

Cisco Virtual Topology System (VTS) 2.6.4 Release Notes

Introduction

This Release Notes document provides information about the new features and enhancements in Cisco VTS 2.6.4 . It also describes how to access information about the known and resolved issues in Cisco VTS 2.6.4 , using Cisco Bug Search Tool.

Cisco VTS 2.6.4 Features and Enhancements

Cisco VTS 2.6.4 supports the following features and enhancements:

- **Supports CISCO VIC 1457**—VTS supports Cisco VIC 1455/1457 on UCS servers. To support Cisco VTC 1457 with VTS, the Cisco VIC must have both UCS and VIC firmware updates.



Note For a proper NIC function, ensure to update the UCS firmware and NIC firmware to the latest UCS HUU version 4.0.2(f). Also, UCS CIMC firmware version should be 4.0.2(f) and VIC firmware version should be 5.0(2b).

- **Support for Static Routes Custom Redistribute Route-Map**—VTS 2.6.4 release supports static route under First Class Object that allows custom route-map to be specified during redistribution in BGP.
- **VTS NED Compatibility**—VTS 2.6.4 is compatible with the following NED versions:
 - ncs-4.6.1.6cisco-vts-2.6.4.tar.gz NSO: 4.6.1.6
 - ncs-5.2.0.1cisco-vts-2.6.4.tar.gz NSO: 5.2.0.1
- **First Class BGP Knobs Enhancement**—VTS 2.6.4 release supports the following features during port-extension:
 - Route-map name for both address families ipv4 and ipv6 under router BGP.
 - Ability to specify a Route-ID under a VRF.
 - Ability to set the maximum paths value for N9k and N7k devices.
 - Ability to add log-neighbor-changes for each router.
- VTS 2.6.4 release supports the following options in first class static route:
 - Ability to track an object in a static route (Track Object).
 - Ability to name a static route.

- Ability to set VRF for port-scoped static route.

- **Performance Enhancements**

Association with Port Extensions for Baremetal and Virtual Machine Ports—To reduce a delay in fetching ports for association with port extensions two filters namely Network and Device are added for Baremetal and Virtual Machine ports. The following functions are supported:

- For the selected specific network, corresponding devices are listed.
- For the selected specific device, view list of matching ports on that device.
- For the select port(s), view the port extension association list to associate with port(s).

Port extensions can be attached to a port only after the VRF has been created. It cannot be attached to a port with no VRF.

Deprecated Features

The following features are deprecated and not tested for VTS 2.6.4 release:

- Device Objects
- Multi VMM
- Mellanox NICs
- Security Groups
- LDAP
- L3 HA
- Segment Routing related features

Limitations and Restrictions

This release has the following limitations/restrictions:

- After an unpublish operation, where a tenant instance remains on OpenStack, if you do a publish operation to the same tenant without cleaning up the tenant instance from OpenStack, the VTS GUI shows it as successful, but the publish operation fails on OpenStack. You must remove the default security group from OpenStack, and also delete the tenant from OpenStack, before you do the publish operation, if you had done an unpublish operation on the same tenant earlier.
- When OpenStack already has a Tenant TA, if you publish the Tenant (TA) with three networks (say, N10, N20, N30) from vCenter to OpenStack, it goes through fine from OpenStack and VTS GUI sides, but Openstack Horizon shows that only N30 is published successfully. OpenStack controller server.log displays an error. This is due to a limitation in OpenStack Horizon.
-
- Multi-site feature has the following constraints
 - VTC HA pair is Global and reachable to all DC sites created in the fabric.
 - VTSr HA pair is Site specific and currently supported as V-deployment in a single site.

