



Installing Cisco VTS

If your Cisco NFVI package includes Cisco Virtual Topology System, the following topics tell you how to install Cisco VTS for use with Cisco NFVI. The Cisco VTS installation procedures are customized for Cisco NFVI from the standard Cisco VTS 2.6.2 installation procedures located on the [Cisco VTS product site](#). You must install Cisco VTS before you install Cisco VIM.

- [Overview to Cisco VTS Installation in Cisco NFVI, on page 1](#)
- [System Requirements for VTC VM, on page 6](#)
- [System Requirements for VTSR VM, on page 7](#)
- [Supported Virtual Machine Managers, on page 7](#)
- [Supported Platforms, on page 7](#)
- [Installing Cisco VTS in Cisco NFVI Environment, on page 9](#)
- [Installing the VTSR VMs, on page 13](#)
- [Verifying Cisco VTS Installation in Cisco NFVI, on page 16](#)
- [Configuring Cisco VTS and VTSR After Installation, on page 18](#)
- [Installing VTS in an HA Configuration, on page 19](#)
- [Sample Cisco VTS Configurations for Cisco NFVI, on page 23](#)


Overview to Cisco VTS Installation in Cisco NFVI

The Cisco Virtual Topology System (VTS) is an overlay management and provisioning system for data center networks. It automates data center overlay fabric provisioning for both physical and virtual workloads. It provides a policy-based approach for overlay provisioning, and can be used for multitenant data centers for cloud services, including Cisco NFVI.

To install Cisco VTS with Cisco NFVI, you must manually install the Cisco VTS Virtual Topology Controller (VTC) and its VTSR VMs before you start the Cisco VIM installation. The VTC and VTSR VMs must be run on an independent pair of servers, that is, not on a Cisco NFVI control, compute, storage, or management node. You set up the networking on those servers as described and outlined in the installation procedures. When you run the Cisco VIM installer, you have to provide the VTC VIP and appropriate VTS credentials.

The following figure shows how Cisco VTS Virtual Topology Controller (VTC) and the VTSR virtual machines (VMs) connect to the Cisco NFVI networks.

Figure 1: Cisco VTS Connectivity to Cisco NFVI



The following table maps Cisco VTS network names to Cisco VIM network names.

Table 1: Cisco VTS to Cisco VIM Network Name Mapping

Cisco VTS VM	Cisco VTS Network Name	Cisco VIM Network Name
VTC	Management Network	API (a)
VTC	Underlay Network	Management or Provision (mx)
VTSR	Management Network	Management or Provision (mx)
VTSR	Underlay Network	Tenant (t)

The following table describes the required IP address allocations for VTS components.

Table 2: Cisco VTS IP Address Allocations

Cisco VIM Network	Required Cisco VTS IP Addresses	Description
API (a)	3 total (1 VIP + 1 IP per VTC VM)	Set up in the VTC config.iso and cluster.conf
Management or Provisioning (mx)	<ul style="list-style-type: none"> 5 total—Three for VTC (one VTC VIP called as VTS_NCS_IP in setup_data and one IP per VTC VM) Two for VTSR: one IP per VTSR VM. 	Set up in VTSR config.iso. Note: VTS component IP addresses cannot overlap with the pool ranges configured in the Cisco VIM setup_data.yaml.

Cisco VIM Network	Required Cisco VTS IP Addresses	Description
Tenant (t)	2 total—(one IP address VTSR VM.	Set up in VTSR config.iso Note: The VTS component IPs cannot overlap with pool ranges that are configured in the Cisco VIM setup_data.yaml.

The following is the VTS IP distribution and setup mechanism.

VIM API network

- VTC1—api (a) network IP1 (associated through the VTC1 config ISO)
- VTC2—api (a) network IP2 (associated through the VTC2 config ISO)
- VTC VIP—api (a) network IP3 (associated through the HA step cluster.conf)

VIM Management/Provisioning network

- VTC1—management/provisioning (mx) network IP1 (associated through the VTC1 config ISO)
- VTC2—management/provisioning (mx) network IP2 (associated through the VTC2 config ISO)
- VTC VIP—management/provisioning (mx) network IP3 (associated through the HA step cluster.conf)
- VTSR 1—management/provisioning (mx) network IP4 (associated through the VTSR-1 config ISO)
- VTSR 2—management/provisioning (mx) network IP5 (associated through the VTSR-2 config ISO)

VIM Tenant network:

- VTSR 1—tenant (t) network IP1 (associated through the VTSR-1 config ISO)
- VTSR 2—tenant (t) network IP2 (associated through the VTSR-2 config ISO)

Cisco VTS Usernames and Passwords in Cisco NFVI

The following table lists the Cisco VTS usernames and passwords that are deployed after you install Cisco VTS in Cisco NFVI.

Table 3: Cisco VTS Usernames and Passwords in Cisco NFVI

Configuration Location	Value Requirements	Description/Comments
CVIM: openstack-configs/setup_data.yaml VTS_PARAMETERS: VTS_USERNAME VTS_PASSWORD VTS_SITE_UUID The following parameters are optional, only required if VTS_DAY0 is enabled. VTC_SSH_PASSWORD VTC_SSH_USERNAME VTS_SITE_UUID Optional: MANAGED	VTS_USERNAME must be admin. VTS_PASSWORD must match VTC UI login password for the admin user. Password must have a minimum of 8 characters and at least one uppercase letter, one digit, and one special character. VTS_SITE_UUID is unique UUID of VTS SITE controlled by Cisco VIM. The VTS_SITE_UUID must be in a generic UUID format (Unique Pod UUID to indicate which pod the VTS is controlling) The VTC_SSH_PASSWORD and VTC_SSH_USERNAME are ssh credentials to login to VTC VMs. MANAGED is either True or False. By default, it is false. If it is True, VTS deployment mode is managed.	Used by VTF to register with the VTC / VTSR.
VTC ISO config.txt : vts-adminPassword AdministrativeUser AdministrativePassword	Must match the Cisco VIM setup_data.yaml VTC_SSH_PASSWORD parameter. AdministrativeUser must match with setup_data.yml VTC_SSH_USERNAME parameter AdministrativePassword matches with VTC_SSH_PASSWORD parameter.	Configures VTC admin user's initial password. SSH username/password for VTC VM.
VTSR ISO: USERNAME PASSWORD		VTSR VM SSH username/password The VTSR adds this in VTS Inventory > Authorization Group > vtsgroup Device User Name associated with VTC admin user

Modes of TOR Configuration with VTS

Cisco VTS supports two modes of TOR configuration:

- **Unmanaged TOR:** It is the default mode of operation for VTS with Cisco VIM. VTS network inventory is added as “Unmanaged” device instead of actual TOR switches. BGP EVPN ingress replication mode mechanism is used for admin domain, but the port configuration does not push configuration to the TOR switches.
- **Managed TOR:** VTS network inventory is added with actual TOR switches. Control and compute nodes information are added with their corresponding interfaces connected with TOR in the VTS host inventory. BGP EVPN multicast replication mode is used for admin domain, while the port configuration enables multicast Internet Group Management Protocol (IGMP) snooping and PIM configuration for Tenant VLAN on actual TOR switches.



