



Release Notes for Cisco UCS Virtual Interface Card Drivers, Release 4.3

First Published: 2023-03-06

Last Modified: 2023-11-14

Release Notes for Cisco UCS Virtual Interface Card Drivers, Release 4.3

Introduction

This document contains information on new features, resolved caveats, open caveats, and workarounds for Cisco UCS Virtual Interface Card (VIC) Drivers, Release 4.3 and later releases. This document also includes the following:

- Updated information after the documentation was originally published.
- Related firmware and BIOS on blade, rack, and modular servers and other Cisco Unified Computing System (UCS) components associated with the release.

The following table shows the online change history for this document.

Revision Date	Description
November 2023	Initial release of VIC drivers for Cisco UCS Software Release 4.3(2c).
August 2023	Initial release of VIC drivers for Cisco UCS Software Release 4.3(2b).
March 2023	Initial release of VIC drivers for Cisco UCS Software Release 4.3(1a).

New Hardware in Release 4.3

New Hardware in Release 4.3(2)

Release 4.3(2c) adds support for the following:

Cisco UCS VIC cards

Following Cisco UCS VIC Cards are supported from release 4.3(2c) onwards:

- Cisco UCS VIC 15230 - The Cisco UCS VIC 15230 is a 4x25-Gbps and 2x100G Ethernet/FCoE capable modular LAN On Motherboard (mLOM) designed exclusively for for Cisco X-series M6/M7 Compute Nodes. The Cisco UCS VIC 15230 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

- Cisco UCS VIC 15427 - The Cisco UCS VIC 15427 is a quad-port small-form-factor pluggable (SFP+/SFP28/SFP56) mLOM card designed for Cisco UCS C-series M6/M7 rack servers. The card supports 10/25/50-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.
- Cisco UCS VIC 15237 - The Cisco UCS VIC 15237 is a dual-port small-form-factor pluggable (QSFP/QSFP28/QSFP56) mLOM card designed for Cisco UCS C-series M6/M7 rack servers. The card supports 40/100/200-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Release 4.3(2b) adds support for the following:

Cisco UCS VIC cards

Following Cisco UCS VIC Cards are supported from release 4.3(2b) onwards:

- Cisco UCS VIC 15425—The Cisco UCS VIC 15425 is a quad-port small-form-factor pluggable (SFP+/SFP28/SFP56) PCIe card designed for Cisco UCS C-series M6/M7 rack servers. The card supports 10/25/50-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.
- Cisco UCS VIC 15235—The Cisco UCS VIC 15235 is a dual-port quad small-form-factor pluggable (QSFP/QSFP28/QSFP56) PCIe card designed for Cisco UCS C-series M6/M7 rack servers. The card supports 40/100/200-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

New Hardware in Release 4.3(1)

Cisco UCS VIC Cards

Following Cisco UCS VIC Cards are supported from release 4.3(1) onwards:

- Cisco UCS VIC 15420—The Cisco UCS VIC 15420 is a 4x25-Gbps Ethernet/FCoE capable modular LAN On Motherboard (mLOM) designed exclusively for Cisco UCS X210c M6/M7 Compute Node. The Cisco UCS VIC 15420 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.
- Cisco UCS VIC 15422—The Cisco UCS VIC 15422 is a 4x25-Gbps Ethernet/FCoE capable mezzanine card (mezz) designed exclusively for Cisco UCS X210c M6/M7 Compute Node. The card enables a policy-based, stateless, agile server infrastructure that can present PCIe standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs.

VIC 15422 is supported with VIC 15420 and can not install VIC 15422 and VIC 15231 together within the same server . This is true for both m6 and M7 blade servers. Also we need to mention to install VIC 15422 mezz card, we need 15420 MLOM + UCSX-V5-BRIDGE

New Features in Release 4.3

New Features in Release 4.3(2)

Release 4.3(2) adds support for the following:

- Receive Side Scaling Version 2 (RSSv2) with Windows NENIC driver for UCS VIC 15000 series adapters.
- QinQ (802.1Q-in-802.1Q) support for Cisco UCS VIC 1400, 14000, and 15000 series adapters.
- Netflow Monitoring support on Cisco UCS 6400 and 6500 series Fabric Interconnects.
- SRIOV support on Cisco UCS Manager.