- VTF's are site specific and currently restricted only to the site which has VTSr registered to.
- Each VTS installation supports either VXLAN or SR fabrics, but not both. For example, For VXLAN or SR deployments of VTS multi-fabric/sites is supported (VTS will support multiple sites of VXLAN or SR).
- The *remote-as* (neighbor > IP > *remote-as*) attribute is not available in the L3 Service Extension templates for both Cisco Nexus 7000 series and Cisco Nexus 9000 series devices. This is now replaced by *inner-remote-as* (neighbor > IP > *inner-remote-as*) attribute. If you are upgrading to VTS 2.6.1 from an earlier release, and your existing template has the *remote-as* attribute, you must use the template migration utility to ensure that the upgrade is successful. See Cisco VTS 2.6.1 Installation Guide for details.
- Cisco VTS does not allow spaces in Tenant names. Also, Cisco VTS does not support spaces for Tenant names when the Tenants are created from VMMs.
- When you attach a Port Extension to a VM port created from OpenStack, and then change the router name (for example, from RT1 to RT2) from OpenStack, the Cisco VTS GUI still shows the older router name (that is, RT1). Also, device configuration stills shows Tenant-RT1 as the VRF name.
- In Cisco VTS deployments with OVS switches and where VTS agent is in use, the outgoing physnet is not correctly configured if the physnet chosen for OVS is named 'tenant'. Because of this, there is no traffic exiting the compute. Any intra-compute switching works fine. When adding VTSPhysicalNet in the HEAT environment file (in neutron-cisco-vts.yaml), ensure that the value is not set to 'tenant'. This also means that there should not be any compute physnet named 'tenant' in the compute.
- Under one of the following conditions, a tenant in OpenStack does not get any project assigned to it:
 - If a tenant already exists on the target OpenStack VMM and you create a publish policy from the source VMM to this target VMM for that tenant and a network.
 - If you had a publish policy from source VMM to an OpenStack target VMM for a given tenant T1/network N1 and you unpublish the same by removing the policy (at which point only network is removed but tenant is not unpublished), and then, again, recreate publish policy for same target VMM or tenant T1 and network N1 (at which point network N1 is published under same tenant T1, but T1 does not get any project assigned to it).

To workaround this issue, you need to assign project to the tenant manually using OpenStack Horizon GUI or OpenStack API. See [CSCvi38836](#) for further details.

- Monitoring features (collectd and Monit) are not supported for Data Plane (VTF) when VTF is running on vCenter (VM mode).
- In an HA setup, when one of VTSRs is not reachable, Monit GUI displays error and the spinner keeps spinning. Also, in an HA setup, when Monit is not running on one of VTSRs, the Cisco VTS GUI logs out the user while trying to access the Monit Pane for Control Plane. In some cases, it displays an error message. You need to ensure that VTSR is reachable and also bring up Monit on VTSR to workround these issues.
- Cisco VTS L3 service extension template with VRRPv3 for Cisco Nexus 9000 device does not work for devices with NX-OS version 7.0(3)I7(2) or later. It works with version 7.0(3)I7(1).
- Out-of-Band configuration reconciliation feature is not supported in case you:
 1. Attach a device template to a TOR, push underlay VLAN configuration
 2. Do a Baremetal port-attach on the above port
 3. Do any out-of-band configuration on this port, and then use the Reconcile option.



Note Baremetal TOR interfaces are configured with "switchport mode trunk" as part of Day Zero configuration. When baremetal port attach is done to this port, VTC allocated VLAN is pushed to the TOR interface. If you do any OOB configuration on this interface and then do a port-detach, it removes the VTC allocated VLAN as well as "switchport mode trunk" (if last port only). This results in mode change of this interface from trunk to access port. When a subsequent port-attach comes (as trunk), the mode will be changed back to trunk. Therefore, there is no operational impact.

- We recommend that you use Out-of-Band configuration reconciliation feature to reconcile configuration that is pushed to the device via ports created from VTS GUI only. Using this feature to reconcile configuration in a VMM integrated VTS setup, where ports are created from the VMM, might cause errors.
- You must ensure the Day Zero configuration on the device does not include configuration that will be pushed using Cisco VTS services or device templates. That is, device Day Zero configuration should not include configuration which would conflict with the configuration that VTS would be pushing into the device either via service configuration or device template configuration.
- In certain cases, if a port detach operation fails, you may need to remove any related out-of-band configurations from device, do an out-of-band reconcile operation from the Cisco VTS GUI, and then try the port detach operation again.
- When you create an L3 Service Extension template which is incomplete or has some issues, and attach it to a router for the first time, the template does not get attached, but Cisco VTS does not display an error. However, when you attach it subsequently, it displays an error.
- Port Scope Static Routes UI will not show ports (both Baremetal and Virtual Server) for a shared Network for a non-owner tenant. That is, if you create a shared network from OpenStack or VTS UI with the Admin tenant (owner) and a non-owner Tenant (say T1), and spawn the ports (Baremetal/Virtual Server) from the non-owner tenant (T1), port attach will go fine, but in the Port Scope Static Routes UI, the ports will not be seen. Ports spawned from the Admin tenant (owner) are visible in the UI.
- You must delete the Port Scope Static Routes first from the Router and then delete the ports from OpenStack. If we delete the ports from OpenStack without removing Port Scope Static Route, it will end up displaying stale port entries in the UI.
- Due to a limitation in the 12xx series Cisco Virtual Interface Cards, hashing of VXLAN traffic is not fairly load balanced across all the CPUs in VTF multi-core deployments.
- If you import a CSV file, with one of the fields wrongly populated, and then reimport the same CSV file, after correcting the error in the file, the UI does not change, and the previously displayed UI error still exists. This is due to a default browser behavior. To resolve this issue, you may navigate away to another screen, and then come back to the import screen, or refresh the page, before you reimport.
- Compute hosts connected via UCS-B will not get discovered via auto discovery. These hosts needed to be added manually via the network inventory or via importing the inventory CSV.
- A limitation in the Cisco Nexus 9000 series switch model N9K-C9372PX does not allow you to use a native VLAN that is part of a VN-Segment.
- On Cisco Nexus 9000 series devices, if you add static routes, the VRF and static routes are added to the device. If you delete these, the VRF and static routes get deleted from the device. If you again create the same static routes, the VRF gets created, but no static routes are created on the device.

