



# 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 the Layer 3 Mode, page 8-4](#)

## Information About the System

The Cisco VSG provides firewall functionality for the VMs that have the 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 (PNSC) and the Cisco VSG data interface MAC address should be seen by the VSM.

The example shows how to display information about the system:

```
vsg# show vsg
Model: VSG
HA ID: 218
VSG Software Version: 4.2(1)VSG1(1) build [4.2(1)VSG1(1)]
PNSC IP: 10.193.77.223
VSG-PERF-1_1#
VSG-PERF-1_1# show nsc-pa status
NSC Policy-Agent status is - Installed Successfully. Version 1.0(1j)-vsg
```

Make sure that the Cisco VSG MAC address is learned by the VSM by entering the **show vservice node detail** command as follows:

```
vsm# show vservice node detail
#Node Information
#Node ID:1      Name:vasatDbd5
  Type:asa      IPAddr:172.8.8.201      Fail:open  Vxlan:bd5555
  Mod State    MAC-Addr                VVer
    4 Alive    00:50:56:b5:37:8f      2
#Node ID:13     Name:vsgl2tD104
  Type:vsg      IPAddr:10.10.10.104     Fail:open  Vlan:504
  Mod State    MAC-Addr                VVer
    4 Alive    00:50:56:b5:6d:36      2
```

```
6 Alive 00:50:56:b5:6d:36 2
```

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

- *Cisco Virtual Security Gateway*
- *Cisco Virtual Network Management Center*
- *Quick Start Guide for Cisco Virtual Security Gateway and Virtual Network Management Center*

## Problems with VM Traffic

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

- Make sure that at least one of the VMware virtual NICs is on the correct DVS port group and is connected.
- If the VMware 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 VMware. You can see the assigned MAC addresses in the .vmx file.

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

- Determine if there is one uplink sharing a VLAN with the VMware 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 8-3](#)

## 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
- **vem-support all**—Displays support information
- **vem status**—Displays status information
- **vem 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
- **vemcmd show learnt**—Displays all of the VMs that have been learned by the VEM

## 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:

```
[root@esx-cos1 ~]# 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 svb_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:

```
[root@esx-cos1 ~]# 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:

```
[root@esx-cos1 ~]# 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 | grep vpath** command.

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

```
~ # vemlog show debug | grep 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 the 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 and Jumbo Frames Fails, page 8-5](#)
- [TCP State Checks, page 8-5](#)
- [Connection Limit in the Cisco VSG, page 8-6](#)
- [Debugging the Traffic Flow Via a Service Chain, page 8-6](#)
- [Troubleshooting the Service Chain by Excluding the Cisco VSG Node, page 8-7](#)
- [VEM/vpath Configured Correctly on a VEthernet Interface for a ServiceChain, page 8-7](#)
- [Cisco VSG on a VXLAN is not working, page 8-7](#)

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

This section includes the following topics:

- [Cisco VSG with a VN Service vmknic in Layer 3 Mode, page 8-4](#)
- [Cisco VSGs with Multiple I3-vn-service vmknics in Layer 3 Mode, page 8-5](#)

### Cisco VSG with a VN Service vmknic in Layer 3 Mode

When encapsulated traffic that is destined to a Cisco VSG is connected to a different subnet other than the vmknic subnet, the VEM does not use the VMware host routing table. Instead, the vmknic initiates an 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, the 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-vn-service vmknics in Layer 3 Mode

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

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

If a router is between the vmknics and the Cisco VSGs, all vmknics 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 vmknic has a route to each Cisco VSG.

To resolve this issue ensure that all I3-vn-service vmknics can reach all the Cisco VSGs in the 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 vmknics.

## 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 vmknic and the Cisco VSG in the 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 vmknic. The vmknic cannot forward the ICMP Too Big message of the router to the client and the vmknic 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.

If L3 pre-frag is enabled, traffic is allowed and fragmentation happens, even when MTU is less than 1582. If L3 pre-frag is disabled, traffic fails for MTU less than 1582.

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

If L3 Pre-frag is enabled, traffic is allowed. If L3 Pre-frag is disabled, traffic fails.

If jumbo frames are enabled in the network and the end-to-end traffic fails, make sure that the MTU of the client and server VMs are 82 bytes smaller than the uplink. For example, if the uplink MTU is 9000, set the MTU of the client and server VMs to 8918 bytes.

## TCP State Checks

By default, TCP state checks are disabled in vPath for the traffic protected by the Cisco VSG. Sometimes, you might see delays in the TCP traffic. To diagnose TCP state checks related issues, you need to enable TCP state checks.