Note As the storage nodes do not have VTF, the switch ports hanging off the storage nodes are configured statically.

To maintain consistency, add the `tor_info` to the storage nodes in the `setup_data` of the pod. .

Listed below is the snippet of the Multicast configuration push to Cisco Nexus 9000, when port is configured with Tenant VLAN ID 111.

```
interface Vlan111
no shutdown
no ip redirects
ip address 22.22.22.200/24
no ipv6 redirects
ip router ospf 100 area 0.0.0.0
ip pim sparse-mode
ip igmp version 3
ip igmp static-oif 239.0.0.1
hsrp 22
ip 22.22.22.1
vlan configuration 111
ip igmp snooping static-group 239.0.0.1 interface port-channel12
ip igmp snooping static-group 239.0.0.1 interface port-channel13
ip igmp snooping static-group 239.0.0.1 interface port-channel14
```



Note Due to limitation of VTS, Tenant VLAN ID needs to be selected as lowest number in the TOR interface. If not, Multicast configuration will be pushed incorrectly.

The following table lists the configurations required to enable the functionality of TORs “managed” through VTS.

Table 4: Cisco VTS Parameters for TORs managed through VTS

Configuration Location	Value Requirements	Description
CVIMmercury: openstack-configs/setup_data.yaml VTS_PARAMETERS: MANAGED:	MANAGED: Set to True or False. By default, it is False.	MANAGED: Must be configured as True, when VTS deployment mode is managed. It is a day-0 configuration, and cannot be enabled as a reconfigure option.

Configuration Location	Value Requirements	Description
TORSWITCHINFO: CONFIGURE_TORS	CONFIGURE_TORS: False	CONFIGURE_TORS value has to be False to indicate that CVIM is not configuring the TORs; this is a way for VTC to know what switches to access and manage
SWITCHDETAILS:	Hostname, ssh_ip, username, and password of the switches for VTC to manage {switch_a_hostname: ethx/y, switch_b_hostname: ethx/y}	Need minimum switch details to access it.
SERVICES: <SERVER_NAME>: tor_info:		For each server, list the tor_info associated to the server, so that VTC can manage the switch ports. Note that the storage nodes do not have VTF and hence switch ports hanging off the storage nodes are configured statically. To maintain consistency, add the tor_info to the storage nodes in the setup_data of the pod.

From an architecture point of view, the following are configured automatically in VTC Node when Managed TOR mode is selected in setup_data.yaml:

- VTS System Settings and Route reflector are configured in VTC.
- Openstack Virtual Machine Manager is configured.
- Global VNI POOL is configured.
- Multicast pools are created to allocate multicast IP address for Tenant VLAN ID.
- Authentication Group is created for device.
- TOR switches are configured under Network Inventory.
- Admin domain is created with BGP EVPN multicast replication mode for L2 and L3 Gateway.
- TOR switches and VTSR are added to L2 and L3 Gateway in admin domain.
- Controller and Compute Node are added under host inventory with corresponding TOR interfaces.
- All VTFS are registered with VTSRs and appear under Virtual Forwarding Groups.

System Requirements for VTC VM

The following table provides information about the minimum system requirements for the VTC virtual machine:

Requirement	Details
Disk space	48 GB

Requirement	Details
CPU	8
Memory	32 GB
Computing host	Certified with Cisco UCS B-series, Cisco UCS C-series Rack Servers

System Requirements for VTSR VM

The following table gives details about the minimum system requirements for the VTSR virtual machine:



Note The VTSR VM serves two purposes. It is required to enable VTS High Availability. It also acts as the control plane for the VTF. You need to install VTSR only if you consider enabling High Availability or if you plan to have a VTF in your set up.

Requirement	Details
Disk Space	Primary disk must be 77 GB.
CPUs	14
Memory	48 GB RAM
Computing Host	Certified with Cisco UCS B-series, Cisco UCS C-series Rack Servers

Supported Virtual Machine Managers

You can install Cisco VTS on the following supported versions of Virtual Machine manager (VMM):

Table 5: Openstack Versions

	OpenStack Liberty	OpenStack Newton
On RHEL	12.0.0; 12.0.1; 12.0.2; 12.0.3; 12.0.4; 12.0.5; 12.0.6	14.0.3 On CentOS
On CentOS	12.0.0; 12.0.1; 12.0.2	N/A

Supported Platforms

The following tables provide information about the platforms that Cisco VTS support, and their roles.



Note VTS supports VXLAN overlays using the BGP EVPN control plane.

Role	Platform Supported
Top-of-rack (ToR) leaf switch	<ul style="list-style-type: none"> • Cisco Nexus 9300TX and 9300PX platform switches • Cisco Nexus 9332PQ and 93128TX switches • Cisco Nexus 9200 platform switches • Cisco Nexus 9500 platform switches
Data center spine	<ul style="list-style-type: none"> • Cisco Nexus 9300TX and 9300PX platform switches • Cisco Nexus 9500 platform switches • Cisco Nexus 9200 platform switches
Border leaf	<ul style="list-style-type: none"> • Cisco Nexus 9300TX and 9300PX platform switches • Cisco Nexus 9500 platform switches • Cisco Nexus 9200 platform switches
Data center interconnect (DCI)	<ul style="list-style-type: none"> • Cisco ASR 9000 Series Aggregation Services routers • Cisco Nexus 9300 platform switches
Virtual machine manager (VMM)	OpenStack Newton on RHEL versions
Hypervisor	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 7.3 with KVM • Red Hat Enterprise Linux 7.5 • CentOS
Virtual forwarders	Cisco Virtual Topology Forwarder (VTF)

The following table lists the software images supported for the different devices.

