



Configuring L3 Forwarding

This chapter contains the following sections:

- [Layer 3 Forwarding Overview, page 1](#)
- [Enabling and Verifying Layer 3 Forwarding, page 3](#)
- [Viewing Layer 3 Forwarding Information, page 4](#)
- [Monitoring Layer 3 Forwarding Statistics, page 6](#)
- [Layer 3 Forwarding Guidelines and Limitations, page 7](#)
- [Feature History for Layer 3 Forwarding, page 8](#)

Layer 3 Forwarding Overview




Note Layer 3 Forwarding requires a Cisco Nexus 1000V Advanced Edition license.

In a typical, centralized Layer 3 forwarding model, a Layer 3 router (virtual and physical) receives packets from a Cisco Nexus 1000V and forwards the traffic across the segments. In this model, the Layer 3 router can become a point of congestion or blockage for the flow of traffic. For example, in the following figure, data

packets from VM1 are routed to the Layer 3 router. The Layer 3 router decides where the data packets need to go and forwards the packets to VM4.


Figure 1: Centralized Layer 3 Forwarding Model



In a distributed forwarding model, the VSM manages all the configurations and the VEMs are instantiated on each host to provide packet switching functionality. In this model, the VSM shares the VM packet routing information with the VEMs, so that the VEMs can forward the packets to the correct host. Distributed forwarding reduces the traffic that is sent to the Layer 3 router because the VEMs send the packets directly to the destination VM. For example, in the following figure, the VEM is aware of VM1 and VM2 routing

information. The VEM automatically directs the traffic from VM1 to VM4 and VM2 to VM3. There is no longer a need to forward the packet information to the Layer 3 router.

Figure 2: Distributed Layer 3 Forwarding Model



402634

Enabling and Verifying Layer 3 Forwarding

Before You Begin

Log in to the CLI in EXEC mode.



Note

Layer 3 Forwarding requires a Cisco Nexus 1000V Advanced Edition license.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature l3forwarding	Enables the Layer 3 forwarding feature.
Step 3	switch(config)# show feature	(Optional) Displays the enabled status for Cisco Nexus 1000V features.

This example shows how to enable the Layer 3 forwarding feature and display the output:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature l3forwarding
switch(config)# show feature
Feature Name           Instance  State
-----
cts                    1        disabled
dhcp-snooping         1        disabled
http-server           1        enabled
lACP                   1        disabled
netflow               1        disabled
network-segmentation 1        enabled
port-profile-roles   1        disabled
private-vlan          1        disabled
segmentation          1        enabled
sshServer             1        enabled
tacacs                1        disabled
telnetServer          1        disabled
vtracker              1        disabled
vxlan-gateway         1        disabled
l3forwarding          1        enabled
switch(config)#
```

Viewing Layer 3 Forwarding Information

Use the following commands to view Layer 3 forwarding information:

**Note**

Make sure that you are logged into the VEM when issuing **vemcmd** commands.

Command	Purpose
vemcmd show ip-forwarding-table	Displays the complete IP forwarding table.
vemcmd show l3-forwarding-table <i>l3-table-id</i>	Displays the Layer 3 forwarding table. If a Layer 3 table ID is not specified, then the complete Layer 3 forwarding table is displayed.
vemcmd show l2 segment <i>segment-id</i>	Displays the router-mac for that segment.

Command	Purpose
show segment statistics module [vlan bridge-domain-name] <i>number</i>	Displays segment statistics for the specified VLAN or bridge domain
show l3-segment-attribute-table [vlan bridge-domain-name] <i>number</i>	Displays the Layer 3 segment attribute table for the specified VLAN or bridge domain.
show interface counters	Displays related interface counter information.

This example shows how to display information about Layer 3 forwarding:

