented and
# will not be used and Node 1 configurations alone will be applied
# For HA , the following Node 2 configurations has to be uncommented
# VTSR VM Network Configuration for Node 2
# NETWORK_IP_ADDRESS, NETWORK_IP_NETMASK, and NETWORK_IP_GATEWAY
# are required to complete the setup. Netmask can be in the form of
# "24" or "255.255.255.0"
# The first network interface configured with the VTC VM will be used for
# underlay connectivity; the second will be used for the management network.
# For both the MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory; they are used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# currently not supported, so if both are specified V6 will take priority (and
# requires VTS_IPV6_ADDRESS to be set).
# The *V6* parameters for the mgmt network are optional. Note that if V6 is used for mgmt
# it must be V6 on both nodes. Netmask must be the prefix length for V6.
#NODE2_MGMT_NETWORK_IP_ADDRESS="19.1.0.21"
#NODE2_MGMT_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_MGMT_NETWORK_IP_GATEWAY="19.1.0.1"
##NODE2_MGMT_NETWORK_IPV6_ADDRESS="a1::21"
##NODE2_MGMT_NETWORK_IPV6_NETMASK="64"
##NODE2_MGMT_NETWORK_IPV6_GATEWAY="a1::1"
#NODE2_UNDERLAY_NETWORK_IP_ADDRESS="19.0.128.21"
#NODE2_UNDERLAY_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_UNDERLAY_NETWORK_IP_GATEWAY="19.0.128.1"
# AUX network is optional
# Although Aux network is optional it should be either present in both nodes
# or not present in both nodes.
# It cannot be present on Node1 and not present on Node2 and vice versa
#NODE2_AUX_NETWORK_IP_ADDRESS="179.254.20.200"
#NODE2_AUX_NETWORK_IP_NETMASK="255.255.255.0"
#NODE2_AUX_NETWORK_IP_GATEWAY="179.254.20.1"
# XR Hostname
#NODE2_XR_HOSTNAME="vtsr02"
# Loopback IP and netmask
#NODE2_LOOPBACK_IP_ADDRESS="130.0.0.1"
#NODE2_LOOPBACK_IP_NETMASK="255.255.255.255"

# VTS site uuid
VTS_SITE_UUID="abcdefab-abcd-abcd-abcd-abcdefabcdef"

```

If there are issues in running the commands, check the `/opt/vts/bin/vts-cli.log` to get more details.

Applying Loopback Template

To apply Loopback template:

Step 1 On VTC (Master VTC in case of an HA setup), go to /opt/vts/bin.

Step 2 Run the following command:

```
admin@VTC1:/opt/vts/bin$ vts-cli.sh -applyTemplate vtsr-underlay-loopback-template
```

This will prompt you to input the parameters. For example:

Note loopback 1 for VTSR device is reserved for VTSR and docker communication. We recommended that you do not use it for VTSR while executing template script.

```
Enter device name: vtsr01
Enter loopback-interface-number: 0
Enter ipaddress: 100.100.100.100
Enter netmask: 255.255.255.255
Template vtsr-underlay-loopback-template successfully applied to device vtsr01
```

In case you have a VTSR HA setup, apply the template on both VTSRs.

The following message is shown if the configuration got applied correctly:

```
Template vtsr-underlay-loopback-template successfully applied to device vtsr01
```

Applying OSPF Template

To apply OSPF template:

Step 1 On VTC (Master VTC in case of an HA setup), go to /opt/vts/bin.

Step 2 Run the following command:

```
admin@VTC1:/opt/vts/bin$ vts-cli.sh -applyTemplate vtsr-underlay-ospf-template
```

This will prompt you to input the parameters. For example:

```
Enter device name: vtsr01
Enter process-name: 100
Enter router-id: 10.10.10.10
Enter area-address: 0.0.0.0
Enter physical-interface: GigabitEthernet0/0/0/0
Enter loopback-interface-number: 0
Enter default-cost: 10
```

In case you have a VTSR HA setup, apply the template on both VTSRs.

The following message is shown if the configuration got applied correctly:

```
Template vtsr-underlay-ospf-template successfully applied to device vtsr01
```

Installing VTF on vCenter

We recommend that you register the VMM via the VTS GUI, before you install VTF to ensure there are no errors later.

Before you install VTF, you must install VTSR and register it to VTS. See [Installing VTSR](#), for details.

Also, verify whether VTSR is in sync with the VTC. If not, use the sync-from operation via VTS-GUI to synchronize the VTS configuration by pulling configuration from the device. See *Synchronizing Configuration* section in the *Cisco VTS User Guide* for more information on this feature.



Note vCenter supports VTF in VTEP mode only.

Before you install VTF, do the following:

- Set additional routes on VTC VM(s)— You need to add routes for all underlay networks into VTC for across-the-ToR underlay communication. For example, if SVI configuration across ToR from VTC is:

```
interface Vlan100
  no shutdown
  no ip redirects
  ip address 33.33.33.1/24
  no ipv6 redirects
  ip router ospf 100 area 0.0.0.0
  ip pim sparse-mode
```

then, below route needs to be added on VTC VM(s):

```
sudo route add -net 33.33.33.0/24 gw 2.2.2.1
```

Where, 2.2.2.1 is the SVI IP address on the local ToR from VTC VM(s).

Step 1 Specify the VTF Mode in the System Settings. Go to **Administration > System Settings** page, select VTEP from the drop down.

Step 2 Go to **Host Inventory** and edit the host on which VTF (VTEP mode) installation needs to be done.

Step 3 On Host Details, fill in all fields.

Ensure that you review the tooltips for important information about the entries.

Step 4 Select the Virtual Switch. You have the following options:

- Not Defined
- DVS
- vtf-vtep

To install VTF, select vtf-vtep

Step 5 Enter the VTF details.

- Underlay VLAN ID
- Underlay bridge/portgroup on DVS—This is the port group towards the fabric to which the VTF underlay interface will be connected. This needs to be created in advance on vCenter.
- Internal Bridge/Portgroup—This is the DvS portgroup towards Virtual Machines and should be also created in advance on vCenter. This portgroup should be setup as trunk, and security policy should allow Promiscuous mode, Mac address changes and Forged transmits (Set to Accept).
- Datastore—This is the datastore where the vmdk of the VTF VM will be stored, specify the datastore on the VTF host that you want to use.

Set up of the underlay ToR and the corresponding port-group on the DVS has to be done manually on vCenter.

- Step 6** Verify the interfaces information.
 - Step 7** Check the **Install VTF on Save** check box, and click Save.
 - Step 8** Check the installation status in the Host Inventory page.
 - Step 9** Check the VTF registration status on **Inventory > Virtual Forwarding Groups** page.
-

Uninstalling VTF in a vCenter Environment

Before you VTF uninstall, go to **Inventory > Virtual Forwarding Groups** to verify that VTF is shown in Virtual Forwarding Groups page.

To uninstall VTF

- Step 1** Go to Host Inventory, and edit the host to change Virtual Switch type from vtf-vtep to not-defined.
 - Step 2** Click **Save**.
 - Step 3** Check the uninstallation status on the Host Inventory page to verify whether Installation status is unchecked and Virtual Switch is not-defined.
 - Step 4** Go to the **Inventory > Virtual Forwarding Groups** page, to verify that it does not show VTF that you uninstalled.
 - Step 5** Go to vCenter using vSphere Web Client.
 - Step 6** Go to Hosts and Clusters, click the VTF VM that got uninstalled from VTS GUI.
 - Step 7** Power off the VTF VM.
 - Step 8** Delete the VTF from disk
-

Verifying VTS Installation

The following sections provide information about how to verify the VTS installation:

- [Verifying VTC VM Installation](#)
- [Verifying VTSR Installation](#)
- [Verifying VTF Installation](#)

Verifying VTC VM Installation

To verify VTC VM installation:

- Step 1** Log in to the VTC VM just created using the VTC VM console.
 - If you have installed the VTC VM in a VMware environment, use the VM console.
 - If you have installed the VTC VM in an RedHat KVM based-OpenStack environment, - telnet 0 <console-port> (The console port is telnet port in the VTC.xml file.)

- Step 2** Ping the management gateway.
In case ping fails, verify the VM networking to the management network.
- Step 3** For the VTC VM CLI, ping the underlay gateway.
In case the ping fails, verify VM networking to the underlay network.
- Note** Underlay network gateway is the switched virtual interface (SVI) created for VTSR and VTF on the leaf where the controller is connected.
- Step 4** Verify whether the VTS UI is reachable, by typing in the VTS management IP in the browser.
-

Verifying VTSR Installation

To verify VTSR installation:

- Step 1** Log in to the VTSR.
- If you have installed the VTC VM in a VMware environment, use the VM console.
 - If you have installed the VTC VM in an RedHat KVM based-OpenStack environment, use virt-manager or VNC based console method to login into the VM. See [Installing VTC VM - Manual Configuration using VNC](#)
- Step 2** Ping the underlay gateway IP address.
In case ping fails, verify underlay networking.
- Step 3** Ping the VTC VM.
- ```
On VTSR262 based on XR 651.we have Mgmt under new Mgmt VRF.So Ping of Mgmt should be within VRF :
vrf vtsr-mgmt-vrf
address-family ipv4 unicast
!
address-family ipv6 unicast
!
RP/0/RP0/CPU0:vtsr01#ping 50.1.1.251 vrf vtsr-mgmt-vrf
Thu Aug 30 13:51:25.873 UTC
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 50.1.1.251, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```
- In case ping fails, verify underlay networking.
- Note** You should be able to ping the gateway IP address for both management and underlay networks, as VTSR registers to the VTC using the management IP address.
- VMM\_ID : It is the VMM-ID local to the site
- Site\_ID : As part of Multi-site support, we have to define the Site-ID which VTSR instance is managing VFG within that site. For more information, see *Multi-site support* section of the *Cisco VTS 2.6.3 User Guide* .
- Step 4** Run **virsh list** to make sure the nested VM is running.
- Step 5** Verify whether the Virtual Forwarding Group (VFG) group is created on VTS GUI, and VTSR is part of the VFG group.
-

