



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

**First Published:** 2023-03-06

**Last Modified:** 2025-04-30

## 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
April 2025	Initial release of VIC drivers for Cisco UCS Software Release 4.3(6a)
March 2025	Updated the <b>New Hardware</b> section - release of VIC drivers for Cisco UCS Software Release 4.3(5a).
October 2024	<ul style="list-style-type: none"><li>• Initial release of VIC drivers for Cisco UCS Software Release 4.3(5a).</li><li>• Updated the <b>Resolved Caveats</b> section - release of VIC drivers for Cisco UCS Software Release 4.3(4c).</li></ul>
June 2024	<ul style="list-style-type: none"><li>• Initial release of VIC drivers for Cisco UCS Software Release 4.3(4b).</li><li>• Initial release of VIC drivers for Cisco UCS Software Release 4.3(4a).</li></ul> <p>Initial release of VIC drivers for Cisco UCS Software Release 4.3(4x).</p>
January 2024	Initial release of VIC drivers for Cisco UCS Software Release 4.3(2d).
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).

Revision Date	Description
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(6)

**Release 4.3(6a) adds support for the following:**

#### Cisco UCS C220 M8 Server

The 1RU, 2-socket Cisco UCS C220 M8 server is designed to meet the needs of customers that choose to deploy high density rack-mount servers. It combines the latest Intel® processors and is a versatile general-purpose application and infrastructure server delivering leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications.

The Cisco UCS C220 M8 server extends the capabilities of the Cisco UCS rack server portfolio incorporating Intel Xeon® 6 CPUs. It improves security, performance, and efficiency while helping achieve sustainability goals with built-in accelerators the such as Intel Trust Domain Extensions (TDX), Intel Data Streaming Accelerator (DSA), Intel QuickAssist Technology (QAT), Intel Advanced Matrix Extensions (AMX), and Intel In-Memory Analytics Accelerator (IAA).

Operating Expenses (OpEx) for power and cooling, management, and maintenance can be decreased by consolidating older servers onto the latest generation of Cisco UCS C220 M8 Rack Servers.

The Cisco UCS C220 M8 server is a dense, fault-tolerant server that provides value, performance, and flexibility for both commercial and enterprise customers.

The Cisco UCS C220 M8 server offers the following:

- Processors: Up to 2x Intel Xeon 6700P or 6500P processors (1 or 2)
- Memory:
  - 32 DIMM slots (16 DIMMS per CPU): 16, 32, 48, 64, 96, 128GB DDR5 at up to 6400 MT/s for up to 4TB of memory
  - 32, 64GB MRDIMMs at up to 8000 MT/s
- PCIe expansion: Up to 3 PCIe 5.0 half-height slots or up to 2 PCIe 5.0 full-height slots plus 1 dedicated 24-Gbps RAID controller slot and 1 dedicated mLOM/OCF 3.0 slot
- RAID controllers:
  - Cisco® 24-Gbps modular tri-mode RAID controller supports SAS 4 or NVMe hardware RAID
  - Cisco 24-Gbps modular tri-mode SAS Host Bus Adapter (HBA)
- Internal storage:
  - Backplane options: Up to 10 SFF SAS/SATA/U.3 NVMe drives through SAS4 tri-mode RAID or HBA controller, with optional up to eight direct-attach U.2/U.3 NVMe drives
  - Up to 16 E3.S 1T direct-attach NVMe drives at PCIe Gen5 x4 each

- mLOM/OCP 3.0:
  - One dedicated PCIe Gen5x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity
  - mLOM slot can flexibly accommodate 10/25/50 and 40/100/200 Gbps Cisco VIC adapters
  - OCP 3.0 slot features full out-of-band manageability that supports Intel X710 OCP dual 10GBase-T through mLOM interposer
- Power supplies: Hot-pluggable, redundant platinum, and titanium options:
  - Platinum: 1050W DC and 1600W AC
  - Titanium: 1200W AC and 2300W AC
- Other storage:
  - Dedicated Baseboard Management Controller (BMC) FlexMMC for utilities (on board)
  - Dual M.2 SATA SSDs (internal or hot-swappable) with HW RAID support
- GPU: Up to three single-wide GPUs supported




---

**Note** Cisco UCS 220 M8 servers support only 15000 series secure boot adapters

---

### Cisco UCS C240 M8 Server

The 2RU, 2-socket Cisco UCS C240 M8 server offers I/O flexibility and larger storage capacity. It combines the fastest Intel® processors and is a versatile general-purpose application and infrastructure server delivering leading performance and efficiency for a wide range of workloads, including AI, big-data analytics, databases, collaboration, virtualization, and high-performance computing.

