



# Release Notes for Cisco UCS Virtual Interface Card Drivers, Release 4.1

**First Published:** 2020-02-20

**Last Modified:** 2021-04-05

## Release Notes for Cisco UCS Virtual Interface Card Drivers, Release 4.1

### Introduction

This document contains information on new features, resolved caveats, open caveats, and workarounds for Cisco UCS Virtual Interface Card (VIC) Drivers, Release 4.1 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
February 21, 2020	Initial release of VIC drivers for Cisco UCS Software Release 4.1(1).
April 28, 2020	Added CSCvs60000 to the list of Open Caveats and Resolved Caveats for Release 4.1(1).
July 30, 2020	Release of VIC drivers for Cisco UCS Software Release 4.1(2).
October 21, 2020	Added INTx Interrupt Mode to Behavior Changes and Known Limitations.
January 12, 2021	Release of VIC drivers for Cisco UCS Software Release 4.1(3).
April 5, 2021	FC-NVMe ESX Configurations and Enabling FC-NVMe with ANA on ESXi 7.0 to Behavior Changes and Known Limitations.

### New Software Features in Release 4.1(1a)

**Release 4.1(1a) adds support for the following:**

- Support for Cisco UCS 64108 Fabric Interconnects that support 96 10/25-Gbps ports, 16 10/25-Gbps unified ports, and 12 40/100-Gbps uplink ports.

- NVMe over Fabrics (NVMeoF) using RDMA over Converged Ethernet version 2 (RoCEv2) on Redhat Enterprise Linux 7.6 with Linux Z-Kernel 3.10.0-957.27.2, for Cisco 14xx Series adapters.
- Support for NVMe over Fibre Channel (FC-NVMe) on SLES12 SP4, SLES15, SLES15 SP1, and RHEL 7.6. This support is available on UCS 6300 series Fabric Interconnects, UCS 6454, and UCS 64108 Fabric Interconnects with Cisco UCS 14xx series adapters. This support is also available on Cisco C220 and C240 M5 Standalone rack servers with Cisco UCS 14xx series adapters.
- RDMA Over Converged Ethernet (RoCE) Version 2 Support with Cisco UCS VIC 1400 Series adapters. It also adds support for Microsoft SMB Direct with RoCEv2 on Microsoft Windows 2019. Refer [UCS Hardware and Software Compatibility](#) for more details about support of Microsoft SMB Direct with RoCEv2 on Microsoft Windows 2019.




---

**Note** Windows RDMA is enabled as a tech preview feature and is disabled by default.

---

- Support for FDMI on Unified fNIC Linux drivers on Red Hat Enterprise Linux 7.6/ 7.7 and SLES 12 SP3, and SLES 12 SP4.
- Support for Red Hat Enterprise Linux 8.1 and on SLES 12.5 multi-queue on Unified fNIC drivers

#### Unsupported Features:

The following features are no longer supported:

- Beginning with Cisco UCS Manager Release 4.1(1), VMware VM-FEX and Hyper-V VM-FEX are no longer supported

#### Behavior Changes on Windows Server 2016 and 2019

For RDMA in 1300 Series versions of the VIC driver, the MTU was derived from either a UCS Manager profile or from Cisco IMC in standalone mode. With VIC 1400 Series adapters, MTU is controlled by the Windows OS Jumbo Packet advanced property. Values derived from UCS Manager and Cisco IMC have no effect.

## New Software Features in Release 4.1(2a)

#### Release 4.1(2a) adds support for the following:

- NVMe over Fabrics (NVMeoF) using IPv4 or IPv6 RDMA over Converged Ethernet version 2 (RoCEv2) is supported on Red Hat Enterprise Linux 7.7 with Linux Z-Kernel-3.10.0-1062.9.1.el7.x86\_64
- RoCEv2 protocol for Windows 2019 NDKPI Mode 1 and Mode 2 with IPV4 and IPV6
- Support for NVMe over Fibre Channel on Red Hat Enterprise Linux 7.7, 8.0 and 8.1
- Generic Network Virtualization Encapsulation (GENEVE Offload) supports on ESX 6.7U3 and ESX 7.0 operating systems
- Support for Red Hat Enterprise Linux 8.2 multi-queue on Unified fNIC drivers
- usNIC is now supported on the Cisco UCS C125 M5 server

- Support for FDMI on Unified fNIC Linux drivers on Red Hat Enterprise Linux 7.8/8.0/ 8.1, SLES 15, and SLES 15 SP1.

