

Configure Accelerated Networking

What is Accelerated Networking

Accelerated networking enables single root I/O virtualization (SR-IOV) on VMs such as a Cisco Catalyst 8000V VM. The accelerated networking path bypasses the virtual switch, increases the speed of network traffic, improves the networking performance, and reduces the network latency and jitter.

Usually, all the networking traffic in and out of the VM traverses the host and the virtual switch. However, with accelerated networking, the network traffic arrives at the virtual machine's network interface (NIC), and is then forwarded to the VM. Thus, all the network policies that the virtual switch applies are now offloaded and applied in the hardware.

For more information about the accelerated networking functionality that is available in Microsoft Azure, see Create a Linux VM With Accelerated Networking Using Azure CLI.

Accelerated networking is available in Cisco Catalyst 8000V public cloud deployments and in government cloud deployments.

Support for Azure-PMD

The Azure-PMD (Poll Mode Driver) functionality on Azure offers a faster, user-space packet processing framework for performance-intensive applications. This framework bypasses the virtual machine's kernel network stack. In a typical packet processing that uses the kernel network stack, the process is interrupt-driven. When the network interface receives the incoming packets, there is an interruption to the kernel to process the packet and a context switch from the kernel space to the user space. Azure-PMD eliminates the context switching and the interrupt-driven method in favor of a user-space implementation that uses poll mode drivers for fast packet processing.

You can enable the Azure-PMD functionality for Cisco Catalyst 8000V running on Microsoft Azure. This functionality increases the performance of the Cisco Catalyst 8000V instance when compared to the previous versions that use accelerated networking.

Supported VM Instance Types

The following VM instance types support the Accelerated Networking functionality:

IOS XE Version	Supported VM Instance Types
17.4.x and later	DS2_v2 / D2_v2 DS3_v2 / D3_v2 DS4_v2 / D4_v2

IOS XE Version	Supported VM Instance Types
	F16s_v2 F32s_v2

Support for Mellanox Hardware

Microsoft Azure cloud has two types of hardware that support the accelerated networking functionality. The following table specifies the Mellanox versions supported for the accelerated networking functionality.

Table 1: Compatibility Matrix of IOS Versions and Accelerated Networking

IOS XE Version	Support for Accelerated Networking	Support for MLX4	Support for MLX5	Support for Azure-PMD
17.4.x and later	Yes	Yes	Yes	Yes

Note Currently, Mellanox ConnectX-3 (CX3) vNIC uses the MLX4 drivers, and ConnectX-4 vNIC (CX4) uses the MLX5 drivers. You cannot specify which NIC Azure must use (MLX4 or MLX5) for your VM deployment.

In the Cisco IOS XE 17.4.1 release, support for the Azure DPDK failsafe/TAP/MLX IOD model was added for both CX3 and CX4 drivers. From the Cisco IOS XE 17.8.1 release, the DPDK failsafe/TAP/MLX I/O model has been replaced with the DPDK NETVSC PMD I/O model. With this update, you experience less overhead while using the accelerated networking functionality.

Note To enable the throughput license performance, you must enable the accelerated networking functionality.

- Enable Accelerated Networking, on page 2
- Disable Accelerated Networking, on page 4
- Verifying Accelerated Networking, on page 4