- You must verify that ARP Suppression is supported on the switches where the network will have ports attached. Cisco Nexus 9000 series devices do not support ARP suppression for Fabric/Host networks when SVI is not created. ARP suppression must not be enabled in cases where ARP is used by applications for keep alive and monitoring.
- If you create a Network in OpenStack, and then attach a Baremetal port from Cisco VTS, you must not delete the Network from OpenStack before all Baremetal ports attached to this network are deleted from Cisco VTS.
- If you attach VTS subnets (Baremetal) to a router from Cisco VTS GUI, and then attach the OpenStack subnets to the same router from the Cisco VTS GUI, all subsequent operations on these subnets need to be done from Cisco VTS.
- Baremetal/SRIOV port attached with Security Groups fails if the necessary TCAM carving is not done as part of Day Zero configuration. See the Day Zero Configuration Examples document for details.
- For Cisco Nexus 9000 series devices, an image that fixes the issue described in [CSCvg48830](#) is required in a pure V deployment (VTEP-to-VTEP) of VTS to work with VXLAN, with multicast replication mode. If that image is not available, you may use the proposed workaround in [CSCvg48830](#).
- Tenant VM comes up Active even though the ToR is out of sync. If you create a fabric static route first, and then spin up a VM from OpenStack, Cisco VTS throws the same exception back to OpenStack. Still, the VM goes to active state without displaying any error. Check OpenStack > DOWN/unbound port.
- Due to a limitation in the software, Cisco VTS does not support IPv4 syslog server over management network. Cisco VTS supports only IPv6 syslog server configuration on management network. IPv4 syslog server configuration is, however, supported on underlay network.
- Disabling ARP suppression for a network under the following conditions throws a JAVA exception:
 1. Create an OpenStack network with DHCP port or a VM.
 2. Go to Cisco VTS and enable ARP suppression.
 3. Create a Baremetal port for the same OpenStack network.
 4. Delete the port and Network from OpenStack. The network will be owned by Cisco VTS due to the Baremetal port and will be removed from the OpenStack model.
- Tenant VM comes up Active even though ToR is out of sync. This problem cannot be fixed for reasons mentioned below. There are two possible cases of failure:
 1. No error is logged when port attach failed due to Static route which has incorrect Next Hop VRF. Check OpenStack > DOWN/unbound port.
 2. If you create a fabric static route first and spin up a VM from OpenStack after that, VTC throws the same exception back to OpenStack.

VTS plugin behaves as expected. For a failed port-attach request to VTC, it brings the port to the DOWN/unbound state in OpenStack. The VM status being ACTIVE cannot be influenced due to the following reasons:

- The VM is active and running, but it does not have connectivity via the failed port.
- The VM might potentially have another port attached, from a different network, and not necessarily is in a failed state.
- The VM might potentially have another port attached, from a different network, and not necessarily is in a failed state.

Such a VM will show ACTIVE in OpenStack, but will not show as ACTIVE in VTS.

Also, Cisco VTS cannot change the plugin handling for the failed port and make it delete the port instead of bringing it DOWN (in which case Nova might potentially show an error for the VM). The reason is that the port can be created separately from the VM and attached to it during the port-attach operation.

Such a port is an independent resource of neutron and cannot be deleted. Cisco VTS cannot differentiate between attach of an existing port and a port creation specifically for the VM.

Check OpenStack > DOWN/unbound port when you have loss of connectivity or ping failure, to determine state of the port.