VIC Driver Updates for Release 4.3

VIC Driver Updates for Release 4.3.2



Note VIC drivers for all operating systems listed in the HCL are cryptographically signed by Cisco.

ESX ENIC Driver Updates

ESX NENIC Version 2.0.11.0

NENIC/NENIC(RDMA) version 2.0.11.0 is supported with ESX 8.0, ESX 8.0U1.

ESX NENIC Version 2.0.10.0

NENIC/NENIC(RDMA) version 2.0.10.0 is supported with ESX 7.0U1, ESX 7.0U2, ESX 7.0U3.

ESX NENIC_ENS Version 1.0.6.0

NENIC_ENS version 1.0.6.0 is supported with ESX 7.0 U1, ESX 7.0 U2, ESX 7.0 U3 and ESX 8.0.

ESX FNIC Driver Updates

Native FNIC Version 5.0.0.41

Native FNIC driver version 5.0.0.41 is supported with ESX 7.0U1, ESX 7.0U2, ESX 7.0U3, ESX 8.0, ESX 8.0U1.



Note Driver version 5.0.0.41 supports both native FC and FC NVME functionality.
 ESX FC NVME is supported with VIC 1400 and 15000 series adapters.
 FDMI is supported with Native FNIC driver version 5.0.0.x on VIC 1400 and 15000 series adapters.
 Interrupt mode INT-x is not supported with ESX nfnic and nenic drivers.

Linux ENIC Driver Updates

ENIC Version 939.x

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.9, 8.2, 8.4, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2
- XenServer 8.2
- SUSE Linux Enterprise Server 12 SP5, 15 SP4, 15 SP5
- Ubuntu Server 20.04, 20.04.1, 20.04.2, 20.04.3, 20.04.4, 20.04.5, 20.04.6, 22.04, 22.04.1, 22.04.2
- CentOS 7.9

Linux FNIC Driver Updates

Unified FNIC Driver 2.0.0.9x

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.9, 8.2, 8.4, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2
- XenServer 8.2
- SUSE Linux Enterprise Server 12 SP5, 15 SP4, 15 SP5
- CentOS 7.9



Note fNIC multiqueue is supported on on RHEL 8.4, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2, SLES 12 SP5, SLES 15 SP4 and SLES 15 SP5.

FC-NVME is supported for VIC 14xx and VIC 15xxx adapters on RHEL 8.4, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2, SLES 12 SP5, SLES 15 SP4 and SLES 15 SP5.

FDMI is supported for VIC 14xx and VIC 15xxx adapters on RHEL 8.4, 8.6, 8.7, 8.8, 9.0, 9.1, 9.2, SLES 12 SP5, SLES 15 SP4 and SLES 15 SP5.

SLES 15 FC-NVMe is supported with DM multi-pathing and native multi-pathing is not supported.

RHEL inbox nvme-cli 1.14 is not working as expected with FC-NVME. Using RHEL 8.4 binaries(nvme-cli 1.12) with RHEL 8.5 FC- NVME is recommended.

Windows 2022, 2019 and 2016 NENIC/ENIC Driver Updates

Windows Server 2022 and 2019 NENIC Version 5.13.24.2

- This driver update provides an VMMQ & RDMA driver for VIC 1400 and 15000 Series Adapters and Supported QoS Changes.

This driver update provides Receive Side Scaling Version 2 (RSSv2) on UCS VIC 15000 series adapter

Windows Server 2016 NENIC Version 5.8.25.9

- This driver update provides an RDMA driver for VIC 1400 and 15000 Series Adapters and Supported QoS Changes.

Windows Server 2022, 2019 and 2016 ENIC Version 4.4.0.12

- This driver update provides a Spectre-compliant driver for VIC 1300 Series adapters.