Enable Accelerated Networking

To enable accelerated networking on a Cisco Catalyst 8000V instance, run the `router# show platform software system hypervisor` command.

```
Router#show platform software system hypervisor
Hypervisor: AZURE
Manufacturer: Microsoft Corporation
Product Name: Virtual Machine
Serial Number: 0000-0016-9163-0690-4834-7207-16
UUID: 80cbc2ea-29e6-cc43-93e9-f541876836f2
Image Variant: None

Cloud Metadata
-------------------
Region: eastus
```

```
Zone:
Instance ID: eac2cb80-e629-43cc-93e9-f541876836f2
Instance Type: Standard_DS4_v2
Version:
Image ID:
Publisher:
Offer:
SKU:

Interface Info
------------------
Interface Number : 0
IPv4 Public IP: 192.168.61.135
IPv4 Private IP: 10.0.0.4
IPv4 Subnet Mask: 255.255.0.0
IPv4 Network: 192.168.0.3
IPv4 Gateway: 10.0.0.1
MAC Address: 000D3A103B48

Interface Number : 1
IPv4 Public IP:
IPv4 Private IP: 10.0.1.4
IPv4 Subnet Mask: 255.255.0.0
IPv4 Network: 192.168.1.3
IPv4 Gateway: 10.0.0.1
MAC Address: 000D3A103348

Interface Number : 2
IPv4 Public IP:
IPv4 Private IP: 10.0.4.4
IPv4 Subnet Mask: 255.255.0.0
IPv4 Network: 192.168.2.3
IPv4 Gateway: 10.0.0.1
MAC Address: 00224827BA0F

Interface Number : 3
IPv4 Public IP:
IPv4 Private IP: 10.0.3.4
IPv4 Subnet Mask: 255.255.0.0
IPv4 Network: 192.168.3.3
IPv4 Gateway: 10.0.0.1
MAC Address: 00224827B2A6

Interface Number : 4
IPv4 Public IP:
IPv4 Private IP: 10.0.4.4
IPv4 Subnet Mask: 255.255.0.0
IPv4 Network: 192.168.4.3
IPv4 Gateway: 10.0.0.1
MAC Address: 00224827B5CB
```



Caution Due to a Microsoft Azure limitation, enabling accelerated networking on all the interfaces of a Cisco Catalyst 8000V router might cause a significant performance drop if packets greater than 1500 bytes are sent across the Azure infrastructure. The performance degradation occurs because Azure starts fragmenting the packets at 1438 bytes and drops out the sequence packets. This is a known issue and a support case is currently opened with Microsoft.

To enable accelerated networking, create or modify a vNIC using the **az network nic** command and the `--accelerated-networking` option. See the Microsoft Azure documentation for the az network nic command and also refer to the following examples.

Note Depending on how you created the Cisco Catalyst 8000V instance, accelerated networking might initially be disabled on the Cisco Catalyst 8000V NICs. If accelerated networking is disabled on the NIC and you want to enable accelerated networking on an interface, use one of the commands as shown in the following examples.

Example 1

This example shows how to create a vNIC "mynic1" and enable accelerated networking using the **az network nic create** command with the `--accelerated-networking true` option.

```
az network nic create -n mynic1 -g "RG1" --accelerated-networking true -l "east us"
--vnet-name "vnetname" --subnet "subnet1"
```

Example 2

This example shows how to create a vNIC "mynic2" and enable accelerated networking using the **az network nic create** command with the `--accelerated-networking true option` option.

```
az network nic create -n "mynic2" -g "RG1" --accelerated-networking true -l "east us"
--vnet-name "vnetname" --subnet "subnet1"
```

Example 3

This example shows how to modify a vNIC "mynic3" to enable accelerated networking using the **az network nic update** command with the `--accelerated-networking true` option.

```
az network nic update -n mynic3 -g rg1 --accelerated-networking true
```

Disable Accelerated Networking

To disable accelerated networking for Cisco Catalyst 8000V, you can create or modify a vNIC using the **az network nic** command and the `--accelerated-networking` option.

For more information about the command, see the Microsoft Azure documentation for the az network nic command.

Example

This example shows how to modify a vNIC "mynic1" to disable Accelerated Networking using the **az network nic update** command with the `--accelerated-networking false` option.

```
az network nic update -n "mynic1" -g rg1 --accelerated-networking false
```

Verifying Accelerated Networking

After Enabling accelerated networking on the NICs, use the following IOS commands to verify whether accelerated networking is enabled on the NIC. The Azure infrastructure uses Mellanox NICs to achieve SR-IOV or accelerated networking.

You can use the following commands to verify Cisco Catalyst 8000V NICs by using the Mellanox kernel drivers as the NIC's I/O drivers to process the packets. In addition, the Mellanox NICs in the HyperV server

of the Azure infrastructure presents a bonded interface to the Cisco Catalyst 8000V guest VM. This VM is used for accelerated networking, and the VM is in a bonded state whenever accelerated networking is enabled.