## New Software Features in Release 4.1(3a)

### Release 4.1(3a) adds support for the following:

- Support for NVMe over Fibre Channel (FC-NVMe) on UCS 6300 series Fabric Interconnects, UCS 6454, and UCS 64108 Fabric Interconnects with Cisco UCS VIC 13xx series adapters on RHEL 7.8, RHEL 7.9, and RHEL 8.2. This support is also available on Cisco C220 and C240 M5 Standalone rack servers with Cisco UCS 13xx series adapters.
- Support for NVMe over Fibre Channel with Cisco UCS 1400 series adapters on RHEL 7.8, RHEL 7.9, RHEL 8.2, and ESXi 7.0.
- Support for NVMe over Fabrics (NVMeoF) using IPv4 or IPv6 RDMA over Converged Ethernet version 2 (RoCEv2) on Red Hat Enterprise Linux 7.8 and 8.2.
- Support for fNIC Multi-Queue on RHEL 7.6, RHEL 7.7, RHEL 7.9, RHEL 8.0, RHEL 8.1, RHEL 8.2, RHEL 8.3, SLES 12 SP5, and SLES15 SP2.
- Support for Enhanced Datapath (ENS) driver on ESX 6.7U3, ESX 7.0, and ESX 7.0U1.
- Generic Network Virtualization Encapsulation (GENEVE Offload) is supported on ESX 7.0U1 Operating system.
- FDMI support on Red Hat Enterprise Linux 7.9/8.2 and SLES 15 SP 2.

## VIC Driver Updates for Release 4.1(1a)

### ESX ENIC Driver Updates

Native ENIC driver versions 1.0.X.0 are for ESXi 6.5 and later releases.

### ESXi NENIC Version 1.0.31.0

Native NENIC driver version 1.0.31.0 is supported with 6.5U3 and 6.7U2/U3.

### ESX FNIC Driver Updates

### Native FNIC Version 4.0.0.48

Native FNIC driver version 4.0.0.48 is supported on ESXi 6.7U2/U3.

## Linux Driver Updates



**Note** VIC drivers for all operating systems listed in the HCL are cryptographically signed by Cisco. However, some older Linux distribution versions such as Red Hat Enterprise Linux (RHEL) 6.x and CentOS 6.x do not support UEFI Secure Boot. Cisco signs the drivers for these distribution versions in case manual verification is desired. This applies to both ENIC and FNIC drivers.



**Note** If UCS servers are not booted with UEFI Secure Boot enabled, Cisco's cryptographic certificates are not loaded in to the Linux kernel keychain. When the Linux kernel loads a Cisco binary driver that was downloaded from [cisco.com](http://cisco.com), it will therefore be unable to verify the authenticity of the driver's cryptographic signature. You may see a warning like this:

```
Request for unknown module key 'Cisco UCS Driver Signing REL Cert: ...'
```

However, since the server was not booted with Secure Boot, the Linux kernel will still load the driver, even though it was unable to validate the driver's cryptographic signature; the message is only a warning.

### ENIC Version 802.21

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 8.0, 8.1
- XenServer 7.1, 7.6, 8.0, 8.1
- SLES 12 SP4, SLES 12 SP5, and SLES 15, SLES 15 SP1
- Ubuntu Server 16.04.5, 16.04.6, 18.04.1, 18.04.2, and 18.04.3
- CentOS 7.6 7.7, and 8.0

### Linux FNIC Driver Updates

#### Unified FNIC Driver Update 2.0.0.59

This driver update contains FC-NVMe support for VIC 14xx drivers on the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.6
- SLES 12.4, SLES 15, SLES 15.1

This driver update adds FDMI driver support for VIC 14xx drivers on the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 8.0, and 8.1
- XenServer 8.0, 8.1
- SLES 12 SP4, SLES 12 SP5, and SLES 15, SLES 15 SP1
- CentOS 7.7, 8.0



---

**Note** fNIC multiqueue is supported on RHEL 7.6, 7.7, 8.0, and 8.1 and SLES 12.5.

---

#### **Non-Unified FNIC Driver Update 1.6.0.51**

- CentOS 7.6
- XenServer 7.1, 7.6



---

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

---



---

**Note** CSCvk34443—After installing the Cisco fNIC driver on a system with SLES 15, the following error message appears:

```
cat: write error: broken pipe
```

The driver is installed correctly, and is operational. The message that appears is an informational message from SUSE, and is not caused by the Cisco fNIC driver.