The Cisco UCS C240 M8 server extends the capabilities of the Cisco UCS rack server portfolio by incorporating Intel Xeon® 6 CPUs. It improves security, performance, and efficiency while helping you achieve your sustainability goals with built-in accelerators such as Intel Trust Domain Extensions (TDX), Intel Data Streaming Accelerator (DSA), Intel QuickAssist Technology (QAT), Intel Advanced Matrix Extensions (AMX), and Intel In-Memory Analytics Accelerator (IAA).

You are able to decrease server Operating Expenses (OpEx) for power and cooling, management, and maintenance by consolidating older servers onto the latest generation of Cisco UCS C240 M8 Server.

The Cisco UCS C240 M8 server offers the following:

- Processors: Up to 2x Intel Xeon 6700P or 6500P processors (1 or 2)
- Memory
  - 32 DIMM slots (16 DIMMS per CPU): 16, 32, 48, 64, 96, 128, 256GB DDR5 at up to 6400 MT/s for up to 8TB of memory
  - 32, 64GB MRDIMMs at up to 8000 MT/s
- PCIe expansion: Up to 8 PCIe 5.0 slots plus 1 dedicated 24-Gbps RAID controller slot and 1 dedicated mLOM/OCP 3.0 slot

- RAID controllers:
  - Cisco 24-Gbps modular tri-mode RAID controller supports SAS 4 or NVMe hardware RAID
  - Cisco 24-Gbps modular tri-mode SAS Host Bus Adapter (HBA)
- Internal storage:
  - Backplane options: Up to 28 SFF SAS/SATA/U.3 NVMe drives through SAS4 tri-mode RAID or HBA controller, with optional up to eight direct-attach U.2/U.3 NVMe drives
  - Up to 36 E3.S 1T direct-attach NVMe drives
  - Up to 16 LFF SAS HDDs plus optional 4 rear SFF HDD/SSDs
- mLOM/OCP 3.0: One dedicated PCIe Gen5x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity
 

mLOM slot can flexibly accommodate 10/25/50 and 40/100/200 Gbps Cisco VIC adapters.

OCP 3.0 slot features full out-of-band manageability that supports Intel X710 OCP Dual 10GBase-T via mLOM interposer
- Power supplies: Hot-pluggable, redundant platinum and titanium options:
  - Platinum: 1050W DC and 1600W AC
  - Titanium: 1200W AC and 2300W AC
- Other storage:
  - Dedicated Baseboard Management Controller (BMC) FlexMMC for utilities (on board)
  - Dual M.2 SATA SSDs (internal or hot-swappable) with HW RAID support
- GPU: Up to three double-wide or eight single-wide GPUs supported




---

**Note** Cisco UCS 240 M8 servers support only 15000 series secure boot adapters

---

### Cisco UCS X210c M8 Compute Node

Cisco UCS X210c M8 Compute Node is the third generation of compute node to integrate into the Cisco UCS X-Series Modular System. It delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads. Up to eight compute nodes can reside in the 7-Rack-Unit (7RU) Cisco UCS X9508 Chassis, offering one of the highest densities of compute, I/O, and storage per rack unit in the industry.

Cisco UCS X210c M8 provides these main features:

- CPU: up to two Intel® Xeon® 6 Scalable Processors with up to 86 cores per processor and up to 336MB of Level 3 cache per CPU
- Memory: up to 8 TB of main memory with 32x 256 GB DDR5-6400 DIMMs and support for MRDIMMs at up to 8000 MT/s

- Storage:
  - Up to nine hot-pluggable EDSFF E3.S NVMe drives with a new passthrough front mezzanine controller option new to the Cisco UCS X210c M8
  - Up to six hot-pluggable, Solid-State Drives (SSDs), or Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of enterprise-class Redundant Array of Independent Disks (RAIDs) or passthrough controllers with four lanes each of PCIe Gen 5 connectivity
  - Up to two M.2 SATA drives or two M.2 NVMe drives for flexible boot and local storage capabilities
- Optional front mezzanine GPU module: the Cisco UCS<sup>®</sup> front mezzanine GPU module is a passive PCIe Gen 4 front mezzanine option with support for up to two U.2 or U.3 NVMe drives and two HHHHL GPUs
- Optional PCIe node connectivity for additional GPU support: the Cisco UCS X210c M8 Compute Node can be paired with the Cisco UCS X440p PCIe Node supporting up to two x16 full-height, full-length dual slot GPUs, or four x8 full-height, full-length single-slot GPUs
- mLOM virtual interface cards:
  - Cisco UCS VIC (Virtual Interface Card) 15420 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 50 Gbps (2x 25Gbps) of unified fabric connectivity to each of the chassis's Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server with secure boot technology
  - Cisco UCS VIC 15230 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis's Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server with secure boot technology
- Optional mezzanine card:
  - Cisco UCS VIC 15422, a 5th Gen virtual interface card, can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric technology. An included bridge card extends this VIC's 4x 25 Gbps of network connections through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per server)
  - Cisco UCS PCI mezzanine card for Cisco UCS X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric modules and enable connectivity to the Cisco UCS X440p PCIe Node
  - All VIC mezzanine cards also provide I/O connections from the Cisco UCS X210c Compute Node to the Cisco UCS X440p PCIe Node
- Security: The server supports an optional Trusted Platform Module (TPM). Additional features include a secure boot FPGA and ACT2 anti-counterfeit provisions

