



# Troubleshooting System Issues

This chapter describes how to troubleshoot Cisco Virtual Security Gateway (VSG) system issues.

This chapter includes the following sections:

- [Information About the System, page 8-1](#)
- [Problems with VM Traffic, page 8-2](#)
- [VEM Troubleshooting Commands, page 8-2](#)
- [VEM Log Commands, page 8-3](#)
- [Troubleshooting the Cisco VSG in Layer 3 Mode, page 8-4](#)

## Information About the System

The Cisco VSG provides firewall functionality for the VMs that have vEths with port profiles created by the Virtual Supervisor Module (VSM). To allow the Cisco VSG to function properly, the Cisco VSG should have registered with a Cisco Prime Network Services Controller (Prime NSC) and the Cisco VSG state should be alive on the Cisco VSM.

The example shows how to display information about the system:

```
vsg# show vsg
Model: VSG
HA ID: 142
VSG software version: 5.2(1)VSG2(1.1a) build [5.2(1)VSG2(1.1a)]
PNSC IP: 10.2.71.139
NSC PA version: 2.1(1e)-vsg
```

```
VSG#
VSG# show nsc-pa status
NSC Policy-Agent status is - Installed Successfully. Version 2.1(1e)-vsg
```

Make sure that cisco VSG is alive by entering the **show vservice node detail** command as follows:

```
vsm# show vservice node detail
#Node Information
Node ID:1      Name:VSG-Node-L3
Type:vsg      IPAddr:70.1.0.75 Fail:close L3
Mod State MAC-Addr VVer
4      Alive  --      2
5      Alive  --      2
```

For more information, see the following documents for your release number:

- *Cisco Virtual Security Gateway*

- *Cisco Prime Network Services Controller*
- *Quick Start Guide for Cisco Virtual Security Gateway and Prime Network Services Controller*

## Problems with VM Traffic

When troubleshooting problems with intrahost VM traffic, follow these guidelines:

- Make sure that at least one of the Hyper-V virtual NICs is on the correct DVS port group and is connected.
- If the Hyper-V virtual NIC is down, determine if there is a conflict between the MAC address configured in the OS and the MAC addresses as that are assigned by Hyper-V. You can see the assigned MAC addresses in the Hyper-V file.

When troubleshooting problems with inter host VM traffic, follow these guidelines:

- Determine if there is one uplink sharing a VLAN with the Hyper-V virtual NIC. If there is more than one uplink, they must be in a port channel.
- Ping an SVI on the upstream switch by entering the **show intX counters** command.

## VEM Troubleshooting Commands

This section includes the following topics:

- [Displaying VEM Information, page 8-2](#)
- [Displaying Miscellaneous VEM Details, page 8-2](#)

### Displaying VEM Information

Use the following commands to display Virtual Ethernet Module (VEM) information:

- **vemlog**—Displays and controls VEM kernel logs.
- **vemcmd**—Displays configuration and status information.
- **vemcmd show vem status**—Displays status information.
- **vemcmd show version**—Displays version information.
- **vemcmd show arp all**—Displays the ARP table on the VEM.
- **vemcmd show vsn config**—Displays all the Cisco VSGs configured on the VEM and the Cisco VSG licensing status (firewall on or off).
- **vemcmd show vsn binding**—Displays all of the VM LTL ports to the Cisco VSG bindings.

### Displaying Miscellaneous VEM Details

These commands provide additional VEM details:

- **vemlog show last number-of-entries**—Displays the circular buffer.

This example shows how to display the number of entries in the circular buffer:

```
PS C:\Program Files (x86)\Cisco\Nexus1000V> vemlog show last 5
```

Timestamp	Entry	CPU	Mod	Lv	Message
Oct 13 13:15:52.615416	1095	1	1	4	Warning vssnet_port_pg_data_ ...
Oct 13 13:15:52.620028	1096	1	1	4	Warning vssnet_port_pg_data_ ...
Oct 13 13:15:52.630377	1097	1	1	4	Warning svcs_switch_state ...
Oct 13 13:15:52.633201	1098	1	1	8	Info vssnet new switch ...
Oct 13 13:16:24.990236	1099	1	0	0	Suspending log

- **vemlog show info**—Displays information about entries in the log.