```

switch# show ip-forwarding-table
Flags: (Rtr)=Router MAC; (L)=Local; (R)=Remote;
VLAN/SEGID|L3 TableID| MAC | IP | Flags
-----+-----+-----+-----+-----
1172 9 FA:16:3E:49:88:D6 192.168.72.65 L,Rtr
1170 9 FA:16:3E:2D:87:5B 192.168.70.101 L
1170 9 FA:16:3E:42:8C:AF 192.168.70.50 L,Rtr
1170 9 FA:16:3E:E4:8D:8A 192.168.70.104 L
1171 9 FA:16:3E:1A:06:0A 192.168.71.2 L

VEM# vemcmd show l3-forwarding-table 1
L3-table-id      IP address      mac address      BD
-----
1                192.168.1.150  bc:16:65:22:ac:42 130
1                192.168.1.48   bc:16:65:22:ac:42 130
1                192.168.1.179  bc:16:65:22:ac:42 130
1                192.168.1.92   bc:16:65:22:ac:42 130

VEM# vemcmd show flow-mgr l3-flows
Flow-id L3-table-id      IP address      mac address      BD
-----
0        5000              10.10.163.20    00:16:3e:a9:03:c8 163
1        5000              10.10.163.64    00:16:3e:20:a9:b4 163
2        5000              10.10.162.63    00:16:3e:20:a9:a3 162
3        5000              10.10.162.10    00:1b:35:ab:45:0e 162

VEM# vemcmd show l2 segment 50001
Bridge domain 11 brtmax 4096, brtcnt 3, timeout 300
Segment ID 50001, swbd 4096, "bd1"
Flags: P - PVLAN S - Secure D - Drop R - Router-mac

Type      MAC Address      LTL  timeout  Flags  PVLAN Remote IP  DSN  Slot
-----
Static    52:54:00:98:b4:ff  65   0        0      0.0.0.0
0
Static    52:54:00:62:12:3a  63   0        0      0.0.0.0
0
Static    52:54:00:61:13:bd  0    0        0      R      0.0.0.0
0

switch# show segment statistics module 3
VLAN/      Rx      Rx      Tx      Tx Missed Missed Dropped Dropped
BD  Pkts  Bytes  Pkts  Bytes  Pkts  Bytes  Pkts  Bytes
1    0    0    0    0    0    0    0    0
3972 0    0    0    0    0    0    0    0
3970 0    0    0    0    0    0    0    0
3968 0    0    0    0    0    0    0    0
3971 0    0    0    0    0    0    0    0

switch# show l3-segment-attribute-table
-----
Segment-id  Segment-type  Attribute  Value
-----

```

```

111127 Vxlan Router IP          45.11.9.1
111127 Vxlan Router MAC       FA:16:3E:8B:59:05
111127 Vxlan SUBNET           0.0.0.0/0
111126 Vxlan Router IP          45.11.8.1
111126 Vxlan Router MAC       FA:16:3E:CD:ED:A1
111126 Vxlan SUBNET           0.0.0.0/0
111125 Vxlan Router IP          45.11.7.1
111125 Vxlan Router MAC       FA:16:3E:B9:A7:D2
111125 Vxlan SUBNET           0.0.0.0/0
111124 Vxlan Router IP          45.11.6.1
111124 Vxlan Router MAC       FA:16:3E:4D:D9:20
111124 Vxlan SUBNET           0.0.0.0/0
111123 Vxlan Router IP          45.11.5.1
111123 Vxlan Router MAC       FA:16:3E:8F:3E:48

```