- After the installation of the Host Agent if neutron-vts-agent service is down on the compute host, check whether the compute host has Python module pycrypto installed. If it does not exist, install this module and restart the neutron-vts-agent.
- You can have multiple VPC pair per one DVS. However, a ToR that is part of one VPC pair cannot be part of another VPC pair. For example, you can have ToR 1 and 2 in a VPC Pair and ToR 3 and 4 in another VPC Pair, but not TOR 2 and 3 in a VPC Pair. You can have a mix of Non-VPC and VPC pair in the same DVS.
- On Cisco Nexus 7000 series devices with FEX, a mix of VPC and non-VPC can be in one DVS. On Cisco Nexus 7000 series devices without FEX, a mix of VPC and non-VPC is not allowed in the same DVS.
- Each VMware ESXi Host can have only one DVS at a time.
- If you have a ToR connected to a Baremetal and is already added in Admin Domain, you need to uncheck the L2/L3 GW check box in Admin Domain before you convert that Baremetal to a VM or add a new ESXi Host to the existing ToR, as the vlan-pool gets changed to DVS.
- **Known limitations in the Security Groups feature:**
 1. Cisco Nexus 7000 series device is not supported due to device limitation. Cisco Nexus 7000 series device does not support ACLs on BD/BDI.
 2. Security Group Rules defined for 'Ingress' direction for SRIOV and Baremetal ports apply only to L3 routed traffic; L2 traffic will pass-through unaffected. This limitation is due to Cisco Nexus 9000 series devices not supporting VLAN ACLs in the 'Egress direction'.
 3. Remote Security Group is not supported for non OVS ports, that is SRIOV, Baremetal, and VTF ports. Default Security Group has 'Ingress' rules that use remote Security Group.
 4. Security Groups feature uses iptables based firewall as the driver for OVS. Native OVS firewall driver cannot be deployed as the firewall-driver. This is because Red Hat has not yet (as of Release 7.4) qualified native OVS firewall driver as deployment ready. iptables-based firewall has the limitation that it cannot support Security Groups for OVS trunk ports ('VLAN aware VMs').
 5. Cisco VTS GUI does not take "Type" and "Code" for ICMP rule.
 6. Cisco VTS GUI allows you to only Add or Delete Security Group rules from within a Security Group. Editing of a given Security Group Rule is not permitted. The behavior is same as that of OpenStack Horizon UI.
 7. By default, OpenStack associates the 'default' Security Group when a port is created, even if you do not choose one. As a result, in Cisco VTS 2.6, after OpenStack VMM registration, all ports should be created on a new tenant, or on an existing tenant from an earlier version that follows the upgrade path. For the existing tenants from Cisco VTS 2.5.2, the 'default' Security Group data would already be present in the database, which will be mapped to the new Security Group model when you upgrade to VTS 2.6.

If you do not prefer creating a new tenant, and instead use the admin/existing tenant, you can create a new security group (say SG1), and associate that new group (SG1) during port creation. Under Launch Instance > Security Groups tab, select the one you just created (that is, SG1) and deselect

the default SG. This is only required for new installations. For existing installations via upgrade path, this is optional.

8. Whenever a new tenant is created from OpenStack, OpenStack adds a default security group to it, which Cisco VTS comes to know of it and creates the tenant and security group in its database. In releases earlier than 2.6.0, Cisco VTS used to know about tenant when the very first network was created, and the tenant used to be deleted from its database when the last network gets deleted for that tenant. The current behavior is such that even after last network is deleted, the tenant will not be deleted from the Cisco VTS database since it will have the default security group added to it.

Therefore, the only way to get the tenant deleted from Cisco VTS, after last network for that tenant is deleted, is to use the OpenStack CLI to delete the default security group of the tenant, at which point Cisco VTS removes the tenant from its database. This is not a Cisco VTS limitation, but the way OpenStack handles objects. It allows the deletion of its child objects even if parent objects are present (children become orphaned). So the manual deletion has to be done from OpenStack backend CLI.

9. The older model of Cisco Nexus 9000 series device (such as C9372PX) does not take "switchport trunk native vlan <vlan_id>" which causes trunk port creation to fail. Models 9200 and 9300-EX/FX should work.
10. Cisco VTS utilizes Cisco Nexus 9000 VACLs to realize SG intent on a per VLAN domain:

VTS pushes the following when SRIOV/BM port attached with SG:

- 1 VACL to accommodate IPv4 Allow Rules
- 1 VACL to accommodate IPv4 Deny Rules
- 1 VACL to accommodate IPv6 Allow Rules
- 1 VACL to accommodate IPv6 Deny Rules

Cisco Nexus 9000 has 64 VACL limit. Owing to this restriction, SRIOV ports running on computes connected to a given ToR cannot span beyond 15 dual-stacked networks or 31 single stacked networks. This limit is per-ToR. By spreading SRIOV ports across ToRs, this limit can be overcome.

11. We recommend that Security Groups are designed such that each Security Group has a maximum of 50 Rules. When the number of Rules within a given Security Group ranges to 100 and more, performance degradation is noticed.
12. For VTC to learn the default SG of any tenant, you need to update the description field of the default SG. This can be done from the OpenStack CLI using the following command:

```
[root@overcloud-controller-1 ~]# openstack
(openstack) security group set <default-sg-id> --project-id <project-id>
--description "Updated default SG description"
```

For OSPD setup, run the command from the OSPD Director:

```
[stack@ospd-director ~]$ source overcloudrc
[stack@ospd-director ~]$ openstack
(openstack) security group set <default-sg-id> --project-id <project-id>
--description "Updated default SG description"
(openstack)
```



Note OpenStack Horizon does not let the user update default SGs except on Packstack setups.