This example shows how to display log entries:

```
PS C:\Program Files (x86)\Cisco\Nexus1000V> vemlog show info
Enabled: Yes
Total Entries: 1092
Wrapped Entries: 0
Lost Entries: 0
Skipped Entries: 0
Available Entries: 6898
Stop After Entry: Not Specified
```

- **vemcmd help**—Displays the type of information you can display.

This example shows how to display the vemcmd help:

```
PS C:\Program Files (x86)\Cisco\Nexus1000V> vemcmd help
show card          Show the card's global info
show vlan [vlan]   Show the VLAN/BD table
show bd [bd]       Show the VLAN/BD table
show l2 <bd-number> Show the L2 table for a given BD/VLAN
show l2 all        Show the L2 table
show port [priv|vsm] Show the port table
show pc            Show the port channel table
show portmac       Show the port table MAC entries
show trunk [priv|vsm] Show the trunk ports in the port table
show stats         Show port stats
```

## VEM Log Commands

Use the following commands to control the vemlog:

- **vemlog stop**—Stops the log.
- **vemlog clear**—Clears the log.
- **vemlog start *number-of-entries***—Starts the log and stops it after the specified number of entries.
- **vemlog stop *number-of-entries***—Stops the log after the next specified number of entries.
- **vemlog resume**—Starts the log but does not clear the stop value.

You can display the list of debug filters by entering the **vemlog show debug | grp vpath** command.

This example shows how to display the list of debug filters:

```
~ # vemlog show debug | Select-String vpath
      vpath          ENWID P ( 95)      ENW      ( 7)
      vpathapi       ENWID P ( 95)      ENW      ( 7)
      vpathfm        ENWID P ( 95)      ENW      ( 7)
      vpathfsm       ENWID P ( 95)      ENW      ( 7)
      vpathutils     ENWID P ( 95)      ENW      ( 7)
      vpathtun       ENWID P ( 95)      ENW      ( 7)
```

# Troubleshooting the Cisco VSG in Layer 3 Mode

This section includes the following topics:

- [show vservice node brief Command Output Indicates Service Node State Is Down](#), page 8-4
- [Traffic with Large Payloads Fails: ICMP Too Big Message Does Not Reach the Client with the Cisco VSG in Layer 3 Mode](#), page 8-5
- [End-to-End Traffic with the Cisco VSG in Layer 3 Mode Fails](#), page 8-5
- [End-to-End Traffic with the Cisco VSG in Layer 3 Mode and Jumbo Frames Fails](#), page 8-5
- [TCP State Checks](#), page 8-6
- [Connection Limit in the Cisco VSG](#), page 8-6

## show vservice node brief Command Output Indicates Service Node State Is Down

This section includes the following topics:

[Cisco VSG with a VN Service Virtual Network Adapter in Layer 3 Mode](#), page 8-4

[Cisco VSGs with Multiple I3-vservice Virtual Network Adapters in Layer 3 Mode](#), page 8-4

### Cisco VSG with a VN Service Virtual Network Adapter in Layer 3 Mode

When encapsulated traffic that is destined to a Cisco VSG is connected to a different subnet other than the virtual network adapter subnet, the VEM does not use the Microsoft Hyper-V host routing table. Instead, the virtual network adapter initiates the Address Resolution Protocol (ARP) for the remote Cisco VSG IP addresses.

You must configure the upstream router to respond by using the proxy ARP feature. If the proxy ARP feature is not configured on the upstream router, ARP fails and the **show vservice node brief** indicates that the service node state is down.

To resolve this issue configure the proxy ARP feature on the router as follows:

```
sg-cat3k-L14-qa(config)# int vlan 3756
sg-cat3k-L14-qa(config-if)# ip proxy-arp
sg-cat3k-L14-qa(config-if)# end
sg-cat3k-L14-qa# sh ip int vlan 3756 | inc Proxy
Proxy ARP is enabled
Local Proxy ARP is disabled
sg-cat3k-L14-qa#
```

### Cisco VSGs with Multiple I3-vservice Virtual Network Adapters in Layer 3 Mode

The data path traffic and the ARP packets for the Cisco VSGs in Layer 3 mode can use any virtual network adapter that is configured on the VEM host for packet forwarding to the Cisco VSG when you enter the **capability I3-vservice** command.