Table 6: Software Images Supported

Cisco Nexus 93xx	NX OS Release 7.0.3.I7.2 or 9.2(1)
Cisco Nexus 95xx	NX OS Release 7.0.3.I7.2 or 9.2(1)
Cisco ASR 9000	Cisco IOS XR Software Release 6.5.1.

The following table lists the VPC modes supported for the different devices.

Note If Cisco Nexus 9000 series ToR is not configured with vPC related configuration, including peer-link, also known as a multichassis etherChannel trunk (MCT), you must not configure feature vpc on the ToR. This may bring loopback interface used for NVE to admin down state.

Table 7: VPC Modes Supported

Cisco Nexus 93xx	Server VPC
Cisco Nexus 95xx	Server VPC

Installing Cisco VTS in Cisco NFVI Environment

Installing Cisco VTS within Cisco NFVI involves installing the Virtual Topology Controller (VTC) VM. You can install the VTC VM using either the automatic or manual configuration options.

- To install the VTC VM using an ISO file (auto configuration), see [Installing VTC VM - Automatic Configuration Using ISO File, on page 9](#).
- To install the VTC VM using the virt-manager application (manual configuration), see [Installing VTC VM - Manual Configuration Using Virt-Manager, on page 10](#).
- To install the VTC VM using VNC (manual configuration), see [Installing VTC VM - Manual Configuration using VNC, on page 12](#)

Installing VTC VM - Automatic Configuration Using ISO File

To install a VTC VM and enable configuration using an ISO file, create a text file with the VM settings, wrap the text file in an ISO file, and then attach the ISO file to the VM CD drive.

-
- Step 1** Connect to the controller node via SSH, and copy the vtc.qcow2 file to /var/lib/libvirt/images/ folder.
- Step 2** Copy the vtc.sample.xml file to your controller. The [Installing Cisco VTS in Cisco NFVI Environment, on page 9](#) topic provides the file contents.
- Step 3** Create a **config.txt** file containing the following parameters:

```

Hostname=vtc
ManagementIPv4Method=Static
ManagementIPv4Address= <VM's a-net IP address in a.b.c.d form>
ManagementIPv4Netmask= <a-net IP mask in a.b.c.d form>
ManagementIPv4Gateway= <a-net gateway IP address in a.b.c.d form>
UnderlayIPv4Method=Static
UnderlayIPv4Address= <VM's mx-net IP address in a.b.c.d form>
UnderlayIPv4Netmask=<mx-net IP mask in a.b.c.d form>
DNSv4=<DNS server--ie. setup_data.yaml::NETWORKING['domain_name_servers'][0]>
Domain=<domain name--ie. setup_data.yaml::NETWORKING['domain_name']>
NTP=<NTP server--ie. setup_data.yaml::NETWORKING['ntp_servers'][0]>
vtc-adminPassword=<password for user 'admin'--setup_data.yaml::VTS_PARAMETERS['VTC_SSH_PASSWORD']>
AdministrativeUser=<VM ssh login user--can be setup_data.yaml::VTS_PARAMETERS['VTC_SSH_USERNAME']>
AdministrativePassword=<VM ssh login user--can be setup_data.yaml::VTS_PARAMETERS['VTC_SSH_PASSWORD']>
ManagementIPv6Method: Unused by NFVI

```

UnderlayIPv6Method: Unused by NFVI

Note *config.txt* file must have a blank line at the end.

Note Before entering the VTS_PASSWORD, review [Cisco VTS Usernames and Passwords in Cisco NFVI, on page 3](#).

Parameter descriptions:

- Hostname—The VM hostname.
- ManagementPv4Method—Whether to use DHCP or static addressing for the Cisco NFVI API network (a-net) interface (eth0).
- ManagementIPv4Address—The api (a) network IPv4 address of the VM (required only for static addressing).
- ManagementIPv4Netmask—The a network IPv4 net mask of the VM (required only for static addressing).
- ManagementIPv4Gateway—The a network API IPv4 gateway of the VM (required only for static addressing).
- UnderlayIPv4Method—Whether to use DHCP or static addressing for the Cisco NFVI management/provisioning (mx) network interface (eth1).
- UnderlayIPv4Address—The mx network IPv4 address of the VM (required only for static addressing).
- UnderlayIPv4Netmask—The mx network IPv4 net mask of the VM (required only for static addressing).
- DNSv4—DNS IPv4 address (required only for static addressing).
- Domain—DNS search domain (required only for static addressing).
- NTPv4—NTP IPv4 address or FQDN (required only for static addressing).
- vts-admin Password—Password for the vts-admin user. This should match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']` or subsequently changed through the VTC UI to match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']`
- Administrative User—New administrative user for login using SSH.
- Administrative Password—Sudo password for the administrative user.

Step 4 Use `mkisofs` to create an ISO file, for example:

```
mkisofs -o config.iso config.txt
```

Step 5 Create the VTC VM using following command:

```
virsh create vtc.sample.xml
```

Installing VTC VM - Manual Configuration Using Virt-Manager

To install the VTC VM configuring it manually using the virt-manager application:

Step 1 Connect to the controller node via SSH, and copy the `vtc.qcow2` file to `/var/lib/libvirt/images/` folder.

Step 2 Copy the Cisco NFVI `vtc.sample.xml` file to your controller. Modify it as per your setup. See [Sample Cisco VTS Configurations for Cisco NFVI, on page 23](#) for examples.

Step 3 Create the VTC VM using following command:

```
virsh create vtc.sample.xml
```

Step 4 Run the command:

```
virsh list --all
```

It should display:

```
Id      Name      State
-----
2 VTC running
```

Step 5 Start virt-manager. Run:

```
virt-manager
```

Step 6 After the virt-manager window opens, click the VTC VM to open up the VTC VM console.

The console displays an installation wizard that takes you through the initial VTC VM configuration.

Step 7 Enter the following:

Note For items that take multiple values, such as DNS and NTP, each value must be separated by a space.