```
switch# show interface counters
```

```

-----
Port                               InOctets                               InUcastPkts
-----
mgmt0                               846142352                               1456395
Eth3/1                              234693677                               48980
Eth4/1                              14229614                                4606
Eth5/1                              198530588                               21751
Eth5/2                              201360061                               35320
Eth6/1                              276841979                               3298
Eth7/1                              72027394                                153
Eth7/2                              74577517                                22113
Po1                                  276808574                               3298
Po2                                  399811656                               57064
Po3                                  146577970                               22259
Veth1                                987879                                  3671
Veth2                                343513                                  2618

```

Monitoring Layer 3 Forwarding Statistics

Use the following commands to view Layer 3 forwarding statistics:



Note

Make sure that you are logged into the VEM when issuing **vemcmd** commands.

Command	Purpose
vemcmd show stats	Displays general Layer 3 forwarding port statistics.
vemcmd show packets	Displays Layer 3 forwarded packets.
vemcmd show bd stats [vlan segment bridge-domain-name] number	Displays Layer 3 forwarded packets per BD.
vemcmd clear bd stats [vlan segment bridge-domain-name] number	Clears the bridge domain statistics for the specified VLAN, segment, or bridge domain.
show l3-segment-attribute-table [vlan bridge-domain-name] number	Displays the Layer 3 segment attribute table for the specified VLAN or bridge domain.

Command	Purpose
show interface counters	Displays related interface counter information.

This example shows how to view Layer 3 forwarding statistics :

```
VEM# vemcmd show stats
LTL Received      Bytes      Sent      Bytes      RxL3frwd      Bytes      TxL3frwd
   Bytes Txflood    Rxdrop    Txdrop    Name          Bytes
52   525  4          50666     483          vnet0         47182      121
   7032  4          0          0            vnet0         7096
53   520  0          50352     478          vnet2         46844      119
   7085  0          0          0            vnet2         7085      119
```

```
VEM# vemcmd show packets
LTL RxUcast TxUcast RxMcast TxMcast RxBcast TxBcast RxL3frwd TxL3frwd
Txflood Rxdrop Txdrop RxJumbo TxJumbo Name
52  2026  2000  16  16  18  0  121  120
16  0  0  0  0  0  0  0  0
53  2026  2000  0  0  16  0  119  119
0  0  0  0  0  0  0  0  0
vnet2
```

```
VEM# vemcmd show bd stats vlan 107
BD L3Rx      Bytes      L3Tx      Bytes      L3Rxmiss      bytes
6  97        6456      95        6359        0              0
```

L3RxMiss - Miss in the L3 hash table for /32 addresses.

Layer 3 Forwarding Guidelines and Limitations

Layer 3 forwarding has the following configuration guidelines and limitations:

- Layer 3 forwarding must be enabled before system host setup or the information in the forwarding tables will be inconsistent. To enable Layer 3 forwarding on active VSMs, you must reload the VSM.
- Layer 3 forwarding is not supported for packets with VXLAN encapsulation received from VMs behind a VEM, such as a VXLAN gateway.
- Same segment Layer 3 forwarding is supported, but ICMP redirect messages are not generated.
- In Anycast forwarding (non-gateway forwarding) mode, external traffic is forwarded using the gateway. Also, packets with a router_mac destination are dropped if there is no matching entry in the Layer 3 forwarding table. An ICMP unreachable message will not be generated.
- There can be only one gateway per segment.
- In Openstack mode, there can be only one subnet in a network. Multiple subnets in one network is not supported.
- QoS and security policies applied to packets on the Layer 3 router are skipped in the distributed Layer 3 forwarding model.
- Destination interface MTU validation is not done in VEM forwarding. There have been no traffic issues observed in testing between VMs on the same VEM.
- There can only be one router per tenant.
- VTEP IP address changes may result in transient packet loss for a brief period.
- The network cannot be changed from shared to non-shared and vice-versa.

- A MAC cannot be associated with multiple IP addresses.
- The following are not supported:
 - PVLAN with Layer 3 forwarding.
 - Localization of Layer 3 forwarding entries in VLAN deployments.
 - IPv6 Layer 3 forwarding.
 - Multicast Layer 3 forwarding.
 - Layer 3 forwarding to and from shared segments.

Feature History for Layer 3 Forwarding

Feature Name	Release	Feature Information
Layer 3 Forwarding	5.2(1)SK3(2.2)	This feature was introduced.