- **Known limitations in Syslog feature:**

- There is no uninstall script to cleanup ConfigureSyslog details, or disable option from VTS CLI to clear syslog config. The only way is specify to syslog server as 0.0.0.0 in Logconfig.ini and reconfigure it.
- All the Logs from different sources of VTC are sent to single source.
- VTSR does not support dualstackIP/Hostname for syslog. If you use VTSR with v6 Management then make sure your syslog server support V6 management. VTSR supports syslog only on management network.

- **Known limitations in Multi VMM feature:**

- If you create a Tenant named "admin", and create Network N1 and Subnet S1 on vCenter VMM (version 6.5), and then perform a publish operation from vCenter VMM (version 6.5) to another vCenter VMM (version 6.0) for the Tenant "admin" and Network "N1", the publish operation for the network and subnet fails with exception. The exception says that the network does not exist for corresponding target subnet that is getting published.

This is due to a limitation in vCenter. Do not publish any tenant (and corresponding workload for the tenant) named "admin". Create a tenant with any other name when you need to publish a workload to vCenter 6.0.

- Publish operation of a Provider Network from OpenStack VMM to vCenter VMM is not supported.
- VTS does not delete Provider Networks from UI. Before deleting the overlay networks from the VMM (OpenStack/VMware), the required dependencies created through Cisco VTS must be removed to synchronize between the VMM (OpenStack/VMware) overlay networks and the Cisco VTS networks. If there is no synchronization between the VMMs (OpenStack/VMware) networks and the networks in Cisco VTS the clean-up process can be done manually via REST APIs.
- After you register vCenter as a VMM, and, for the first time, perform a publish operation to publish a tenant and multiple networks to this vCenter VMM, the tenant and networks fail to get published to the VMM. The error next to the policy certificate shows exception related to SSL handshake. Click the Retry button to get the tenant and networks published to the VMM.
- Upon publishing, Cisco VTS does not create the users for a tenant that it creates in OpenStack. To view the tenant project, user has to be assigned to the project. The OpenStack user has to attach a user to the tenant
- Cisco VTS publishes networks to OpenStack as network type = vxlan. Before performing a publish operation, make sure that the plugin.ini, which is located at /etc/neutron/plugin.ini, has the following properties with network type vxlan as one of the values, for example:

```
type_drivers = vxlan, <network_type2>, <network_type3> ... <network_type_n> [comma
separated list of network types]
tenant_network_types = vxlan, <network_type2>, <network_type3> ...<network_type_n>
[comma separated list of network types]
```

- In order to delete a published network/subnet, you have to first unpublish the network, and then perform the delete operation.
- When you create an overlay network/subnet from VTS, publish from VTS to a target VMM (OpenStack Liberty), perform port attach operations, and then unpublish the network from the target VMM, the network and subnet are deleted from target VMM. Cisco VTS throws an exception during the deletion. You can ignore this exception.
- External networks are not eligible to be Multi VMM networks.

- You cannot delete a network or subnet from Cisco VTS after a publish operation. You need to delete the publish operation before you change network or subnet from the source VMM or VTS. If you update from source VMM, the target VMM will not get affected. If you update from the VTS GUI, the update will fail.
- If you have a Multi-VMM setup with two vCenter VMMs, Static Multi Homing across VMMs is not supported.
- Cisco VTS supports only one domain setup in OpenStack. Otherwise, there might be confusing operation done to other domains in the same OpenStack environment.
- For vCenter-based setups, Cisco VTS supports only discovery using the CSV option. Auto Discovery using seed IP is not supported for vCenter-based setups.
- VTF L2 mode is supported only on OpenStack Newton.
- Migration from OVS to VTF is not supported.
- FEX interface group range reverts back to default range when you edit the interface group by adding new modules or devices. However, in physical interface group, the custom range is retained even though you update the interface group with additional devices.
- Cursor does not show up in username and password text boxes on the Cisco VTS login page. This happens only when Cisco VTS UI is accessed using Google Chrome browser, after Cisco VTS is restarted or upgraded, or when it comes up for the first time. Waiting for few minutes gets the focus on the credentials text box on login page. Also, you can still go ahead and enter the credentials which will still show up in the appropriate box.
- Detaching multiple templates attached to devices, which have interdependency between them in terms of configuration, fails, as the order of deletion does not continue till the last template in the list of templates to be detached. You must be aware of the interdependencies in configuration among the templates that you are attempting to detach.