- VTS Hostname
- DHCP / Static IP configuration for static IP
- Management IP address for VTC—This is the Cisco NFVI api (a) network IP address.
- Management IP Netmask (api network)
- Management Gateway address (api network)
- DNS Address—One of the DNS servers in `setup_data.yaml::NETWORKING['domain_name_servers']`
- DNS Search domain—`setup_data.yaml::NETWORKING['domain_name']`
- Underlay IP address—This is the IP address for Cisco NFVI management/provisioning (mx) network.
- Underlay IP Netmask (mx network)
- NTP address—One of the `setup_data.yaml::NETWORKING['ntp_servers']` addresses
- Password change for user `vtc-admin`—Enter the default user `vtc-admin` password. The `vtc-admin` user is used for password recovery and to revisit a configuration screen if you make a mistake or need to change the information. If you log in to the VTC VM using `vtc-admin` username and password again, you get the same dialog to go through the VTC VM setup again. The password must match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']` or subsequently changed through the VTC UI to match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']`
Before entering the `VTS_PASSWORD`, reviewing [Cisco VTS Usernames and Passwords in Cisco NFVI, on page 3](#) is recommended.
- Administrator User—Enter administrative username and password. This username and password are used to login to the VM via SSH.

- Password for administrator user

VTC VM reboots at this time. Wait for two minutes for the VTC VM to be up. You can ping the IP address given for VTC VM in the setup process to verify whether the VTC VM is up.

Step 8 SSH into VTC VM using the IP address, administrative username/password given in the setup process (not vts-admin user).

Installing VTC VM - Manual Configuration using VNC

If the server where you install VTC is in a remote location with network latency or low bandwidth, you can use VNC to access the VTC VM and manually configure it using the CTC VM graphic console. To do this:

Step 1 Connect to the controller node via SSH, and copy the vtc.qcow2 file to /var/lib/libvirt/images/ folder.

Step 2 Copy the vtc.sample.xml file to your controller. Modify it as per your setup. The sample VTC XML file output is provided in [Sample Cisco VTS Configurations for Cisco NFVI, on page 23](#).

Step 3 Replace the following sections of the vtc.sample.xml file:

```
<graphics type='spice' port='5900' autoport='yes' listen='127.0.0.1'>
  <listen type='address' address='127.0.0.1' />
</graphics>
```

with the following:

```
<graphics type='vnc' port='5900' autoport='yes' listen='0.0.0.0'>
  <listen type='address' address='0.0.0.0' />
</graphics>
```

Note Setting the listen address to 0.0.0.0 allows external clients to connect to the VNC port (5900). You have to make sure that iptables configuration (if any) allows inbound TCP port 5900 connections.

Step 4 Create the VTC VM using following command:

```
virsh create vtc.sample.xml
```

You should now be able to use a VNC client to connect to the VTC VM graphic console and continue the setup.

Step 5 Enter the following:

Note For items that take multiple values, such as DNS and NTP, use a space to separate each value.

- VTS Hostname
- DHCP/Static IP configuration for static IP
- Management IP address for VTC—This is the Cisco NFVI api (a) network IP address.
- Management IP Netmask (api network)
- Management Gateway address (api network)
- DNS Address—One of the DNS servers in setup_data.yaml::NETWORKING['domain_name_servers']
- DNS Search domain—-- setup_data.yaml::NETWORKING['domain_name']
- Underlay IP address—This is the IP address for Cisco NFVI management/provisioning (mx) network.

- Underlay IP Netmask (mx network)
- NTP address—One of the `setup_data.yaml::NETWORKING['ntp_servers']` addresses
- Password change for user `vts-admin`—Enter the default user `vts-admin` password. The `vts-admin` user is used for password recovery and to revisit a configuration screen if you make a mistake or need to change the information. If you log into the VTC VM using `vts-admin` username and password again, you get the same dialog to go through the VTC VM setup again. This should match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']` or subsequently changed through the VTC UI to match the value in `setup_data.yaml::VTS_PARAMETERS['VTS_PASSWORD']`
- Administrator User—Enter administrative username and password. This username and password are used to login to the VM via SSH.
- Password for administrator user.

VTC VM reboots at this time. Wait for two minutes for the VTC VM to come up. You can ping the IP address given for VTC VM in the setup process to verify whether the VTC VM is up.

Step 6 SSH into VTC VM using the IP address, administrative username/password given in the setup process (not `vts-admin` user).

Installing the VTSR VMs

Before you can install Cisco VTS for Cisco NFVI, you must install the VTSR VM and register it to VTS. VTSR VM is the control plane VM. Installing and registering the VTSR VM requires you to complete the following procedures:

- [Creating VTSR VM , on page 13](#)
- [Creating an ISO for IOS VTSR, on page 14](#)

Creating VTSR VM

The VTSR VM is essential to the Virtual VTEP topology. The VTSR VM contains a nested VM so VTSR must enable nesting.

Before you begin

You must complete a VTS VM installation, and the VTC UI initial password must be changed to the password that you enter for Cisco VIM when you install Cisco VIM. This password is set in `setup_data.yaml` or the Cisco VIM Insight. Login to VTC UI and create a site with Unique UUID and EVPN VxLAN Type. Then, update the site UUID in `setup_data.yaml` as `VTS_SITE_UUID`.

Bringing up the KVM-based VTSR VM

Step 1 Create the VTSR VM XML referring the Cisco NFVI sample (VTSR.XML).

Step 2 Generate an ISO file for the VTSR. See [Creating an ISO for IOS VTSR, on page 14](#) .

Step 3 Create the VM using the XML.

```
virsh create VTSR.xml
```

Creating an ISO for IOS VTSR

To create an ISO file for VTSR:

Step 1 Create the system.cfg file based on the sample below.

Note Verify that the configuration file has no spaces or extra characters.

Note Before you enter the VTS_USERNAME and VTS_PASSWORD, review [Cisco VTS Usernames and Passwords in Cisco NFVI, on page 3](#).

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

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

# The following are the common configurations for VTSR
# VTS Registration Information:
# VTS_ADDRESS should be the VTS IP. 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 is used for
# underlay connectivity, while the second interface is used for the management network.
# For both MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory and used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# not supported. 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, 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 is used for
# underlay connectivity, while the second interface is used for the management network.