## Verifying VTF Installation

To verify VTF installation:

- 
- Step 1** Log in to the VTF VM / vhost.
- If you have installed the VTC VM in a VMware environment, use the VM console.
  - If you have installed the VTC VM in an RedHat KVM based-OpenStack environment, use virt-manager or VNC based console method to login into the VM. See [Installing VTC VM - Manual Configuration using VNC](#)
  - For vhost mode, connect to the compute and checkvpfa/vpp services or RPM packages.
- If registration is successful, you should see the newly registered VTF IP (underlay IP) under VTF list (**Inventory > Virtual Forwarding Groups**).
- Step 2** Ping the underlay gateway IP address.
- In case ping fails, verify underlay networking.
- Step 3** Ping the VTC VM underlay IP address.
- In case ping fails, verify underlay networking.
- In case VTC and VTF are on different subnets, then verify whether routes to VTS are configured on the compute.
- Step 4** Verify whether the VTF CLI is available . To do this, run:
- ```
'sudo vppctl
```
- If the o/p command fails, run the following command to identify whether vpfa service is up and running:
- ```
sudo service vpfa status
```
- If there are errors, try restarting the service.
- ```
sudo service vpfa restart
```
- Step 5** Verify whether the VTF is part of the VFG on VTS GUI (**Inventory > Virtual Forwarding Groups**).
-

Changing Password for Cisco VTS from VTS GUI

The GUI password change will trigger the updating of password on all host agents which are running on the Physical computes. And if there are VTFs in your setup, then the VTSR and VTF passwords will also get updated.

**Important**

- Traffic disruption will happen only if you have VTFs installed (Virtual deployment) and it happens because of the vpfa process restart.
In case of a Physical deployment there will not be any traffic disruption.
- For Baremetal ports there is no impact.
- The password change from the GUI will change only the host agent password. Not the Linux password. So, we cannot use the command 'passwd'
- If you are changing the Linux password of a Physical or Virtual host then you should also update the VTC host inventory with correct password. Changing the Linux password will not impact any traffic.
- If you setup two nodes with different GUI Password and try setting up L2 HA, it will fail. You need to make sure that both the nodes have same password before setting up L2 HA.
- If you already have L2 HA, you can change the GUI Password from the GUI by logging in with VIP IP. This will change the GUI Password on both Master and Slave nodes. Changing the GUI password on master and slave nodes separately is not supported.

Step 1 Log in to VTS GUI and click on settings icon on the top-right corner and click **Change Passphrase**.

Step 2 Enter the current password, new password, then click **Change Passphrase**.

Step 3 Click **OK** in the Confirm Change Passphrase popup, to confirm.

Note The message in the Confirm Change Passphrase window is just a generic message. See important notes above for details about possible traffic disruption.

Changing Password for Cisco VTS Linux VM

You can use the Linux command 'passwd' to change the VTC VM password. After changing the password, you should use the new password for the subsequent SSH session to the VTC VM.

For example:

```
root@vts:~# passwd admin
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
```

For an HA installation you must change the password on both Master and Slave with the command 'passwd'. In an HA setup, you can change the passwords without uninstalling HA. This password change will not impact the HA setup as HA uses the GUI Password, which needs to be same on both Master and Slave nodes.



Note You can set different admin password on both the nodes, but make sure you remember and use the correct password to log in to the respective nodes.

Troubleshooting VTS Login Issues

When you are unable to log in into VTS via CLI, check for the following points:

As a part of security compliance, it is recommended that you should remember your admin password. If you forget the admin password, then you should use another user account and change the password as a root user.

VTS installs an SSH Guard application on the system by default.

When you enter a wrong password consecutively for 4 times, the SSH Guard application temporarily blocks the IP that has authentication failures for approximately 7 to 10.5 minutes. If you are trying to enter the wrong password again, the blocking time doubles for each set of 4 failed logins.

For example, a host with IP of 192.0.2.1 is blocked when you enter a wrong password for 4 times within a period of 20 minutes interval. The IP is unblocked from the first set of login failures within 7 to 10.5 minutes due to periodic interval checks. When you try to enter the wrong password again for 4 consecutive times, the blocking time doubles upto 14 minutes.

