



VXLAN Load Balancing

- [VXLAN Load Balancing, on page 1](#)

VXLAN Load Balancing

Virtual extensible LAN (VXLAN) load balancing ensures that data moves efficiently between Cisco Application Centric Infrastructure (ACI) Virtual Edge and the leaf switch over multiple network interfaces under both of the following circumstances:

- When you have a MAC pinning policy and use VXLAN encapsulation
- When the MAC pinning policy and VXLAN encapsulation are enabled on the Cisco ACI Virtual Edge virtual machine manager (VMM) domain in the Cisco Application Policy Infrastructure Controller (APIC)

Beginning with this Cisco ACI Virtual Edge release, VXLAN load balancing is enabled by default. This Cisco ACI Virtual Edge release adds interfaces to accommodate VXLAN load-balancing and improve overall performance.



Note VXLAN load balancing is not supported for Cisco ACI Virtual Edge when Cisco ACI Virtual Edge is part of Cisco ACI Virtual Pod (vPod mode).

In previous releases, Cisco ACI Virtual Edge had three interfaces: one management, one internal, and one external. In VMware vCenter, there were two port groups, **internal** and **external**. Cisco ACI Virtual Edge now has the following:

- Two internal interfaces: Handle data traffic from the virtual machines (VMs). Traffic from private VLANs (PVLANS) is split evenly between two new internal port groups—**ave-internal-1** and **ave-internal-2** in VMware vCenter.
- Two external VXLAN interfaces: Load balance the VXLAN traffic. There are two new port groups in VMware vCenter—**ave-external-vxlan-1** and **ave-external-vxlan-2**, one for each interface. Infra VLAN, used by OpFlex, also uses the two external VXLAN interfaces.
- One external VLAN interface: Handles all VLAN-tagged traffic except for infra VLAN. It has its own VMware vCenter port group **ave-external-vlan**, which allows all Cisco ACI fabric VLANs, based on the VMM configuration.
- Management interface: Unchanged from previous releases.

- Two virtual tunnel endpoints (VTEPS) (kni interfaces) are created automatically and are pinned to each of the external-vxlan interfaces.



Note The names of the new VMware vCenter port groups are assigned automatically. Do not use these new port groups—which begin with **ave**—for tenant traffic.

In VMware vCenter, the **internal** and **external** port groups are still present. These port groups remain to accommodate Cisco ACI vPod and for upgrade and downgrade compatibility.

Verify VXLAN Load Balancing

You can verify that VXLAN load balancing is enabled by checking whether the Cisco Application Centric Infrastructure (ACI) Virtual Edge has received a DHCP IP address and whether OpFlex is up.

You run **ifconfig** commands to view kernel NIC information and check that Cisco ACI Virtual Edge has received a DHCP address. You run the **vemcmd show opflex** command to see if OpFlex is up.

View Kernel NIC Information

You can view information about the kernel NIC.

Enter the following commands: **ifconfig kni0** and **ifconfig kni2**, and check that kni0 and kni2 have been assigned IP addresses:

Example:

```
cisco-ave_198.51.100.62_AVE-FI:~$ ifconfig kni0
kni0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 198.51.100.111 netmask 255.255.0.0 broadcast 198.51.100.255
    inet6 fe80::250:56ff:feaf:807b prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:af:80:7b txqueuelen 1000 (Ethernet)
    RX packets 528552 bytes 50610919 (48.2 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 285294 bytes 44487029 (42.4 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
cisco-ave_198.51.100.62_AVE-FI:~$ ifconfig kni2

kni2: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 198.51.100.121 netmask 255.255.0.0 broadcast 198.51.100.255
    inet6 fe80::250:56ff:feaf:3dc9 prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:af:3d:c9 txqueuelen 1000 (Ethernet)
    RX packets 285152 bytes 17116682 (16.3 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 10873 bytes 2921194 (2.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
cisco-ave_198.51.100.62_AVE-FI:~$
```

The output displays information about the two virtual TEP internal uplinks.

Note You can enter the **ifconfig** command to view complete interface information, including information about ens160, the management interface, and kni1, the ave-ctrl interface for vMotion.

View OpFlex Information

You can see if OpFlex is online and view its runtime status.

Enter the following command: **vemcmd show opflex**.

Example:

```
cisco-ave_198.51.100.62_AVE-FI:~$ vemcmd show opflex
Status: 12 (Active)
Channel0: 12 (Active), Channel1: 12 (Active)
Dvs name: comp/prov-VMware/ctrlr-[AVE-FI]-vC-191/sw-dvs-413
Remote IP: 192.0.2.11 Port: 8000
Infra vlan: 5
FTEP IP: 192.0.2.20
Switching Mode: LS
Encap Type: VXLAN
NS GIPO: 225.10.10.1
cisco-ave_198.51.100.62_AVE-FI:~$
```