---

#### **Windows 2019 and 2016 NENIC Driver Updates**

##### **Windows Server 2019 NENIC Version 5.5.22.3**

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

##### **Windows Server 2019 and 2016 ENIC Version 4.2.0.5**

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

#### **Windows 2019 and 2016 FNIC Driver Updates**

##### **Windows Server 2019 and 2016 FNIC Version 3.2.0.14**

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

## **VIC Driver Updates for Release 4.1(2a)**

#### **ESX ENIC Driver Updates**

Native ENIC driver versions 1.0.X.0 are for ESXi 6.5 and later releases.

##### **ESXi NENIC Version 1.0.33.0**

Native NENIC driver version 1.0.33.0 is supported with ESXi6.5U3, ESXi6.7U2/U3 and ESXi 7.0.

### ESX FNIC Driver Updates

#### ESXi FNIC Version 1.6.0.52

ESXi FNIC Version 1.6.0.52 is supported for ESXi6.5U3.

#### Native FNIC Version 4.0.0.56

Native FNIC driver version 4.0.0.56 is supported on ESXi6.7U2/U3 and ESXi7.0.

### Linux Driver Updates

#### ENIC Version 802.43

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 7.8, 8.0, 8.1, and 8.2
- XenServer 7.1, 7.6, 8.0, 8.1
- SLES 12 SP4, SLES 12 SP5, SLES 15, and SLES 15 SP1
- Ubuntu Server 16.04.5, 16.04.6, 18.04.1, 18.04.2, 18.04.3, 18.04.4. and 20.04
- CentOS 7.6 7.7, 7.8, 8.0, and 8.1

#### ENIC Version 802.44

This driver supports the following Operation System versions:

- XenServer 8.2
- CentOS 8.2

### Linux FNIC Driver Updates

#### Unified FNIC Driver Update 2.0.0.63

This driver supports the following Operation System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 7.8, 8.0, 8.1, 8.2
- CentOS 7.7, 7.8, 8.0, 8.1, and 8.2
- SLES 12 SP4, SLES 12 SP5, SLES 15, and SLES 15 SP1
- XenServer 8.0, 8.1, and 8.2

#### Non-Unified FNIC Driver Update 1.6.0.51

- CentOS 7.6
- XenServer 7.1, 7.6



---

**Note** FC-NVMe is supported on RHEL 7.6, 7.7, 8.0, and 8.1 with Cisco VIC 14xx adapters.  
FC-NVMe is supported on SLES 15 with DM multi-pathing only and native multi-pathing is not supported.  
FNIC Multi-Queue is supported on RHEL 7.6, 7.7, 8.0, 8.1, 8.2, and SLES 12 SP5.

---



**Note** CSCvk34443—After installing the Cisco fNIC driver on a system with SLES 15, the following error message appears:

```
cat: write error: broken pipe
```

The driver is installed correctly, and is operational. The message that appears is an informational message from SUSE, and is not caused by the Cisco fNIC driver.

---

### **Windows 2019 and 2016 NENIC Driver Updates**

#### **Windows server 2016 NENIC version 5.5.22.3**

- Because RDMA is not supported on Windows 2016, this driver does not support using RDMA. It is provided for Legacy support purposes.

#### **Windows Server 2019 NENIC Version 5.5.22.3**

- This driver is for VIC 14xx adapters. RDMA is supported on Windows 2019.

#### **Windows Server 2019 and 2016 ENIC Version 4.2.0.5**

- This driver supports VIC 13xx and earlier adapters.