Windows 2022, 2019 and 2016 FNIC Driver Updates

Windows Server 2022, 2019 and 2016 FNIC Version 3.3.0.11

- This driver update provides a Spectre-compliant fNIC driver for VIC 15XXX, 14XX and VIC 13XX adapters.

VIC Driver Updates for Release 4.3.1



Note VIC drivers for all operating systems listed in the HCL are cryptographically signed by Cisco.

ESX ENIC Driver Updates

ESX NENIC Version 1.0.45.0

NENIC Version 1.0.45.0 is supported with ESX 7.0U1, ESX 7.0U2, ESX 7.0U3, ESX 8.0.

ESX NENIC_ENS Version 1.0.6.0

NENIC_ENS version 1.0.6.0 is supported with ESX 7.0 U1, ESX 7.0 U2, ESX 7.0 U3 and ESX 8.0.

ESX FNIC Driver Updates

Native FNIC Version 5.0.0.37

Native FNIC driver version 5.0.0.37 is supported on ESX 7.0 U1, ESX 7.0 U2, ESX 7.0 U3 and ESX 8.0.



Note Driver version 5.0.0.37 supports both native FC and FC NVME functionality.

ESX FC NVME is supported with VIC 1400 and 15000 series adapters.

FDMI is supported with Native FNIC driver version 5.0.0.x on VIC 1400 and 15000 series adapters.

Interrupt mode INT-x is not supported with ESX nfnic and nenic drivers.

RDMA Support Driver Updates

ESX NENIC_RDMA Version 2.0.4.0 NENIC_RDMA version 2.0.4.0 is supported with ESX 7.0U3, and ESX 8.0.

Linux ENIC Driver Updates**ENIC Version 930.x**

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.9, 8.2, 8.4, 8.6, 8.7, 9.0, 9.1
- XenServer 8.2
- SuSE Linux Enterprise Server 12 SP5 15 SP4
- Ubuntu Server 20.04, 20.04.1, 20.04.2, 20.04.3, 20.04.4, 20.04.5, 22.04, 22.04.1
- CentOS 7.9

Linux FNIC Driver Updates**Unified FNIC Driver Update 2.0.0.89**

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.9, 8.2, 8.4, 8.6, 8.7, 9.0, 9.1
- XenServer 8.2
- SuSE Linux Enterprise Server 12 SP5, 15 SP1, 15 SP2, 15 SP3, 15 SP4
- CentOS 7.9



Note fNIC multiqueue is supported on RHEL 7.9, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 9.0 and 9.1, and SLES 12 SP5, SLES 15 SP1, SLES 15 SP2, SLES 15 SP3 and SLES 15 SP4.

FC-NVME is supported for VIC 14xx and VIC 15xxx adapters on RHEL 7.9, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 9.0 and 9.1, and SLES 12 SP5, 15 SP1, 15 SP2, SLES 15 SP3 and SLES 15 SP4.

FDMI is supported for VIC 14xx and VIC 15xxx adapters on RHEL 7.6, 7.7, 7.8, 7.9, 8.0, 8.1, 8.2, 8.3, 8.4, 8.6, 8.7, 9.0 and 9.1, and SLES 12 SP3, SLES 12 SP4, SLES 12 SP5, SLES 15, SLES 15 SP1, SLES 15 SP2, SLES 15 SP3 and SLES 15 SP4.

SLES 15 FC-NVMe is supported with DM multi-pathing and native multi-pathing is not supported. RHEL inbox nvme-cli 1.14 is not working as expected with FC-NVME. Using RHEL 8.4 binaries(nvme-cli 1.12) with RHEL 8.5 FC- NVME is recommended.

Windows 2022, 2019 and 2016 NENIC/ENIC Driver Updates**Windows Server 2022 and 2019 NENIC Version 5.12.9.1**

- This driver update provides an VMMQ & RDMA driver for VIC 1400 and 15000 Series Adapters and Supported QoS Changes.

Windows Server 2016 NENIC Version 5.8.25.9

- This driver update provides an RDMA driver for VIC 1400 and 15000 Series Adapters and Supported QoS Changes.