## New Hardware in Release 4.3(5)

### Release 4.3(5a) adds support for the following:

The Cisco UCS X-Series Modular System simplifies your data center, adapting to the unpredictable needs of modern applications while also providing for traditional scale-out and enterprise workloads. It reduces the number of server types to maintain, helping to improve operational efficiency and agility as it helps reduce complexity. Powered by the Cisco Intersight™ cloud operations platform, it shifts your thinking from

administrative details to business outcomes with hybrid cloud infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized.

The Cisco UCS X215c M8 Compute Node is integrate into the Cisco UCS X-Series Modular System. Up to eight compute nodes can reside in the 7-Rack-Unit (7RU) Cisco UCS X9508 Chassis, offering one of the highest densities of compute, IO, and storage per rack unit in the industry.

The Cisco UCS X215c M8 Compute Node offers the following:

- CPU: Up to 2x 4th Generation AMD EPYC Processors with up to 128 cores per processors
- Memory:
  - 24 DIMM slots (12 DIMMs per CPU socket), up to 4800 MT/s DDR5.
  - Up to 6 TB of capacity.
- Storage: Up to 6 hot-pluggable, Solid-State Drives (SSDs), or Non-Volatile Memory Express (NVMe) 2.5-inch drives with a choice of enterprise-class Redundant Array of Independent Disks (RAID) or pass-through controllers with four lanes each of PCIe Gen 4 connectivity and up to 2 M.2 SATA or NVMe drives for flexible boot and local storage capabilities.
- Optional Front Mezzanine GPU module: The Cisco UCS Front Mezzanine GPU module is a passive PCIe Gen 4 front mezzanine option with support for up to two U.2 or U.3 NVMe drives and two HHHH GPUs.
- mLOM virtual interface cards:
  - Cisco UCS Virtual Interface Card (VIC) 15420 occupies the server's Modular LAN on Motherboard (mLOM) slot, enabling up to 50Gbps (2 x25Gbps) of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100Gbps connectivity per server.
  - Cisco UCS Virtual Interface Card (VIC) 15230 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server with secure boot capability.
- Optional Mezzanine card:
  - Cisco UCS Virtual Interface Card (VIC) 15422 can occupy the server's mezzanine slot at the bottom rear of the chassis. An included bridge card extends this VIC's 100Gbps (4 x 25Gbps) of network connections through IFM connectors, bringing the total bandwidth to 100Gbps per VIC 15420 and 15422 (for a total of 200Gbps per server). In addition to IFM connectivity, the VIC 15422 I/O connectors link to Cisco UCS X-Fabric technology.
  - Cisco UCS PCI Mezz card for X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric modules and enable connectivity to the X440p PCIe Node.
- Security: Includes secure boot silicon root of trust FPGA, ACT2 anti-counterfeit provisions, and optional Trusted Platform Model (TPM).

For complete list of supported peripherals for Cisco UCS X215c M8 Compute Node, see [Cisco UCS X215c M8 Compute Node Spec Sheet](#).



**Note** Cisco UCS X215c M8 Compute Node supports only 15000 Series secure boot VIC adapters.

## New Hardware in Release 4.3(4)

**Release 4.3(4b) adds support for the following:**

### Cisco UCS C245 M8 Server

The Cisco UCS C245 M8 server is perfectly suited for a wide range of storage and I/O-intensive applications such as big data analytics, databases, collaboration, virtualization, consolidation, AI/ML and high-performance computing supporting up to two AMD<sup>®</sup> CPUs in a 2RU form factor.

The Cisco UCS C245 M8 Server extends the capabilities of the Cisco UCS server portfolio. It powers 4<sup>th</sup> Gen AMD<sup>®</sup> EPYC<sup>™</sup> Processors with 100 percent more cores per socket designed using AMD's chiplet architecture. With advanced features like AMD Infinity Guard, compute-intensive applications will see significant performance improvements and will reap other benefits such as power and cost efficiencies.