### **Windows 2019 and 2016 FNIC Driver Updates**

#### **Windows Server 2019 and 2016 FNIC Version 3.2.0.14**

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

## **VIC Driver Updates for Release 4.1(3a)**

### **ESX ENIC Driver Updates**

Native ENIC driver versions 1.0.X.0 are for ESXi 6.5 and later releases.

#### **ESXi NENIC Version 1.0.35.0**

Native NENIC driver version 1.0.35.0 is supported with ESXi6.5U3, ESXi6.7U2/U3, ESXi7.0, and ESXi7.0U1.

### **ESX FNIC Driver Updates**

#### **ESX FNIC Version 1.6.0.52**

ESX FNIC Version 1.6.0.52 is supported for ESXi6.5 and ESXi6.5U3.

#### **Native FNIC Version 4.0.0.65**

Native FNIC driver version 4.0.0.65 is supported on ESXi6.7U2/U3 and ESXi7.0, and ESXi7.0U1.

#### **Asynchronous Native FNIC Version 5.0.0.11**

Asynchronous Native FNIC driver version 5.0.0.11 is supported on ESXi7.0, and ESXi7.0U1.

**ESX NENIC\_ENS Version 1.0.2.0**

NENIC\_ENS version 1.0.2.0 is the initial ENS driver release supported on ESXi 6.7U3.

**ESX NENIC\_ENS Version 1.0.4.0**

NENIC\_ENS version 1.0.4.0 is the initial ENS driver release and is supported on ESXi7.0, and ESXi7.0U1.

**Linux Driver Updates****ENIC Version 802.74**

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 7.8, 7.9, 8.0, 8.1, 8.2 and 8.3
- XenServer 7.1, 8.1, and 8.2
- SLES 12 SP4, SLES 12 SP5, SLES 15, SLES 15 SP1 and SLES 15 SP2
- Ubuntu Server 16.04.5, 16.04.6, 16.04.7, 18.04.1, 18.04.2, 18.04.3, 18.04.4, 18.04.5, 20.04 and 20.04.1
- CentOS 7.6 7.7, 7.8, 8.0, 8.1, and 8.2

**Linux FNIC Driver Updates****Unified FNIC Driver Update 2.0.0.69**

This driver supports the following Operation System versions:

- Red Hat Enterprise Linux 7.6, 7.7, 7.8, 7.9, 8.0, 8.1, 8.2, and 8.3
- CentOS 7.7, 7.8, 8.0, 8.1, and 8.2
- SLES 12 SP4, SLES 12 SP5, SLES 15, SLES 15 SP1 and SLES 15 SP2
- XenServer 8.1, and 8.2

**Non-Unified FNIC Driver Update 1.6.0.51**

- CentOS 7.6
- XenServer 7.1

**Note**

---

FC-NVMe is supported on SLES12 SP4 with Cisco VIC 14xx adapters.

FC-NVMe is supported on SLES 15 SP1 with DM multi-pathing only and native multi-pathing is not supported.

FNIC Multi-Queue is supported on RHEL 7.6, 7.7, 8.0, 8.1, 8.2, and SLES 12 SP5.

---





**Note** CSCvk34443—After installing the Cisco fNIC driver on a system with SLES 15, the following error message appears:

```
cat: write error: broken pipe
```

The driver is installed correctly, and is operational. The message that appears is an informational message from SUSE, and is not caused by the Cisco fNIC driver.

### Windows 2019 and 2016 NENIC Driver Updates

#### Windows server 2016 NENIC version 5.6.30.3

- Because RDMA is not supported on Windows 2016, this driver does not support using RDMA. It is provided for Legacy support purposes.

#### Windows Server 2019 NENIC Version 5.6.30.3

- This driver is for VIC 14xx adapters. RDMA is supported on Windows 2019.

#### Windows Server 2019 and 2016 ENIC Version 4.3.7.4

- This driver supports VIC 13xx and earlier adapters.

### Windows 2019 and 2016 FNIC Driver Updates

#### Windows Server 2019 and 2016 FNIC Version 3.2.0.14

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

## Resolved Caveats

The following table lists the resolved caveats in Release 4.1.