However, detaching multiple templates which do not have interdependent configuration, works fine.

- You may encounter a string conversion error in `ncs-java-vm.log`, while adding a physical device to a device group. This error occurs when Cisco VTS checks whether a device is a VTF and identifies that it is not, and, therefore, not defined by an IP address. You may ignore this error.
- You must create separate DVS Switches on vCenter for port interfaces connected in V-host and P-host. Otherwise, you may face P2V connectivity issues with the port group with Cisco Nexus 7000 devices. Cisco VTS does not support VM migration across DVS.
- IPv4/IPv6 LLDP support is not available for Cisco ASR 9000 devices. Cisco VTS auto discovery does not work for Cisco ASR 9000 devices with IPv4/IPv6.
- For VMM registration, Cisco VTS does not support dual stack configuration on VTC and VMM. Ensure that the same IP version (either IPv4 or IPv6) is used on VTC and the VMM.
- The IPv6 hostname added in the vCenter should match with the hostname you added to the inventory through using CSV or Discovery. If it does not match, you will encounter mac-binding issues when you do a port attach.
- If you add a TACACS+ server with IPv6 address in the Cisco VTS UI, and on the TACACS+ server if IPv6 TACACS port is disabled, the TACACS+ server will not be reachable. Even if the IPv4 TACACS port is enabled on the server, the server will be unreachable as there is no support for roll back to IPv4 when IPv6 fails.
- VRF name change from VTS GUI is not supported for VTSR. Cisco VTS does not allow changing the name of a router if it connects to a port on a V node.
- If you are using VTSR, then BGP ASN value you set should be between 0 and 65535.

- While you specify the VTF credentials when you install VTF via Host Inventory, you must not use *root* as the username. Choosing *root* as username will not allow you to log in to the VTF, after installation. You may choose a username other than *root*.
- When you add bulk VMs on a vhost compute, one VM fails to spawn because of lack of free hugepage. You will see the below log in `/var/log/neutron/server.log`:

```
[Insufficient free host memory pages available to allocate guest RAM]
```

To support vhost-user mode, `numa_nodes` and `mem_page_size` in the OpenStack flavor configuration should be changed as follows:

```
# nova flavor-key m1.medium set hw:numa_nodes=2 (This number is based on the numa nodes
you have on the compute, if you have one set it to 1, if you have two nodes set it to
2)
# nova flavor-key m1.medium set hw:mem_page_size=large (You can do this customized
setting with any flavor)
```

For example:

```
# nova flavor-list
| 9592ec23-0118-4d91-8a71-51375de9e025 | m1.medium | 2048 | 40 | 0 |
| 2 | 1.0 | True |
+-----+-----+-----+-----+-----+
# nova flavor-show 9592ec23-0118-4d91-8a71-51375de9e025
+-----+-----+-----+-----+-----+
| Property | Value |
+-----+-----+-----+-----+-----+
| OS-FLV-DISABLED:disabled | False |
| OS-FLV-EXT-DATA:ephemeral | 0 |
| disk | 40 |
| extra_specs | {"hw:mem_page_size": "2048", "hw:numa_nodes": "2"} |
| id | 9592ec23-0118-4d91-8a71-51375de9e025 |
| name | m1.medium |
| os-flavor-access:is_public | True |
| ram | 2048 |
| rxtx_factor | 1.0 |
| swap | |
| vcpus | 2 |
+-----+-----+-----+-----+-----+
```

Compute log:

```
# ls /sys/devices/system/node/node*
node0 node1
# cat /sys/devices/system/node/node*/hugepages/hugepages-2048kB/free_hugepages
9985
9985
```

- When you add multiple TACACS+ servers via **Administration > Remote Authentication Settings**, and click the **Save**, some times the Cisco VTS UI does not show all the TACACS+ servers you have added. You may need to refresh the page to view all the TACACS+ servers added. Clear the browser cache to solve this issue.
- When you log in using the vCenter VTC plugin, ensure that you log in as a Cisco VTS local admin user, even if you have enabled TACACS+ based external authentication and authorization and have users with Cisco VTS admin privileges configured in the TACACS+ server. Only a Cisco VTS admin user present in the Cisco VTS local database is allowed to log in via the vCenter plugin.
- If TACACS+ server IP, port, and key attributes are updated through REST API, the changes will not have any effect on the AAA functions. You need to update these parameters via the Cisco VTS UI (**Administration > Remote Authentication Settings**).