You can deploy the Cisco UCS C-Series servers as part of the Cisco Unified Computing System<sup>™</sup> managed by Cisco Intersight<sup>®</sup> or Cisco UCS Manager<sup>®</sup> to take advantage of Cisco<sup>®</sup> standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility or as standalone servers.

The Cisco UCS C245 M8 Server brings many innovations to the UCS AMD<sup>®</sup> rack server line. With the introduction of PCIe Gen 5.0 expansion slots for high-speed I/O, a DDR5 memory bus, and expanded storage capabilities, the server delivers significant performance and efficiency gains that will greatly enhance application performance. Features include:

- Support for up to two 4<sup>th</sup> Gen AMD EPYC<sup>™</sup> CPUs in a server designed to drive as much as 256 CPU cores (128 cores per socket)
- Up to 24 DDR5 DIMM slots, yielding up to 6 TB of capacity, using 256 GB DIMMs (12 DIMMs per socket)
- Up to 4800 MT/s DDR5 memory
- Up to 8 x PCIe Gen 4.0 slots or up to 4 x PCIe Gen 5.0 slots, plus a hybrid modular LAN on motherboard (mLOM) /OCP 3.0 slot (details below)
- Support for Cisco UCS VIC 15000 Series secure boot adapters as well as a host of third-party NIC options
- Up to 28 hot-swappable small-form-factor (SFF) SAS/SATA or NVMe drives (with up to 8 direct-attach NVMe drives) and New tri-mode RAID controller supports SAS4 plus NVMe hardware RAID.
- M.2 boot options
  - Up to two 960GB SATA M.2 drives with hardware RAID support
  - Up to two 960GB NVMe M.2 drives with NVMe hardware RAID
- Support for up to Eight GPUs
- Modular LOM / OCP 3.0
  - One dedicated PCIe Gen4x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity

- mLOM slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot, supporting quad-port 10/25/50 Gbps or dual-port 40/100/200 Gbps network connectivity
- OCP 3.0 slot that features full out-of-band management for select adapters

Cisco IMC supports all the peripherals supported by Cisco UCS C245 M8 Server. For complete list of supported peripherals for Cisco UCS C245 M8 Server, see [Cisco UCS C245 M8 SFF Rack Server Spec Sheet](#).

## 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 Software Features in Release 4.3(6)

#### Release 4.3(6) adds support for the following:

- Support for PXE Boot in IPv6-only environments using UEFI mode for LAN Boot. This enhancement allows administrators to operate in modern network settings while minimizing the dependency on IPv4 addresses, ensuring efficient server provisioning and management.

### New Software Features in Release 4.3(5)

#### Release 4.3(5) adds support for the following:

FPIN Support in NPV mode: Cisco UCS VIC driver provides FPIN support in NPV mode on the ESXi server. This feature is a fabric notification feature for enhancing failover mechanisms between the FC switch and the target.

Using the FPIN support in NPV mode, FNIC can now inform the ESXi stack to perform failover in case of any protocol level errors or link issues between the FC switch and the target like loss of signal.

Using this feature, the VIC driver can also perform debugging, wherein the driver prints all the bit notification from the FC switch and TAC.

#### Operating System and Software Requirements

Below listed are the supported software required on the ESXi server to enable this feature:

- OS: ESX 8.0U3 and onwards
- FC Switch (MDS) version - 9.4
- NetApp Data ONTAP OS version - 9.13.1P2
- Pure Storage version 6.4
- Native FNIC Version 5.0.0.45-10EM.803.0.0.24022510

## New Features in Release 4.3(2)

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

- 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 4.3.6



---

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

---

#### ESX ENIC Driver Updates

##### ESX NENIC Version 2.0.17.0

NENIC/NENIC(RDMA) version 2.0.17.0 is supported with ESX 7.0 U1, ESX 7.0 U2, ESX 7.0 U3, ESX 8.0 and above.

##### ESX NENIC\_ENS Version 1.0.9.0

NENIC\_ENS version 1.0.9.0 is supported with ESX 7.0 U1, ESX 7.0 U2, ESX 7.0 U3, ESX 8.0 and above.

#### ESX FNIC Driver Updates

##### Native FNIC Version 5.0.0.45

Native FNIC driver version 5.0.0.45 is supported with ESX 8.0 and above.

##### Native FNIC Version 5.0.0.46

Native FNIC driver version 5.0.0.46 is supported with ESX 7.0 U1, ESX 7.0 U2 and ESX 7.0 U3.