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 slows 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
```

## Debugging the Traffic Flow Via a Service Chain

When configured, the service-chain functionality enables traffic to flow through the Cisco VSG and the Cisco ASA 1000V cloud firewall. The Cisco VSG monitors the data packets and authorizes its flow from the VM to the destination ports. The VM and Cisco ASA 1000V are always in the same broadcast domain, that is, either a VLAN or a Virtual Extensible Local Area Network (VXLAN).

To debug the traffic flow via the service chain, follow these steps:

- 
- Step 1** Make sure that the VM's default gateway is set to the ASA 1000V inside interface and is reachable.
  - Step 2** On the VSM, ensure that the Cisco VSG and ASA 1000V are alive, which ensures that the vPath is able to reach the service nodes.

```
vsm# show vservice node brief
#Node Information
ID Name Type IP-Address Mode State Module
 2 VSG vsq 192.168.10.1 v-140 Alive 3,4
 6 ASA asa 3.3.3.1 v-200 Alive 3,4
```

- Step 3** On the VSM, check a connection's status of action (SAct).

```
vsm# show vservice connection
# Module 1
Proto SrcIP[:Port] SAct DstIP[:Port] DAct Flags Bytes
icmp 192.168.10.15 Pp 192.168.11.15 882
```

In the SAct value Pp, the uppercase 'P' indicates the action that is initiated by the Cisco VSG, while the lowercase 'p' indicates the action that is deduced based on the returning traffic from the ASA V1000. If the SAct value is 'rr,' it indicates that the traffic is redirecting to either the Cisco VSG or the ASA V1000 but no response is being received.

- Step 4** On the VSM, verify that the service node version information (VVer) is '2' so that it works in the service-chain.

```
vsm# show vservice node detail
#Node Information
#Node ID:2      Name:VSG
Type:asa       IPAddr:192.168.10.1   Fail:open  Vlan:140
Mod  State      MAC-Addr           VVer
 3  Alive      00:50:56:a6:02:a5   2
 4  Alive      00:50:56:a6:02:a5   2
#Node ID:6      Name:ASA
Type:asa       IPAddr:3.3.3.1       Fail:open  Vlan:200
Mod  State      MAC-Addr           VVer
 3  Alive      00:50:56:a6:02:6d   2
 4  Alive      00:50:56:a6:02:6d   2
```

## Troubleshooting the Service Chain by Excluding the Cisco VSG Node

The service-chain configuration has the Cisco VSG and ASA 1000V nodes in its service path for a given traffic flow. For debugging purposes, the Cisco VSG can be removed temporarily from the node configuration to isolate a problem. Thus, a user can verify the traffic flow with just the ASA 1000V. Later, the Cisco VSG can be added again to restore the original service-chain configuration using the two said service nodes.

## VEM/vpath Configured Correctly on a VEthernet Interface for a ServiceChain

You can use the `module vem vem-num execute vemcmd show vservice bindings` command on the VSM to ensure that the bindings are correctly configured on the VEM for a service chain. Two entries appear for a single LTL—one of each service node must be displayed.

```
vsm# module vem 3 execute vemcmd show vsn bindings
VSG Services Enabled | VSG Licenses Available 2
ASA Services Enabled | ASA Licenses Available 2
  LTL  PATH  VSN  SWBD   IP          P-TYPE  P-ID
 60   6     2   3756  10.10.10.202 1 49
 60   6     33  3770  172.31.2.11  2 52 >> two service node bindings for LTL
60 and 62
 62   6     2   3756  10.10.10.202 1 49
 62   6     33  3770  172.31.2.11  2 52
```

## Cisco VSG on a VXLAN is not working

The Cisco VSG node can be configured with a VXLAN in the Layer 2 mode only. Make sure that the adjacency is correctly defined as Layer 2 and the bridge-domain configuration is valid. The `show service node brief` command can be used to check a service node's state with respect to the vPath.

This example shows a Cisco VSG node configuration for a VXLAN:

```
vservice node VSG-vxlan-33071 type vsg
 ip address 20.20.20.182
 adjacency l2 vxlan bridge-domain 33071
 fail-mode close
```

This example shows how to display the node status in a VXLAN:

```
vsm# show vservice node brief
#Node Information
```

## ■ Troubleshooting the Cisco VSG in the Layer 3 Mode

ID	Name	Type	IP-Address	Mode	State	Module
35	VSG-vxlan-33071	vsg	20.20.20.182	vxlan	Alive	3,4,5