- In Cisco VTS, you can add two entries for the same TACACS+ server—one with the IP address, and the other with the Hostname. However, you can enable accounting only on one of these servers.
- Cisco Nexus 7000 TORs with TACACS+ configuration is not supported on Cisco VTS. This is due to a limitation in the platform.
- In an OpenStack Liberty environment, bulk port attach causes errors and multiple VMs fails to spawn while you try to attach multiple ports to networks. You need to edit the libvirtd.conf file to increase the keep alive interval. The following configuration is recommended for creation of ten VMs.

```
# vi /etc/libvirt/libvirtd.conf

    keepalive_interval = 5

    keepalive_count = 100
```

- While discovering Cisco ASR 9000 series routers using the Discovery feature, Cisco VTS gets the IP address of the interface which is connected to the neighbor device (spine or border leaf), and not the management interface IP address. You need to manually edit the discovery table to provide the management interface IP address, to ensure that the device is added to the inventory.
- When performing bulk port attach for tenant VMs spawned across multiple hosts with vCenter Web Client using the networking tab, some of the port attach events are missed and not captured or processed. Bulk port attach for a maximum of 8 tenant VMs in the same scenario works without any events being missed. If you need to perform a bulk port attach for more than eight tenant VMs, it would work only if this is being done for tenant VMs on the same host. This is due to a vCenter networking issue, and a case has been opened with VMware to resolve this issue
- Do not use special characters while creating Tenant/Network/Router via vCenter VTS plugin. VTS GUI does not support this, and the name will not show up in the GUI.
- For V-side tenant VMs, if V2V migrations are performed, you need to perform ARP flush and then ping the gateway for MAC-to-IP binding to be complete.
- In the VTS GUI, Loopback IP address is retrieved automatically in network inventory using the Loopback number.
- Cisco Nexus 9000 Series switches running NX-OS version 7.0(3)I2(1) and later do not support VTEP connected to FEX host interface ports.
- After migrating a VM to a different host, Cisco Nexus 9000 switch still shows old host details in the MAC table. However, the BGP routing table (show bgp l2vpn evpn) has the correct details. See [CSCuy77657](#) for more details.
- **BGP Peering limitations with DCI for virtual side:**
 - When BFD session on NH goes down, on the local VTF, NH as well all dependent port scoped static routes are unresolved. On the remote VTFs the NH status is not set to down even though all the port scoped static routes over the NH are deleted. Also, MAC and MAC-IP RT2 are not withdrawn in the BGP when the NH is unresolved.
 - VTSRXR BGP advertises EVPN Route Type 5 with two Router MAC EXTCOMM with one of the Router MAC with value 0000.0000.0000. This issue is seen when local IP VRF routes are advertised as EVPN Route Type 5 NLRI towards EVPN peers configured with encapsulation vxlan. If the remote peer is a Cisco Nexus platform, You can configure the following neighbor inbound route-map on the remote Nexus peer, to drop the duplicate Router MAC EXTCOMM with value 0000.0000.0000:

```
ip extcommunity-list expanded exp_zero permit "0000.0000.0000"
route-map rmac0 permit 10
  set extcomm-list exp_zero delete
route-map vts-subnet-policy permit 10
```

```
router bgp 100
  router-id 80.70.10.1
  address-family ipv4 unicast
  address-family ipv6 unicast
  address-family l2vpn evpn
  neighbor 80.3.2.1
    remote-as 100
  update-source loopback0
address-family l2vpn evpn
  send-community both
  route-reflector-client
  route-map rmac0 in
neighbor 80.3.2.2
  remote-as 100
  update-source loopback0
  address-family l2vpn evpn
  send-community both
  route-reflector-client
  route-map rmac0 in
```

- Sometimes EVPN Route-type 2 with MAC+IP with two labels is not accepted and IP Address not imported to IP VPN VRF under following conditions:
 - BGP Update message has both EVPN Route-type 2 with MAC only and MAC+IP NLRI's.
 - The first NLRI in the Update message is a EVPN Route-type 2 with MAC only (without IP Address) and it is not imported to L2 VRF (bridge-domain) either because L2VRF is not configured or import of L2 VRF does not match the RT EXTCOMM of EVPN Route-type 2.
 - And EVPN Route-type 2 with MAC+IP is present in the same BGP Update message.
- The issue can be avoided by having different BGP attributes for EVPN Route-type 2 with MAC and MAC+IP NLRIs, so that MAC routes and MAC+IP routes are sent in the different BGP Update message. The following configuration can be applied to both both VTSRs to avoid the issue:

```
prefix-set default-route-prefix-set
  0.0.0.0/0 le 32
end-set
!
prefix-set default-route-v6-prefix-set
  ::/0 le 128
end-set

route-policy set_community_xrivr_out
  set extcommunity soo vts additive
  if evpn-route-type is 2 and destination in default-route-prefix-set or destination
  in default-route-v6-prefix-set then
    pass
  elseif evpn-route-type is 2 then
    set community (100:9600) additive
    pass
  else
    pass
  endif
end-policy
```