Verifying Accelerated Networking for Cisco Catalyst 8000V 17.4.x (With Azure-PMD)

After enabling accelerated networking on the NICs, use the following IOS commands to verify whether accelerated networking with Azure-PMD is enabled on NIC. The Azure infrastructure uses Mellanox NICs to achieve SR-IOV or accelerated networking.

Use the following commands to verify the Cisco Catalyst 8000V NICs by using the Mellanox Azure-PMD drivers as the NIC's I/O drivers to process the packets. In addition, the Mellanox NICs in the HyperV server of the Azure infrastructure presents a bonded interface to the Cisco Catalyst 8000V guest VM. This VM is used for accelerated networking, and the VM is in a bonded state while accelerated networking is enabled. Note that the bonded interfaces share the same MAC address. The aggregate counters appear on Gi interfaces, while the non-accelerated packets counters appear on the net_tap interfaces. The accelerated packets counters appear on the net_mlx interfaces.

In the following example, the interface Gi2 indicates that a majority of the packets are flowing over the net_mlx interface.

```
Router#show platform hard qfp act dat pmd controllers | inc NIC|good_packets
NIC extended stats for port 0  (Gi1) net_failsafe 000d.3a8f.1bf1 xstats count 13
  rx_good_packets: 411
  tx_good_packets: 326
NIC extended stats for port 1  (Bonded) net_mlx5 000d.3a8f.1bf1 xstats count 35
  rx_good_packets: 389
  tx_good_packets: 326
NIC extended stats for port 2  (Bonded) net_tap 000d.3a8f.1bf1 xstats count 13
  rx_good_packets: 22
  tx_good_packets: 0
NIC extended stats for port 3  (Gi2) net_failsafe 000d.3a8f.1040 xstats count 13
  rx_good_packets: 10638289
  tx_good_packets: 3634525
NIC extended stats for port 4  (Bonded) net_mlx5 000d.3a8f.1040 xstats count 35
  rx_good_packets: 10639534.  ==>>> This verifies Accelerated Networking is working properly
 for RX
  tx_good_packets: 3636099    ==>>> This verifies Accelerated Networking is working properly
 for TX
NIC extended stats for port 5  (Bonded) net_tap 000d.3a8f.1040 xstats count 13
  rx_good_packets: 291
  tx_good_packets: 0
NIC extended stats for port 6  (Gi3) net_failsafe 000d.3a8f.1a90 xstats count 13
  rx_good_packets: 3637187
  tx_good_packets: 10522981
NIC extended stats for port 7  (Bonded) net_mlx5 000d.3a8f.1a90 xstats count 35
  rx_good_packets: 3638631
  tx_good_packets: 10524554
NIC extended stats for port 8  (Bonded) net_tap 000d.3a8f.1a90 xstats count 13
  rx_good_packets: 28
  tx_good_packets: 0
```

Verifying Accelerated Networking for Cisco Catalyst 8000V 17.8.x (With Azure PMD)

From the Cisco IOS XE 17.8.1 release, the previous DPDK failsafe/TAP/MLX I/O model has been replaced with the DPDK NETVSC PMD I/O model. Use the following commands to verify the accelerated networking functionality on a Cisco Catalyst 8000V running on Cisco IOS XE Release 17.8.x.