# For both MGMT and UNDERLAY networks, a <net-name>_NETWORK_IP_GATEWAY
# variable is mandatory and used for monitoring purposes.
#
# V6 is only supported on the mgmt network and dual stack is
# not supported.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 2 Copy your VTSR system.cfg files to the same path where the script resides. For example:

```
admin:/opt/cisco/package/vts/bin$ ls -l
total 1432
-rwxr-xr-x 1 vts-admin vts-admin 4767 Sep 29 16:40 build_vts_config_iso.sh
-rw-r--r-- 1 root root 1242 Sep 29 23:54 system.cfg
```

Step 3 Create the ISO file as shown below (you need to log in as root):

```
root:/opt/cisco/package/vts/bin# ./build_vts_config_iso.sh vtsr system.cfg.
Validating input.
Generating ISO File. Done!
```

Step 4 Spawn the VTSR VM with the ISO connected to it.

Step 5 Power on the VM.

In case you spawn a new VTSR VM later, it comes up with VTSR Day Zero configuration and get re-registered with the VTC. Use the **sync-to** option available in the Config Sync feature to synchronize the configuration with the latest VTC configuration. See the *Synchronizing Configuration* section for more information on this feature.

Verifying Cisco VTS Installation in Cisco NFVI

The following procedures provide information about how to verify the Cisco VTS installation in Cisco NFVI.

Verifying VTSR VM Installation

To verify VTSR VM installation:

Before you begin

Ensure the tenant network (t) gateway and management network (mx) gateway are reachable from the VTSR server.

Step 1 Log into the VTSR VM using the VTC VM console.

- If you installed the VTC VM in an RedHat KVM based-OpenStack environment, use virt-manager or VNC console to log into the VM. See [Installing VTC VM - Manual Configuration using VNC, on page 12](#)

Step 2 Ping the Cisco NFVI tenant (t) network gateway IP address.

In case ping fails, verify Cisco NFVI tenant network.

Step 3 Ping the VTC Cisco NFVI management/provisioning (mx) network IP address.

In case ping fails, verify the mx network.

Note You should be able to ping the gateway IP address for both Cisco NFVI mx and t networks, as VTSR registers to the VTC using the VTC mx network IP address.

Verifying VTC VM Installation

To verify VTC VM installation:

-
- Step 1** Log into the VTC VM just created using the VTC VM console.
- If you installed the VTC VM in an RedHat KVM based-OpenStack environment, - telnet 0 <console-port> (The console port is the Telnet port in the VTC.xml file.)
- Step 2** Ping the Cisco NFVI api network gateway.
- If ping fails, verify the VM networking to the Cisco NFVI api network.
- Step 3** For the VTC VM CLI, ping the Cisco NFVI management/provisioning (mx) network gateway.
- If ping fails, verify VM networking to the mx network.
- Note** Underlay network gateway is the switched virtual interface (SVI) created for IOSXRv and VTF on the leaf where the controller is connected.
- Step 4** After a few minutes, verify whether the VTS UI is reachable by typing in the VTS api network IP in the browser.
-

Troubleshooting VTF Registration

If VTF registration issues arise, you can use the following commands to find the VTF registration logs on each Cisco NFVI compute node:

```
[root@devstack-71 neutron]# docker exec -it neutron_vtf_4269 bash
[root@devstack-71 /]# cd /var/log/vpfa
[root@devstack-71 vpfa]# ls
vpfa_err.log  vpfa_med.log  vpfa_server.log          vpfa_server_frequent.log  vpfa_stdout.log

vpfa_freq.log  vpfa_reg.log  vpfa_server_errors.log  vpfa_server_slow.log
[root@devstack-71 vpfa]# tail vpfa_reg.log
2016-06-23 02:47:22,860:INFO:VTF-REG: Sent PATCH {"vtf": {"username": "admin",
"vpp-client-name": "devstack-71", "ip": "34.34.34.5", "binding-host-name": "devstack-71",
"gateway-ip": "34.34.34.1", "local-mac": "00:3a:7d:6a:13:c9"}} to
https://172.18.96.15:8888/api/running/cisco-vts/vtfs/vtf
2016-06-23 02:47:23,050:INFO:VTF-REG-ERR: Failure:400!!!
```

A successful log example is shown below:

```
[root@devstack-71 vpfa]# tail vpfa_reg.log
2016-06-23 15:27:57,338:INFO:AUTH: Successful Login - User: admin
URI:/yang-api/datastore/interfaces Host:IPv4Address(TCP, '34.34.34.5', 21345) Method:GET
2016-06-23 15:28:07,340:INFO:AUTH: Successful Login - User: admin
URI:/yang-api/datastore/interfaces Host:IPv4Address(TCP, '34.34.34.5', 21345) Method:GET
```

If a VTF registration fails, check the following:

- IP network connectivity between the compute nodes and the VTC and VTSR VMs (Cisco NFVI tenant and management/provisioning networks)
- VTS_PARAMETERS—The VTS_USERNAME must be admin.
- The VTC and VTSR must be up and the VTS configurations must be applied. The VTSR must be registered with VTC.
- Check that the VTS UI shows "vtsgroup3" in Inventory->Authorization Groups.
- Check that the VTC Admin Username is admin and Device Username is what was set for XRVR_USERNAME in the VTSR config ISO.

Configuring Cisco VTS and VTSR After Installation

The following steps cover the Cisco VTS configurations you need to provision after installation.

Step 1 If you changed the Cisco VTS username/password when you configured the VTS HA configuration, continue with Step 3. If not, log into the Cisco VTS GUI using the default username/password admin/admin.

Step 2 Change the Cisco VTS password using the UI Change Password tab.

Note Before you enter the Cisco VTS password, review [Cisco VTS Usernames and Passwords in Cisco NFVI](#), on page 3.

Step 3 Log into the VTC VM by running the following command:

```
cd /opt/vts/bin
sudo ./vts-cli.sh -applyTemplate vtsr-underlay-loopback-template
```

```
./vts-cli.sh -applyTemplate vtsr-underlay-loopback-template command is applyTemplate and template
name is vtsr-underlay-loopback-template
Enter device name: <hostname of vtsr>
Enter loopback-interface: <loopback interface name>
Enter ipaddress: <loopback interface ip>
Enter netmask: <loopback interface netmask>
```

Similarly configure IGP config in VTSR

Step 4 Log into the VTC VM by running the following command:

```
cd /opt/vts/bin
sudo ./vts-cli.sh -applyTemplate vtsr-underlay-ospf-template
```

```
./vts-cli.sh -applyTemplate vtsr-underlay-ospf-template command is applyTemplate and template name
is vtsr-underlay-ospf-template
Enter device name: <hostname of vtsr>
Enter process-name: <ospf process id >
Enter router-id: <ospf router id>
Enter area-address: <ospf area address>
Enter physical-interface: <VTSR interface connected to NFVI t-network>
Enter loopback-interface: <vtsr loopback interface>
Enter default-cost: <ospf default >
```

Installing VTS in an HA Configuration

Complete the following steps to install Cisco VTS in a Layer 2 HA configuration.