Windows Server 2022, 2019 and 2016 ENIC Version 4.4.0.12

- This driver update provides a Spectre-compliant driver for VIC 1300 Series adapters.

Windows 2022, 2019 and 2016 FNIC Driver Updates**Windows Server 2022, 2019 and 2016 FNIC Version 3.3.0.11**

- This driver update provides a Spectre-compliant fNIC driver for VIC 15XXX, 14XX and VIC 13XX adapters.

Resolved Caveats

The following table lists the resolved caveats in Release 4.2.

There are no resolved caveats in Release 4.2(1d).

Defect ID	Description	First Bundle Affected	Resolved In
CSCvq02558	The VIC 1400 Series Windows drivers on Cisco UCS B-Series and C-Series servers could not support more than 2 RDMA engines per adapter. Windows could only support RDMA on 4 vPorts on each RDMA Engine. You can Enable RDMA with the PS command on more than 4 vPorts on each RDMA Engine, but the driver would not allocate RDMA resources to more than 4 vPorts per engine. Executing a Get-NetAdapterRdma command on the host could show additional vPorts with RDMA Capable Flag as True. Using the Get-smbclientNetworkInterface command shows the actual number of RDMA vPort resources available for use. This issue is resolved.	4.0(3.51)B and C	4.2(1i)B and C
CSCvy11532	The Windows neNIC Driver failed to load (Yellow Bang) on VIC 14XX Series adapter on Cisco C245 M6 (AMD Based) Rack Servers with SMT / X2APIC features enabled. This issue is resolved.	4.2(0.232)C	4.2(1d)
CSCvx37120	When no BIOS Policy was used in the Service Profile for Cisco UCS M6 servers, the "\$" sign appeared in CDN Names for network interfaces in OS. This issue is resolved.	4.2(1a)A	4.2(1i)A

Defect ID	Description	First Bundle Affected	Resolved In
CSCvy75588	Call trace was seen on RHEL 8.4 when fc-nvme name space was not configured.	VIC FW 5.2(1a) Driver version 2.0.0.72-189.0	4.2(2a)A
CSCvz51592	SLES 15.3 intermittently crashed during sanboot with inbox driver.	Inbox fnic 1.6.0.53 Unfied fnic 2.0.0.74-198.0	4.2(2a)A
CSCwa67341	NENIC warning message with Event ID 10 in the windows Event Log. When the warning is posted the QoS on this adapter is disabled.	4.2(1.147)C	4.2(2a)A

Open Caveats

The following table lists the open caveats.

Defect ID	Description	Workaround	First Release Affected
CSCwh50478	Microsoft Windows 2022 OS resulted in bugcheck 0x50 when the interrupt count is configured to a value greater than 256	<p>Perform the following steps to avoid the OS crash with <code>bugcheck 0x50</code> on Microsoft Windows 2019 or Windows 2022:</p> <ol style="list-style-type: none"> Go to device manager. Select the appropriate network adapters. Record the driver key. Go to registry editor. Select <code>Computer\HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\Class\above driver key</code> Add new dword "RSSv2Disable" with hex value 1. Reboot Microsoft Windows OS. 	4.3(2c)

Defect ID	Description	Workaround	First Release Affected
CSCvy16861	In a Windows Hyper-V environment the VMQ feature is enabled. Event ID 113 is logged in the system event viewer when VMs are powered ON.	It has been determined that this issue does not have any functional or performance impact on the functioning of the VMQ feature. This issue will be investigated in a future release.	4.2(0.193)B
CSCvv76888	On a Cisco VIC 1300 Series adapters, using neNIC Driver version 4.3.0.6 and a 1300 Series adapter with a VMQ policy, a yellow bang appears when configured with a VMQ sub-vNIC value of 10 or less.	When a VMQ policy is created, ensure that there are at least 32 interrupts, even though the number of VMQs in the policy is lower. This will enable the driver to load and function correctly.	4.1(2.13)B
CSCvx81384	In a UCS Manager service profile where vHBAs are assigned with an FC adapter policy that has more than one I/O Queues, BSOD will be observed after loading the fNIC driver on Windows 2019. The issue is observed on VIC 1400 Series adapters on SAN and Local boot. The server showed BSOD or the below error: Stop code: PAGE FAULT IN NONPAGED AREA	Modify the FC adapter policy and set the I/O Queues to 1.	2.4(08)
CSCvr63930	On ESXi, a Cisco UCS B-series blade server and Cisco UCS C-series rack server with a a Cisco VIC 1440, VIC 1480,,1455, 1457, or 1467 adapter, the port link speed output is not updating after an uplink is Down /Up.	<ul style="list-style-type: none"> • Reboot the ESXi host to get the correct port speed reporting. • If a new UIF link needs to be added or deleted to the existing system, a reboot is required to reflect the correct speed in ESXi. 	3.1(1.152)B 2.1(2.56)A