Therefore, all virtual network adapters that are on a VEM host must be able to reach all Cisco VSGs in Layer 3 mode.

If a router is between the virtual network adapters and the Cisco VSGs, all virtual network adapters must have an interface in the router network (VLAN), and all the Cisco VSGs in the Layer 3 mode must have an interface in the router network (VLAN) to ensure that each virtual network adapter has a route to each Cisco VSG.

To resolve this issue, ensure that all l3-vservice virtual network adapters can reach all the Cisco VSGs in Layer 3 mode that are used by the VEM host.

**Note**

You must enable Proxy ARP on all the interfaces of the router that is alongside the virtual network adapters.

## Traffic with Large Payloads Fails: ICMP Too Big Message Does Not Reach the Client with the Cisco VSG in Layer 3 Mode

If a router lies between the virtual network adapter and the Cisco VSG device in Layer 3 mode, and the router receives a packet that it cannot forward due to a large packet size, the router generates an ICMP Too Big message for the virtual network adapter. The virtual network adapter cannot forward the ICMP Too Big message of the router to the client and the virtual network adapter drops the message. The client never receives the ICMP Too Big message and cannot refragment the packet for successful end-to-end traffic and the end-to-end traffic fails. This problem is typically seen if the router interface to the VEM is set at a higher maximum transmission unit (MTU) than the router interface to the Cisco VSG. For example, the router interface to the VEM has an MTU of 1600 and the interface to the Cisco VSG has an MTU of 1500.

This problem can be seen as an increase in the ICMP Too Big Rcvd counter in the **show vservice statistics** command.

To resolve this issue, configure an oversized MTU (for example, 1600) on both of the router interfaces.

## End-to-End Traffic with the Cisco VSG in Layer 3 Mode Fails

When the VEM communicates with the Cisco VSG in Layer 3 mode, an additional header with 82 bytes is added to the original packet. The VEM does not support fragmentation in Layer 3 mode and the ports or network elements (which carry a vPath encapsulated packet) must be configured in such a way that the vPath overhead is accommodated.

If end-to-end traffic fails in Cisco VSG Layer 3 mode, set the uplink MTU to 1582 bytes to accommodate the additional overhead. This solution assumes that the client and server VM MTUs are at the default of 1500 bytes.

## End-to-End Traffic with the Cisco VSG in Layer 3 Mode and Jumbo Frames Fails

Traffic with the Layer 3 encapsulation fails even with the uplink MTU set to 9000 bytes.

VSG on Microsoft Hyper-V does not support jumbo frames.

## TCP State Checks

By default, TCP state checks are disabled in Cisco vPath for the traffic protected by the Cisco VSG. Sometimes, you might see delays in the TCP traffic. You can enable TCP state checks to diagnose the issue.

Check the following counters at the VSM in the **show vservice statistics** output:

```
vsm# show vservice statistics | grep "TCP chkfail"
      TCP chkfail InvalACK    0          TCP chkfail SeqPstWnd    0
      TCP chkfail WndVari          0
```

This example shows how to enable the TCP state checks on a VSM:

```
VSM(config)# vservice global type vsg
VSM(config-vsn)# tcp state-checks
VSM(config-vsn)#
```

## Connection Limit in the Cisco VSG

The Cisco VSG can have up to 256,000 active connections at any given point of time. If for some reason new connections slow down or connections see too many failures, you can check the Cisco VSG for any connection limits that it experiences. If the VEM-to-Cisco VSG connection is not smooth or have some issues that indicates that the Cisco VSG might have missed a few updates from vPath which results in an accumulation of large active connections in its flow table.

This example shows how to check the active connection count on the Cisco VSG:

```
vsg# show service-path statistics | inc "Active Connections"
Active Flows          48 Active Connections          24
```