- Step 1** Create two VTC VMs. (In the following steps, these are referred to as VTC1 and VTC2.) When you create the VMs, reserve three IP addresses for each Cisco VIM network to which the VTC VM are connected as described in [Overview to Cisco VTS Installation in Cisco NFVI, on page 1](#) .
- Step 2** If you changed the initial VTC password in a previous installation step, proceed to Step 4. If not, log into the VTC GUI using the default username/password admin/admin.
- Step 3** Change the VTC password using the UI Change Password tab. See [Cisco VTS Usernames and Passwords in Cisco NFVI, on page 3](#) for information about Cisco VTS usernames and passwords.
- Step 4** Edit the cluster.conf file on VTC1 and VTC2 located in /opt/vts/etc/. Both VTCs must have identical information in the cluster.conf file. Parameters includes:
- vip_public—VIP address used for the Cisco VIM API (a) network.
 - vip_private—VIP address used for VTS on the Cisco VIM management/provisioning (mx) network. Cisco VIM uses VTFs, so this field must be entered. The vip_private field is the VIP for the VTS master private interface
 - master_name—Enter the name of the VTC you want to be the primary one in the HA configuration.
 - master_ip—The master VTC IP address used for the Cisco NFVI API network.
 - slave_name—Enter the name of the VTC you want to be the secondary one in the HA configuration.
 - slave_ip—The secondary VTC IP address used for the Cisco NFVI API network.
 - external_ip—The external IP address. This comes from the Cisco VIM setup_data.yaml file after you complete the Cisco VIM installation and the [Cisco VIM Configurations for Cisco VTS Installation](#) procedure.

```
###Virtual Ip of VTC Master on the public interface. Must fill in at least 1
vip_public=
vip_public_ipv6=

###VTC1 Information. Must fill in at least 1 ip address
master_name=
master_ip=
master_ipv6=

###VTC2 Information. Must fill in at least 1 ip address
slave_name=
slave_ip=
slave_ipv6=

###In the event that a network failure occurs evenly between the two routers, the cluster needs an
outside ip to determine where the failure lies
###This can be any external ip such as your vmm ip or a dns but it is recommended to be a stable ip
within your environment
###Must fill in at least 1 ip address
external_ip=
external_ipv6=

#####
### Non-mandatory fields ###
#####
```

```
###If you intend to use a virtual topology forwarder (VTF) in your environment, please fill in the
vip for the underlay as well as the underlay gateway. Otherwise leave blank.
###Virtual Ip of VTC Master on the private interface. You can fill in ipv4 configuration, ipv6, or
both if you use both
vip_private=
private_gateway=

vip_private_ipv6=
private_gateway_ipv6=

###If you have your vtc's in different subnets, xrivr needs to be configured to route traffic and the
below section needs to be filled in
###If you have your vtc's on the same subnet, the below section has be skipped

###Name of your vrf. Example: VTS_VIP
vrf_name=

###Ip of your first Xrivr. Example: 11.1.1.5
xrivr1_mgmt_ip=

###List of neighbors for xrivr1, separated by comma. Example: 11.1.1.1,11.1.1.2
xrivr1_bgp_neighbors=
xrivr1_bgp_neighbors_ipv6=

###Ip of your second Xrivr. Example: 12.1.1.5
xrivr2_mgmt_ip=

###List of neighbors for xrivr2, separated by comma. Example: 12.1.1.1,12.1.1.2
xrivr2_bgp_neighbors=
xrivr2_bgp_neighbors_ipv6=

###Username for Xrivr
xrivr_user=

###Xrivr ASN information
remote_ASN=
local_ASN=

###Xrivr BGP information
bgp_keepalive=
bgp_hold=

###Update source for Xrivr1 (i.e. loopback)
xrivr1_update_source=

###Update source for Xrivr2 (i.e. loopback)
xrivr2_update_source=

###Router BGP Id for Xrivr1
xrivr1_router_id=

###Router BGP Id for Xrivr2
xrivr2_router_id=

###XRVR1 name
xrivr1_name=

###XRVR2 name
xrivr2_name=

###If you plan on having your VTC's on different subnets and intend to use a virtual topology forwarder
(VTF) in your environment,
### please fill out the following fields. Otherwise, leave blank
```

```

###List of neighbors for xrvr1, separated by comma. Example: 2.2.2.2,2.2.2.3
xrvr1_underlay_neighbors=
xrvr1_underlay_neighbors_ipv6=

###List of neighbors for xrvr2, separated by comma. Example: 3.3.3.2,3.3.3.3
xrvr2_underlay_neighbors=
xrvr2_underlay_neighbors_ipv6=

###Directly connected Tor information for Xrvr1
xrvr1_directly_connected_device_ip=
xrvr1_directly_connected_device_ipv6=
xrvr1_directly_connected_device_user=
xrvr1_directly_connected_device_neighbors=
xrvr1_directly_connected_device_neighbors_ipv6=
xrvr1_directly_connected_ospf=
xrvr1_directly_connected_router_id=
xrvr1_directly_connected_update_source=

###Directly connected Tor information for Xrvr2
xrvr2_directly_connected_device_ip=
xrvr2_directly_connected_device_user=
xrvr2_directly_connected_device_neighbors=
xrvr2_directly_connected_device_neighbors_ipv6=
xrvr2_directly_connected_ospf=
xrvr2_directly_connected_router_id=
xrvr2_directly_connected_update_source=

###VPC Peer information if any. Otherwise leave blank
xrvr1_vpc_peer_ip=
xrvr1_vpc_peer_user=
xrvr1_vpc_peer_ospf=
xrvr1_vpc_peer_router_id=
xrvr1_vpc_peer_update_source=

xrvr2_vpc_peer_ip=
xrvr2_vpc_peer_user=
xrvr2_vpc_peer_ospf=
xrvr2_vpc_peer_router_id=
xrvr2_vpc_peer_update_source=

###VTC Underlay Addresses
vtc1_underlay=
vtc2_underlay=
vtc1_underlay_ipv6=
vtc2_underlay_ipv6=

##Gateway of secondary L3 underlay
vtc2_private_gateway=
vtc2_private_gateway_ipv6=