Defect ID	Description	Workaround	First Release Affected
CSCvt66474	<p>On Cisco VIC 1400 Series adapters, the neNIC driver for Windows 2019 can be installed on Windows 2016 and the Windows 2016 driver can be installed on Windows 2019. However, this is an unsupported configuration.</p> <p>If the Windows 2019 neNIC driver is installed on Windows 2016, RDMA is not supported.</p> <p>If the Windows 2016 neNIC driver is installed on Windows 2019, the RDMA feature that is supposed to be enabled on Windows 2019 will be disabled.</p>	The driver binaries for WS2016 and WS2019 are in folders that are named accordingly. Install the right binary on the platform that is being built or upgraded.	4.1(1.49)C
CSCvz57245	On a B200 M6 blade server with UCS VIC 14425 adapter configured for SAN boot and with 4 vHBAs, LUNs go offline when one of the controller nodes is down or stuck and multiple reboots have occurred.	<p>Perform the following steps to bring the LUN back online:</p> <ol style="list-style-type: none"> 1. On Windows Server, go to Disk Management, 2. Right click on the disk that has gone offline and click on Online. This should return the LUN to service. 	4.2(1a)A
CSCwa93556	On M5 Blade and Rack servers with VIC 1440 and 1480 adapters, ESXi OS installation fails with FC boot when the adapter policy is set to INTX mode.	No workaround.	4.2(1.151)A
CSCwb79770	On a UCS C3260 standalone server with a VIC 15000 Series adapter, Vport connectivity fails when 16k RX RingSize is configured during Initial Configuration. This issue happens ONLY when the RX ring side is configured to values above 4K when first setting up the initial configuration. Once the host is rebooted or the interface is enabled or disabled, the issue disappears.	Disable, then re-enable the interface from Power Shell in Hyper-V.	3.3(0.11)A
CSCwa56085	A system assertion occurred on a VIC 1400 Series adapter while TCP traffic was running on the enic interfaces during the scan for hardware changes.	No workaround.	3.0(0.1)A

Behavior Changes and Known Limitations

Virtual Machine Multi Queue (14xx and 15xxx vNICs)

Disabling VMMQ state for vport is not taking effect.

Workaround

Use the power shell command to disable and assign queue to 1.

```
#Get-VMNetworkAdapter -vmname * | Set-VMNetworkAdapter -VmmqEnabled $false -VmmqQueuePairs 1
```

Cisco UCS VIC adapters with Cisco UCS VIC firmware version 4.1(2b) and later do not support Third Party Transceivers

Cisco UCS VIC adapters with Cisco UCS VIC firmware version 4.1(2b) and later do not support third party transceivers.

Use Cisco qualified transceivers or cabling for the physical links after 4.1(2b).

When LUNs per Target is set to more than 1024 in FC adapter policy of vHBAs, but the actual value deployed in FC vNIC is capped to 1024.

In Cisco UCS Manager 4.2(3c) release or later, when **LUNs Per Target** field is set to more than 1024 in FC adapter policy of vHBAs of Service Profile, the actual value deployed in FC vNIC is capped to 1024.

This issue occurs because the firmware version on the VIC adapter is old and does not support more than 1024 value for **LUNs Per Target**.