The **show platform hardware qfp act dat pmd controllers** command displays the devices that are bonded to the net_netvsc ports.

```
Router#show platform hardware qfp active datapath pmd controllers | inc NIC |good_packets
NIC extended stats for port 0 (Gi2) net_netvsc 000d.3a10.3348 xstats count 56
rx_good_packets: 411
tx_good_packets: 350
tx_q0_good_packets: 311
rx_q0_good_packets: 100
vf_rx_good_packets: 487
vf_tx_good_packets: 350
NIC extended stats for port 1 (Gi1) net_netvsc 000d.3a10.3b48 xstats count 56
rx_good_packets: 60359
tx_good_packets: 55464
tx_q0_good_packets: 6579
rx_q0_good_packets: 5633
vf_rx_good_packets: 53780  ==>>> This verifies Accelerated Networking is working properly
for RX
vf_tx_good_packets: 49831  ==>>> This verifies Accelerated Networking is working properly
for TX
NIC extended stats for port 2 (Gi4) net_netvsc 0022.4827.b2a6 xstats count 56
rx_good_packets: 0
tx_good_packets: 0
tx_q0_good_packets: 0
rx_q0_good_packets: 0
vf_rx_good_packets: 0
vf_tx_good_packets: 0
NIC extended stats for port 3 (Gi5) net_netvsc 0022.4827.b5cb xstats count 56
rx_good_packets: 0
tx_good_packets: 0
tx_q0_good_packets: 0
rx_q0_good_packets: 0
vf_rx_good_packets: 0
vf_tx_good_packets: 0
NIC extended stats for port 4 (Gi3) net_netvsc 0022.4827.ba0f xstats count 56
rx_good_packets: 0
tx_good_packets: 0
tx_q0_good_packets: 0
rx_q0_good_packets: 0
vf_rx_good_packets: 0
vf_tx_good_packets: 0
NIC extended stats for port 5 (Bonded) net_mlx4 0022.4827.b2a6 xstats count 13
rx_good_packets: 0
tx_good_packets: 0
NIC extended stats for port 6 (Bonded) net_mlx4 0022.4827.b5cb xstats count 13
rx_good_packets: 0
tx_good_packets: 0
NIC extended stats for port 7 (Bonded) net_mlx4 000d.3a10.3b48 xstats count 13
rx_good_packets: 54726
tx_good_packets: 65464
NIC extended stats for port 8 (Bonded) net_mlx4 0022.4827.ba0f xstats count 13
rx_good_packets: 363863
tx_good_packets: 105245
NIC extended stats for port 9 (Bonded) net_mlx4 000d.3a10.3348 xstats count 13
rx_good_packets: 0
tx_good_packets: 0
```

The **show platform software vnic-if interface-mapping** command indicates that net_netvsc driver is used from the Cisco IOS XE 17.8.1 release.

```
show platform software vnic-if interface-mapping
------------------------------------------------------------
 Interface Name        Driver Name         Mac Addr
------------------------------------------------------------
 GigabitEthernet3       net_netvsc         000d.3a4e.7542
 GigabitEthernet2       net_netvsc         000d.3a4e.7163
 GigabitEthernet1       net_netvsc         000d.3a4e.757d
------------------------------------------------------------
```

The **show platform software vnic database** command indicates whether MLX4 or MLX5 is present and also indicates the PMD that is used.

```
show platform software vnic-if database
vNIC Database
  eth00_1572882209232255500
    Device Name : eth0
    Driver Name : mlx5_pci
    MAC Address : 000d.3a4e.757d
    PCI DBDF    : b421:00:02.0
    Server      : IFDEV_SERVER_KERN
    Management  : no
    Status      : bonded
  eth01_1572882212261074300
    Device Name : eth1
    Driver Name : mlx5_pci
    MAC Address : 000d.3a4e.7542
    PCI DBDF    : 83e2:00:02.0
    Server      : IFDEV_SERVER_KERN
    Management  : no
    Status      : bonded
  eth02_1572882215293497600
    Device Name : eth2
    Driver Name : mlx5_pci
    MAC Address : 000d.3a4e.7163
    PCI DBDF    : be1d:00:02.0
    Server      : IFDEV_SERVER_KERN
    Management  : no
    Status      : bonded
  eth_15__1572882218326526600
    Device Name : Gi1
    Driver Name : hv_netvsc
    MAC Address : 000d.3a4e.757d
    PCI DBDF    : 000d3a1f-26f8-000d-3a1f-26f8000d3a1f
    Server      : IFDEV_SERVER_UIO
    Management  : no
    Status      : supported
  eth_16__1572882223436559900
    Device Name : Gi2
    Driver Name : hv_netvsc
    MAC Address : 000d.3a4e.7163
    PCI DBDF    : 000d3a1f-26f8-000d-3a1f-26f8000d3a1f
    Server      : IFDEV_SERVER_UIO
    Management  : no
    Status      : supported
  eth_17__1572882228553741500
    Device Name : Gi3
    Driver Name : hv_netvsc
    MAC Address : 000d.3a4e.7542
    PCI DBDF    : 000d3a1f-26f8-000d-3a1f-26f8000d3a1f
    Server      : IFDEV_SERVER_UIO
    Management  : no
    Status      : supported
```