```

Step 5 Execute the cluster installer script, `cluster_install.sh`, located in `/opt/vts/bin/` on VTC1 and VTC2. Do not run the script until have completed Steps 1-5.

```

admin@vtc1:/opt/vts/bin$ sudo ./cluster_install.sh
[sudo] password for admin:
Change made to ncs.conf file.
Need to restart ncs
Created symlink from /etc/systemd/system/multi-user.target.wants/pacemaker.service to
/lib/systemd/system/pacemaker.service.
Created symlink from /etc/systemd/system/multi-user.target.wants/corosync.service to
/lib/systemd/system/corosync.service.
Please run cluster_install.sh on vtc2.waits until finished Both nodes are online.

```

```
Configuring master Configuring Pacemaker resources
Master node configuration finished
HA cluster is installed
```

Note In order for HA to run, the `cluster_install.sh` script updates `/etc/hosts` with the VTC information. If run on the node you specified as master, it completes the basic cluster setup, then wait for the slave to complete. Once the slave is finished, the master completes the remainder of the setup.

When the `cluster_install` script is finished on the master, you can see both the public and private vip using 'ip addr'. If you use VTFs, now that the VIP is up, both VTSRs completes their auto-registration.

Step 6 Verify the HA Status:

```
admin@vtc1:/opt/cisco/package/vtc/bin$ sudo crm status
Last updated: Wed May 4 00:00:28 2016
Last change: Wed May 4 00:00:10 2016 via crm_attribute on vtc2
Stack: corosync
Current DC: vtc2 (739533872) - partition with quorum
Version: 1.1.10-42f2063
2 Nodes configured
4 Resources configured

Online: [ vtc1 vtc2 ]

ClusterIP (ocf::heartbeat:IPaddr2): Started vtc1
Master/Slave Set: ms_vtc_ha [vtc_ha]
Masters: [ vtc1 ]
Slaves: [ vtc2 ]
ClusterIP2 (ocf::heartbeat:IPaddr2): Started vtc1

admin@vtc1:/opt/cisco/package/vtc/bin$ sudo ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 52:54:00:00:bd:0f brd ff:ff:ff:ff:ff:ff
    inet 11.1.1.4/24 brd 11.1.1.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet 11.1.1.2/32 brd 11.1.1.2 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 2001:420:10e:2010:5054:ff:fe00:bd0f/64 scope global dynamic
        valid_lft 2591955sec preferred_lft 604755sec
    inet6 fe80::5054:ff:fe00:bd0f/64 scope link
        valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 52:54:00:4c:11:13 brd ff:ff:ff:ff:ff:ff
    inet 15.15.15.4/24 brd 11.1.1.255 scope global eth1
        valid_lft forever preferred_lft forever
    inet 15.15.15.20/32 brd 11.1.1.20 scope global eth1
```

Completing VTSR HA Configuration

Complete the following steps to set up the VTSR HA configuration:

Before you begin

You must complete a VTS VM installation and change the VTC UI initial password to the password that you enter for Cisco VIM when you install Cisco VIM. This password is set in `setup_data.yaml` or the Cisco VIM Insight.

Login to VTC UI and create a site with Unique UUID and EVPN VxLAN Type. Update this UUID as `VTS_SITE_UUID` in `setup_data.yaml`.

Ensure the tenant network (t) gateway and management network (mx) gateway are reachable from the VTSR server.

Power on the 2 VTSR VM 's as per the VTSR install step. The VTSR VM comes up in active/active HA mode.

Uninstalling VTC HA

To move VTC back to it's original pre-HA state, run the following script on both the active and standby nodes.

```
sudo /opt/vts/bin/cluster_uninstall.sh
```

Sample Cisco VTS Configurations for Cisco NFVI

Sample VTC VM libvirt Domain Configuration

```
<domain type='kvm' id='1332'>
  <name>VTC-release2.1</name>
  <uuid>5789b2bb-df35-4154-ald3-e38cefc856a3</uuid>
  <memory unit='KiB'>32389120</memory>
  <currentMemory unit='KiB'>32388608</currentMemory>
  <vcpu placement='static'>8</vcpu>
  <resource>
    <partition>/machine</partition>
  </resource>
  <os>
    <type arch='x86_64' machine='pc-i440fx-rhel7.0.0'>hvm</type>
    <boot dev='hd'>/>
  </os>
  <features>
    <acpi/>
    <apic/>
    <paef/>
  </features>
  <cpu mode='custom' match='exact'>
    <model fallback='allow'>Westmere</model>
    <feature policy='require' name='vmx'>/>
  </cpu>
  <clock offset='utc'>/>
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
    <emulator>/usr/libexec/qemu-kvm</emulator>
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' cache='none'>/>
      <source file='/home/cisco/VTS2.1/vtc.qcow2'>/>
      <target dev='vda' bus='virtio'>/>
    </disk>
  </devices>
</domain>
```

```

    <alias name='virtio-disk0'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0'/>
</disk>
<controller type='usb' index='0'>
    <alias name='usb0'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x2'/>
</controller>
<controller type='pci' index='0' model='pci-root'>
    <alias name='pci.0'/>
</controller>
<controller type='virtio-serial' index='0'>
    <alias name='virtio-serial0'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0'/>
</controller>
<interface type='bridge'>
    <mac address='52:54:00:5b:12:3a'/>
    <source bridge='br-ex'/>
    <virtualport type='openvswitch'>
        <parameters interfaceid='263c1aa6-8f7d-46f0-b0a3-bdbdad40fe41'/>
    </virtualport>
    <target dev='vnet0'/>
    <model type='virtio'/>
    <alias name='net0'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0'/>
</interface>
<interface type='bridge'>
    <mac address='52:54:00:8d:75:75'/>
    <source bridge='br-control'/>
    <virtualport type='openvswitch'>
        <parameters interfaceid='d0b0020d-7898-419e-93c8-15dd7a08eebd'/>
    </virtualport>
    <target dev='vnet1'/>
    <model type='virtio'/>
    <alias name='net1'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x0b' function='0x0'/>
</interface>
<serial type='tcp'>
    <source mode='bind' host='127.0.0.1' service='4888'/>
    <protocol type='telnet'/>
    <target port='0'/>
    <alias name='serial0'/>
</serial>
<console type='tcp'>
    <source mode='bind' host='127.0.0.1' service='4888'/>
    <protocol type='telnet'/>
    <target type='serial' port='0'/>
    <alias name='serial0'/>
</console>
<channel type='spicevmc'>
    <target type='virtio' name='com.redhat.spice.0'/>
    <alias name='channel0'/>
    <address type='virtio-serial' controller='0' bus='0' port='1'/>
</channel>
<input type='mouse' bus='ps2'/>
<graphics type='spice' port='5900' autoport='yes' listen='127.0.0.1'>
    <listen type='address' address='127.0.0.1'/>
</graphics>
<sound model='ich6'>
    <alias name='sound0'/>
    <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/>
</sound>
<video>
    <model type='qxl' ram='65536' vram='65536' heads='1'/>
    <alias name='video0'/>