Defect ID	Description	FirstVersion Affected	Resolved In
CSCvs60000	Improper handling of a memory allocation failure in the eNIC driver could result in network connectivity issues.  For example, when the issue affected a TX queue, it lead to a WATCHDOG timeout, which triggered an interface reset. When the issue affected an RX queue, it lead to ingress packet drops (reflected in the ethtool counter <b>rx_no_buf</b> ).  This defect is now resolved.	4.0(4d)A	Resolved in 4.1(1a)A  4.1(1.40)A  neNIC 4.0.0.8
CSCvq50787	An OS reset occurred when servers with Linux fNIC drivers running UCS Manager were scanning the /sys/ directory.	4.0(4b)A	Resolved in 4.1(1a)A
CSCvn28299	On a SET switch created using two VMMQ capable NICs, network traffic sometimes stopped when removing and re-adding VMMQ capable NICs.	4.0(1.107)A	Resolved in 4.1(1a)A

Defect ID	Description	FirstVersion Affected	Resolved In
CSCvo18110	NVMe over Fibre Channel lost namespace paths during port flap.	4.0(0.9)A	Resolved in 4.1(1a)A
CSCvo00914	IO operations failed on NVMe over Fibre Channel namespaces during link flaps on the path between the UCS VIC server and Fibre Channel NVMe target.	4.0(0.9)A	Resolved in 4.1(1a)A
CSCvp21853	IO failure occurred on the FC-NVMe namespace.	2.0(0.37)A	Resolved in 4.1(1a)A
CSCvp35462	An eNIC running RHEL 7.6 failed boot into the OS when the server received invalid multicast packets and fails to boot into the OS.	4.0(3.94)A	Resolved in 4.1(1a)A
CSCvo36323	On a 13xx Series VIC card in a standalone C220 M5 Server, the rx_no_bufs counters on all connected vmNIC interfaces on ESXi 6.5/6.7 incremented whenever the neNIC was in use. This issue is now resolved.	2.1(2.22)	Resolved in 4.1(1a)A
CSCvr96728	An update to the VIC driver firmware caused rack servers with asynchronous connections to lose network connections and reboot.	4.0(2c)C	Resolved in 4.1(1a)A
CSCvt99638	Multiple storage errors were seen after QUEUE_FULL messages during fibre channel traffic. This issue is now resolved.	4.0(4g)	Resolved in 4.1(2a)A
CSCvu25233	On a 6400 Series Fabric Interconnect connected to VIC 1455/1457 adapter using SFP-H25G-CU3M or SFP-H25G-CU5M cables or on a VIC 1455/1457 adapter connected to 2232PP using a SFP-10GB-CUxM cable, link-flapping and links down occurred on some ports.	4.0(1a)A	Resolved in 4.1(2a)A
CSCvu87940	After importing a VNIC config file on a standalone C-series VIC adapter, when the host was rebooted, VNICs did not receive a link-up, resulting in loss of network to the host OS. This occurred when all of the following conditions are met: <ul style="list-style-type: none"> <li>• The user imported a VNIC configuration file that was exported when VIC was configured with <b>VNTAG mode</b> enabled.</li> <li>• VIC network ports are connected to Cisco Nexus switches supporting network interface virtualization.</li> <li>• The switch ports and/or portchannel are configured with <code>switchport mode vntag</code>.</li> </ul>	4.0(4h)C	Resolved in 4.1(2a)A

Defect ID	Description	FirstVersion Affected	Resolved In
CSCvs04971	When too many ports are present in a zone, LUNs are not discovered. Fixed by installing fNIC driver version 4.0.0.47 or above.	Driver version 4.0(0.9)	Resolved in 4.1(2a)A
CSCvs36209	On a VIC 1340 adapter nfNIC driver running ESXi 6.7u3, the nfNIC driver didn't reply to ADISC sent from IBM StorWise following a Zone Change. Fixed by installing fNIC driver version 4.0.0.56 or above.	4.0(4f)A	Resolved in 4.1(2a)A
CSCvt97063	On UCS servers connected to Fabric Interconnects with ESXi hosts using the Native fNIC driver (NFNIC), the following symptoms were observed immediately after system QoS changes, including enable/disable of non-default classes: <ul style="list-style-type: none"> <li>• Significant Abort errors were seen on both the host and VIC sysym logs.</li> <li>• ICMP/Ping loss to VMs were observed, but the ESXi host didn't experience an ICMP/ping issue.</li> </ul> This issue is now resolved.	4.0(4g)A	NFNIC driver version 4.0.0.56
CSCvv69526	When the vHBA on a VIC 1400 series adapter is disabled and rebooted, the vHBA remains disabled in link-down state.	4.1(2a)A	Resolved in 4.1(3a)A
CSCvi75867 CSCvv78557	Improper handling of SCSI Inquiry commands causes the fNIC driver to become unresponsive. When the failure is encountered, the system displays the error message: IRQL_NOT_LESS_OR_EQUAL (D1) BSOD.	Driver version 3.0.0.8	Resolved in 4.1(3a)A driver version 3.2.0.14
CSCvu84344	After a Hyper-V switch is created, VMQ queues added on a WS2016 server are not listed when the command <b>Get-NetAdapterVmqQueue</b> is run and there is no output. BSoD is seen when live migrating VMs.	enic6x64 4.0.0.3	enic6x64 4.3.7.3