RHEL 8.7 boots to emergency shell when LUNs Per Target is set to greater than 1024

If **LUNs Per Target** is set to greater than 1024 with multiple paths running RHEL 8.7, the OS takes a long time to scan all the paths. Eventually, the scan fails and the OS boots to the emergency shell.

Reduce the number of **LUNs Per Target** (paths) to be scanned by the OS.

System Crashes When the SFP Module is Hot Swapped with the VIC Management Driver Installed

On UCS C220 M5 servers, when the SFP module is hot swapped on VIC 1495 or VIC 1497 adapters, a Blue Screen of Death (BSOD) appears and the system reboots. This happens only with the VIC management driver on Microsoft Windows.

Support for VIC Management Interface Driver for Windows 2016 and later is deprecated

Installing the driver iso containing VIC Management Driver on Windows 2016 or later might lead to system BSOD in certain scenarios.

Perform the following steps to uninstall the VIC management driver and avoid any potential issues:

1. In the Windows host, go to the **Device Manager** and select the appropriate Cisco VIC Internet Interface.
2. Right-click the appropriate Cisco VIC Internet Interface and select the **Uninstall Device** option.

A new window is displayed.

3. On the **Uninstall Device** window, select the **Delete the driver software for the device** checkbox to uninstall and delete the interface from the system device.
The Cisco VIC management interface is uninstalled and deleted.
4. Scan the hardware changes and ensure that the uninstalled VIC management interface is shown under the other devices as **PCI Device**.
You might get Yellow Bang error in Windows Device Manager, indicating that the driver is not loaded.



Note If the BSOD error persists, then reboot the operating system in the safe mode to uninstall the VIC management driver by performing the above steps.

Downloading the driver iso that does not contain VIC Management Driver causes Yellow Bang in Windows Device Manager, indicating that the driver is not loaded. This is expected and there is no functional impact.

Q-in-Q Forwarding (14xx and 15xxx VNICs)

For double tagged frames (1Q + 1Q) generated by the host are sent out by the VICs, you must configure the following commands on the Linux host.

1. Disable VLAN TX offload on the 14xx or 15xxx VNICs that need to transmit out double tagged (1Q + 1Q) frames.

Perform this from the host and enter the following `ethtool` command:

```
ethtool -K <interface_name> txvlan off
```

2. To verify that the VLAN TX offload feature has been turned off, enter the following command:

```
ethtool -k <interface_name> | grep tx-vlan-offload
```

Windows : Default adapter policy win-HPN-SMBd to be changed to 512+ for large logical processors value

Modify the interrupt value to 514 and re-deploy the updated setting.

Support for Physical NIC Mode

Beginning from release 4.2(3b), physical NIC mode is supported completely and the term **Experimental** is removed from Physical NIC mode for Cisco UCS C-Series Rack Servers.

Link Speed on ESXCLI is not updated at Runtime after Link Down/UP

This issue occurs when VMware API is not updating the Link status to Driver.

To avoid this, run the following command at FI or Up Link switch:

```
sh interface port-channel (uplink Po)
```

vNIC MTU Configuration

MTU on VIC 1400 Series adapters in Windows is now derived from the Jumbo Packet advanced property rather than from the UCS configuration

For VIC 14xx adapters, you can change the MTU size of the vNIC from the host interface settings. The new value must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets could be dropped during data transmission.

RDMA Limitations

- The VIC 1400 Series Windows drivers on Blade and Rack servers do not support more than 2 RDMA engines per adaptor. Currently, Windows can only support RDMA on 4 VPorts on each RDMA Engine.
- RoCE version 1 is not supported with any fourth generation Cisco UCS VIC 1400 Series adapters.
- UCS Manager does not support fabric failover for vNICs with RoCEv2 enabled.
- RoCEv2 cannot be used on the same vNIC interface as NVGRE, NetFlow, and VMQ features.
- RoCEv2 cannot be used with usNIC.
- RoCEv2 cannot be used with GENEVE offload.