```



```

        <address type='pci' domain='0x0000' bus='0x00' slot='0x02' function='0x0' />
    </video>
    <memballoon model='virtio'>
        <alias name='balloon0' />
        <address type='pci' domain='0x0000' bus='0x00' slot='0x07' function='0x0' />
    </memballoon>
</devices>
<seclabel type='dynamic' model='selinux' relabel='yes'>
    <label>system_u:system_r:svirt_t:s0:c26,c784</label>
    <imagelabel>system_u:object_r:svirt_image_t:s0:c26,c784</imagelabel>
</seclabel>
</domain>

```

Sample VTSR VM libvirt Domain Configuration

```

<domain type='kvm' id='20'>
  <name>SAMPLE-VTSR-1</name>
  <memory unit='GiB'>48</memory>
  <cpu mode='host-passthrough' />
  <vcpu placement='static'>14</vcpu>
  <resource>
    <partition>/machine</partition>
  </resource>

  <os>
    <type arch='x86_64' machine='pc-i440fx-rhel7.0.0'>hvm</type>
    <boot dev='hd' />
    <boot dev='cdrom' />
  </os>
  <features>
    <acpi />
    <apic />
    <pae />
  </features>
  <clock offset='localtime' />
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
    <emulator>/usr/libexec/qemu-kvm</emulator>

    <disk type='file' device='cdrom'>
      <driver name='qemu' />
      <source file='/home/admin/VTS20/images/vtsr_node1_cfg.iso' />
      <target dev='hda' bus='ide' />
      <readonly />
    </disk>

    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' />
      <source file='/home/admin/VTS20/images/vtsr.qcow2' />
      <target dev='vda' bus='virtio' />
      <alias name='virtio-disk0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x09' function='0x0' />
    </disk>

    <controller type='usb' index='0'>
      <alias name='usb0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x2' />
    </controller>
    <controller type='ide' index='0'>
      <alias name='ide0' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x1' />
    </controller>
    <controller type='pci' index='0' model='pci-root'>

```

```

    <alias name='pci.0' />
  </controller>

  <interface type='bridge'>
    <source bridge='br-ex' />
    <virtualport type='openvswitch'>
      <parameters interfaceid='4ffa64df-0d57-4d63-b85c-78b17fcac60a' />
    </virtualport>
    <target dev='vtsr-dummy-mgmt' />
    <model type='virtio' />
    <alias name='vnet1' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x02' function='0x0' />
  </interface>

  <interface type='bridge'>
    <source bridge='br-inst' />
    <virtualport type='openvswitch'>
      <parameters interfaceid='4ffa64df-0d67-4d63-b85c-68b17fcac60a' />
    </virtualport>
    <target dev='vtsr-dummy-2' />
    <model type='virtio' />
    <alias name='vnet1' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />
  </interface>

  <interface type='bridge'>
    <source bridge='br-inst' />
    <virtualport type='openvswitch'>
      <parameters interfaceid='4ffa64df-0f47-4d63-b85c-68b17fcac70a' />
    </virtualport>
    <target dev='vtsr-dummy-3' />
    <model type='virtio' />
    <alias name='vnet1' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0' />
  </interface>

  <interface type='bridge'>
    <source bridge='br-inst' />
    <virtualport type='openvswitch'>
      <parameters interfaceid='4ffa64df-0d47-4d63-b85c-58b17fcac60a' />
    </virtualport>
    <vlan>
      <tag id='800' />
    </vlan>
    <target dev='vtsr-gig-0' />
    <model type='virtio' />
    <alias name='vnet1' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0' />
  </interface>

  <interface type='bridge'>
    <source bridge='br-ex' />
    <virtualport type='openvswitch'>
      <parameters interfaceid='3ffa64df-0d47-4d63-b85c-58b17fcac60a' />
    </virtualport>
    <target dev='vtsr-gig-1' />
    <model type='virtio' />
    <alias name='vnet1' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0' />
  </interface>

  <interface type='bridge'>

```

```
<source bridge='br-inst' />
<virtualport type='openvswitch'>
  <parameters interfaceid='a2f3e85a-4de3-4ca9-b3df-3277136c4054' />
</virtualport>
<vlan>
  <tag id='800' />
</vlan>
<target dev='vtsr-gig-2' />
<model type='virtio' />
<alias name='vnet3' />
<address type='pci' domain='0x0000' bus='0x00' slot='0x07' function='0x0' />
</interface>

<serial type='pty'>
  <source path='/dev/pts/0' />
  <target port='0' />
  <alias name='serial0' />
</serial>
<console type='pty' tty='/dev/pts/0'>
  <source path='/dev/pts/0' />
  <target type='serial' port='0' />
  <alias name='serial0' />
</console>
<input type='tablet' bus='usb'>
  <alias name='input0' />
</input>
<input type='mouse' bus='ps2' />
<graphics type='vnc' port='5900' autoport='yes' listen='0.0.0.0' keymap='en-us'>
  <listen type='address' address='0.0.0.0' />
</graphics>
<video>
  <model type='cirrus' vram='9216' heads='1' />
  <alias name='video0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x08' function='0x0' />
</video>
<memballoon model='virtio'>
  <alias name='balloon0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x0a' function='0x0' />
</memballoon>
</devices>
</domain>
```