## Open Caveats

The following table lists the open caveats in Release 4.1.

Defect ID	Description	Workaround	First Release Affected
CSCvq02558	<p>The VIC 1400 Series Windows drivers on B-Series and C-Series servers do not support more than 2 RDMA engines per adapter. Currently, Windows can only support RDMA on 4 VPorts on each RDMA Engine. Currently, you can Enable RDMA with the PS command on more than 4 vPorts on each RDMA Engine, but the driver will not allocate RDMA resources to more than 4 vPorts per engine. Executing a <b>Get-NetAdapterRdma</b> command on the host could show additional vPorts with RDMA Capable Flag as True. Using the <b>Get-smbclientNetworkInterfce</b> command shows the actual number of RDMA vPort resources available for use.</p>	Use the <b>Get-smbclientNetworkInterfce</b> command instead of the <b>Get-NetAdapterRdma</b> command to confirm the number of effective RDMA vPorts.	4.0(3.51)B and C
CSCvr67129	<p>An error occurs when system IOMMU is enabled and RDMA read response packets belonging to an already completed IO are retransmitted. The following error message appears in the host DMESG log, indicating a DMAR error happens when a VIC adapter tries to DMA from a host buffer.</p> <pre>DMAR: DRHD: handling fault status reg 2  DMAR: [DMA Read] Request device [62:00.1] fault addr xxxxxxxx [fault reason 06] PTE Read access is not set.</pre>	Disable IOMMU.	5.0(388)VC
CSCvp48149	<p>On Windows operating systems, yellow bang warning icons may appear for a VIC management device exposed for legacy reasons. There is no functional impact when this interface is exposed.</p>	None	3.1(1a)

Defect ID	Description	Workaround	First Release Affected
CSCvt66474	<p>On 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>Case 1 : Installing Windows 2019 neNIC driver on Windows 2016 succeeds...but on Windows 2016 RDMA is not supported.</p> <p>Case 2 : Installing Windows 2016 neNIC driver on Windows 2019 succeeds...but on Windows 2019 RDMA comes with default disabled state, instead of enabled state.</p>	The driver binaries for Windows 2016 and Windows 2019 are in folders that are named accordingly. Install the correct binary on the platform that is being built/upgraded.	4.1(1a)C
CSCvt99638	Multiple storage errors are seen after QUEUE_FULL messages.	Reboot the host.	4.0(4g)A Resolved in 4.1(2a)A
CSCvt97063	<p>On UCS servers connected to Fabric Interconnects with ESXi hosts using the Native fNIC driver (NFNIC), the following symptoms were observed immediately after system QoS changes, including enable/disable of non-default classes:</p> <ul style="list-style-type: none"> <li>• Significant Abort errors were seen on both the host and VIC sysym logs.</li> <li>• ICMP/Ping loss to VMs were observed, but the ESXi host didn't experience an ICMP/ping issue.</li> </ul> <p>This issue is now resolved.</p>	<p>Enable/disable vHBA on affected servers.</p> <p>Apply the fix in NFNIC driver 4.0.0.56 and later.</p>	4.0(4g)A
CSCvu25233	On a 6400 Series Fabric Interconnect connected to VIC 1455/1457 adapter using SFP-H25G-CU3M or SFP-H25G-CU5M cables or on a VIC 1455/1457 adapter connected to 2232PP using a SFP-10GB-CUxM cable, link-flapping and links down can occur on some ports.	<ul style="list-style-type: none"> <li>• Use 25G optical transceiver with optical cables or 25G AOC cables.</li> <li>• Use 10G optical cable (10G AOC cable or 10G optical transceiver with optical cable).</li> </ul>	4.0(1a)A Resolved in 4.1(2a)A