Configuration Fails When 16 vHBAs are Configured with Maximum I/O Queues

Cisco UCS Manager supports a maximum of 64 I/O Queues for each vHBA. However, when you configure 16 vHBAs, the maximum number of I/O Queues supported for each vHBA becomes 59. In Cisco UCS Manager Release 4.0(2), if you try to configure 16 vHBAs with more than 59 I/O queues per vHBA, the configuration fails.

VM-FEX

ESX VM-FEX and Windows VM-FEX are no longer supported.

Auto-negotiation

When a `palo_get_an_status mptool` command is issued, it now shows that auto-negotiation is turned on all the time.

Link Training

The Link Training option is not configurable from CIMC for VIC 13xx adapters.

INTx Interrupt Mode

INTx interrupt mode is not supported with the ESX nenic driver and nfnic driver.

FC-NVMe Failover

To protect against host and network failures, you must zone multiple initiators to both of the active controller ports. Passive paths will only become active if controller fails, and will not initiate a port flap. On operating systems based on older kernels that do not support ANA, dm multi-path will not handle the passive paths correctly and could send IOs to a passive path. These IO operations will fail.

FC-NVMe Namespaces

Starting with RHEL 8.5 nvme-cli version 1.14, the nvme list command will not display fc-nvme namespaces. Use nvme-cli from RHEL 8.4 or a nvme-cli version of 1.15 or later to view fc-nvme namespaces.

FC-NVMe ESX Configurations

VIC 15000 and 1400 Series adapters running ESXi currently only support a maximum FC-NVMe namespace block size 512B, while some vendors use a default 4KB block size for ESXi 7.0. The target FC-NVMe namespace block must therefore be specifically configured to 512B. Under Storage go to NVMe and change the Block Size in NVMe from 4KB to 512B.

Configuration changes are also required to avoid a decrease in I/O throughput and/or BUS BUSY errors, caused by a mismatch between FC-NVMe Target controller Queue-depth and VM Device Queue-depth. To avoid this, run the following command to display all controllers discovered from ESXi Host:

```
# Esxcli nvme controller list
```

Check the list of controller queues and queue size for the controllers:

```
# vsish -e get /vmkModules/vmknvme/controllers/controller number/info
```

All Controllers on the same target will support same queue size, for exaple:

```
Number of Queues:4
Queue Size:32
```

To tune the VMs, change the `queue_depth` of all NVMe devices on the VMs to match the controller Queue Size. For example, if you are running a RHE VM, enter the command:

```
# Echo 32 > /sys/block/sdb/ device/queue_depth
```

Verify that the `queue_depth` was set to 32 by running the command:

```
# cat /sys/block/sdb/ device/queue_depth
```



Note This change is not persistent after reboot.



Note For additional driver configuration, it may be necessary to set the Adapter Policy to **FCNVMeInitiator** to create a FC-NVMe Adapter.

Adapter Policy can be found under **Server > Service Profile>Policies> Adapter Policies Create FC Adapter Policy>**

Adapter Policy can be found under **Server > Service Profile> Storage> Modify vHBAs**

Enabling FC-NVMe with ANA on ESXi 7.0

In ESXi 7.0, ANA is not enabled for FC-NVMe. This can cause failure of the Target side Path failover.

For a procedure to enable ANA, go to the following URL: https://docs.netapp.com/us-en/ontap-sanhost/nvme_esxi_7.html#validating-nvmefc

Related Cisco UCS Documentation

Documentation Roadmaps

For a complete list of all B-Series documentation, see the *Cisco UCS B-Series Servers Documentation Roadmap* available at the following URL: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/UCS_roadmap.html

For a complete list of all C-Series documentation, see the *Cisco UCS C-Series Servers Documentation Roadmap* available at the following URL: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/overview/guide/ucs_rack_roadmap.html.

For information on supported firmware versions and supported UCS Manager versions for the rack servers that are integrated with the UCS Manager for management, refer to [Release Bundle Contents for Cisco UCS Software](#).

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly [What's New in Cisco Product Documentation](#), which also lists all new and revised Cisco technical documentation.

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Follow [Cisco UCS Docs on Twitter](#) to receive document update notifications.