And apply this policy to all VTSR neighbors in outbound direction:

Sample is below:

```
router bgp 100
neighbor 80.70.10.1
  remote-as 100
  update-source Loopback0
address-family l2vpn evpn
  route-policy recvrt_filter_comm_xrvr_in in
  encapsulation-type vxlan
  route-policy set_community_xrvr_out out    ==> apply policy in out direction
  advertise vpnv4 unicast re-originated
  advertise vpnv6 unicast re-originated
```

- VTF-Vm mode is deprecated or no longer supported in any OpenStack or vCENTER deployments from VTS262 onwards.
- Verify the new requirements when deploying 2.6.2 (especially important for OpenStack users). The new VTC VM resource requirements are:
 - RAM: 32GB
 - Disk: 64GB
- **Issue:** You may notice, devices goes Out Of Service after upgrade from VTS 2.6.1 to 2.6.2 release provided you have device template with class-map->qos configuration in VTS 2.6.1 and that template is attached to a device.

Workaround: After upgrade to vts 2.6.2, if you notice device out of sync with the difference of prematch match-all, then perform the sync-from device operation to bring that config from device into VTS cdb.

Known and Resolved Caveats

You can get details related to the known and resolved issues in Cisco VTS 2.6.4, using the Cisco Bug Search tool. See [Using the Cisco Bug Search Tool](#) for information about how to search for bugs.

Table 1: Resolved Caveats in VTS 2.6.4

Bug ID	Description
CSCvn06073	DpTrans Exception when connecting L2GW & DCGW in Admin Domain
CSCvp42015	vts-bm-port model missing
CSCvq47380	VTS deletes all interfaces in the router instead of the selected interfaces
CSCvp12066	Port Attach fails when Bare Metal is connected via VPC interface but one switch connection is down
CSCvp19795	VTS - BGP summary routes problem

Known Caveats

Issues While Using 25G Links

If you are using the 25G cables, the link may not be established correctly.

To resolve the issue, you need to disable the FEC settings on the following components:

- TOR
- Compute's BIOS settings for Cavium

Configuring FEC Setting on TOR

Run the following commands on TOR (N9K):

```
configure
int e1/33
fec off
```

Configuring FEC on BIOS

Follow these steps to configure the BIOS settings:

1. Choose **BIOS > Advanced**.
2. Scroll to **Cavium** settings.
3. Click **Port Level Configuration**.
4. Set the **Link Mode** to **25G**.
5. Reboot your computer and go to **BIOS > Advanced**.
6. Set **FEC** to **OFF** and reboot to RHEL.
7. Repeat steps 1 through 6 for both the ports.

Partial binding of NIC Ports

A NIC has two ports. A mixed mode, where one port is bound to Kernel and another is bound to DPDK, is not supported.

If you want to bind only one port to VPP, you must unbind the second port from the kernel mode qede driver.

Table 2: Resolved Caveat

Bug ID	Description
CSCvj20903	<p>Host Inventory Redesign:</p> <p>VTC inventory performance improvement introduced a behavior change compare to previous releases. VTC will not create automatic device group if the VPC host is not added to the host inventory.</p>

Using the Cisco Bug Search Tool

Use the Bug Search tool to search for a specific bug or to search for all bugs in a release.

Procedure

Step 1 Go to [Bug Search Tools & Resources](#) on Cisco.com.

Step 2 At the Log In screen, enter your registered Cisco.com username and password; then, click **Log In**. The Bug Search page opens.

Note If you do not have a Cisco.com username and password, you can register for them at <https://tools.cisco.com/IDREG/guestRegistration.do>.

Step 3 To search for a specific bug, enter the bug ID in the Search For field and press **Return**.

Step 4 To search for bugs in the current release:

Click the Search Bugs tab and specify the following criteria:

- a) In the Search For field, enter product name and press Return. (Leave the other fields empty.)
- b) When the search results are displayed, use the filter tools to find the types of bugs you are looking for. You can search for bugs by status, severity, modified date, and so forth.

To export the results to a spreadsheet, click the **Export All to Spreadsheet** link.

For more details on the tool overview and functionalities, check out the help page, located at <http://www.cisco.com/web/applicat/cbsshelp/help.html>.

Related Documentation

The Cisco VTS documentation set consists of:

- *Cisco Virtual Topology System Installation Guide*
- *Cisco Virtual Topology System Release Notes*

These docs are available on Cisco.com.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly What's New in Cisco Product Documentation, which also lists all new and revised Cisco technical documentation, at <http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>.

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS Version 2.0.