Defect ID	Description	Workaround	First Release Affected
CSCvu87940	<p>After importing a VNIC config file on a standalone C-series VIC adapter, when the host is rebooted, VNICs may not receive a link-up, resulting in loss of network to the host OS. This occurs when all of the following conditions are met:</p> <ul style="list-style-type: none"> <li>The user imported a VNIC configuration file that was exported when VIC was configured with <b>VNTAG mode</b> enabled.</li> <li>VIC network ports are connected to Cisco Nexus switches supporting network interface virtualization.</li> <li>The switch ports and/or portchannel are configured with <code>switchport mode vntag</code>.</li> </ul>	<p>Force the adapter to re-generate the UUID by disabling and re-enabling VNTAG mode for the adapter as follows:</p> <ol style="list-style-type: none"> <li>Go to the CIMC screen where the adapter is attached and go to the <b>General</b> tab for the adapter.</li> <li>De-select the <b>Enable VNTAG Mode</b> button and click <b>Save Changes</b>.</li> <li>Select the "Enable VNTAG Mode" button and click "Save Changes".</li> </ol>	4.0(4h)C Resolved in 4.1(2a)A
CSCvs04971	<p>When too many ports are present in a zone, LUNs are not discovered.</p> <p>Fixed by installing fNIC driver version 4.0.0.47 or above.</p>	Scale down the number of zone members around 120.	Driver version 4.0(0.9)
CSCvs36209	<p>On a VIC 1340 adapter nfNIC driver running ESXi 6.7u3, the nfNIC driver doesn't reply to ADISC sent from IBM StorWise following a zone change.</p> <p>Fixed by installing nfNIC driver version 4.0.0.56 or above.</p>	Suppressing the RSCN on the storage port prevents the other hosts from losing access, but this is only a short term workaround.	4.0(4f)A
CSCvv42176	Installation of the SLES 15.2 OS using inbox fNIC driver 1.6.0.47 fails to complete.	<p>To work around the issue, disable vHBA through UCS Manager, inject the async driver and proceed with the installation.</p> <p>Once installation is done, enable vHBA from UCSM.</p> <p>No workaround is available for booting with sanboot. Installation will fail.</p>	4.1(3.106)B
CSCvv22616	Flogi fails with vHBA interface if interrupt mode is configured to INT-x on the adapter policy.	None	5.1(2.23)VC

Defect ID	Description	Workaround	First Release Affected
CSCvw65797	On a server connected to an NVMe storage target through a VIC 1400 Series adapter and running RHEL 7.8, removing an enic_rdma module when the link is down might cause the server to become unresponsive.	Do not remove the enic_rdma module when the link is down.	1.0(0.0)
CSCvw83070	When Physical NIC mode is enabled and FIP disabled in a configuration file for a standalone S3260 VIC 1455 adapter, after importing the saved configuration file, FIP gets enabled at the VIC adapter.  VIC level configuration shows FIP is enabled and shows FIP enabled message logs at the VIC console.	Enable FIP using the standalone CIMC GUI and then disable FIP from the CIMC GUI. This will result in FIP being disabled at the VIC adapter.	4.1(3)520A

## Behavior Changes and Known Limitations

### 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 1440, 1480, 1495, 1497, 1455, 1457 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, & 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.

### **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.

### **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 Windows nenic driver. INTx interrupt mode is not supported when the enic driver has RoCEv2 enabled or IOMMU enabled.

MSI interrupt mode on Fibre Channel interfaces is not supported. If the user configures MSI interrupt mode for Fibre Channel interface, Fibre Channel interfaces will come up in MSix mode.

### **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 ESX Configurations**

VIC 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](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](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](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](#).

### Other Documentation Resources

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

## 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.