---

**Note** Driver version 5.0.0.46 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 1128.x

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 8.8, 8.10, 9.2, 9.4, 9.5
- SUSE Linux Enterprise Server 12 SP5, 15 SP4, 15 SP5, 15 SP6
- Ubuntu Server 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, 22.04.3, 22.04.4, 24.04, 24.04.1 with 6.8.0-51 kernel

### ENIC Version 939.x

This driver supports the following Linux Operating System versions:

- -Citrix Hypervisor 8.4 LTSR

### Linux FNIC Driver Updates

#### Unified FNIC Driver 2.0.0.10x

This driver supports the following Linux Operating System versions:

- Red Hat Enterprise Linux 8.8, 8.10, 9.2, 9.4, 9.5
- SUSE Linux Enterprise Server 12 SP5, 15 SP4, 15 SP5, 15 SP6

#### Unified FNIC Driver 2.0.0.9x

This driver supports the following Linux Operating System versions:

- Citrix Hypervisor 8.4 LTSR



#### Note

- RDMA and ENS features are not supported with VIC 13xx adapters from the release 4.3(4a) and later releases.
- RDMA and ENS Features are supported with VIC 14xx adapters onwards.
- fnic multiqueue is supported on RHEL 8.4+, 9.0+ , SLES 12SP5+
- FC-NVME is supported for Cisco UCS VIC 14xx and VIC 15xxx adapters on RHEL 8.4+, RHEL 9.0+, SLES 15SP12+
- 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.

For the latest set of software and hardware, check the support matrix: <https://www.cisco.com/c/en/us/products/servers-unified-computing/interoperability.html>.

### Windows 2025, 2022 and 2019 NENIC/ENIC Driver Updates

#### Windows Server 2025 and 2022 NENIC Version 5.15.17.4

- This driver update provides an VMMQ & RDMA driver for VIC 1400, 14000 and 15000 Series Adapters and Supported QoS Changes.

**Windows Server 2022 and 2019 ENIC Version 4.4.0.15**

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

**Windows 2025, 2022 and 2019 FNIC Driver Updates****Windows Server 2025, 2022 and 2019 FNIC Version 3.3.0.24**

- This driver update provides a Spectre-compliant fNIC driver for VIC 15000, 1400, 14000 and VIC 1300 adapters.

**VIC Management Driver for Standalone Rack Server PCIe Interface Support for Windows 2025, 2022 and 2019****Windows Server 2025, 2022 and 2019 VIC management driver version 1.0.0.1**

This driver update provides for VIC 1400 and 15000 series adapters.

**VIC Driver Updates for Release 4.3.4**


---

**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 Linux 7.9+ , 8.2+, 9.0+
- Citrix Hypervisor 8.2 LTSR
- 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



---

**Note** For the latest set of hardware and software, see the compatibility matrix here: <https://www.cisco.com/c/en/us/products/servers-unified-computing/interoperability.html>,

---

## 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
- Citrix Hypervisor 8.2 LTSR
- SUSE Linux Enterprise Server 12 SP5, 15 SP4, 15 SP5
- CentOS 7.9



---

**Note**

- RDMA and ENS features are not supported with VIC 13xx adapters from the release 4.3(4a) and later releases.  
RDMA and ENS Features are supported with VIC 14xx adapters onwards.  
2.0.xx.0 NENIC driver has support for both NENIC and RDMA capabilities.
- 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 3.3.0.24**

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

**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
- Citrix Hypervisor 8.2 LTSR
- 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
- Citrix Hypervisor 8.2 LTSR
- 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.

## Resolved Caveats

The following table lists the resolved caveats in Release 4.3.

**Table 1: Release 4.3**

Defect ID	Description	First Bundle Affected	Resolved In
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.	3.0(0.1)A	4.3(6a)
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.	3.3(0.11)A	4.3(6a)
CSCwj66629	When QinQ is enabled on vNIC (eth0 or eth1), with service profile having iSCSI policy (on vNIC eth2 or eth3), the native untagged traffic does not work through vNIC (eth0 or eth1).	4.3(4a)	4.3(6a)



Defect ID	Description	First Bundle Affected	Resolved In
CSCwk37506	When Cisco UCS servers with 1400 or 15000 series adapters have multiple paths for SANboot configured, and one path has issues in discovering the LUN while another path is successful, the clean-up done by fnic driver causes crash when the OS is loaded.  This issue is resolved.	4.3(4c)	4.3(4c)
CSCwh50478	Microsoft Windows 2022 OS resulted in bugcheck 0x50 when the interrupt count is configured to a value greater than 256  This issue is resolved.	4.3(2c)	4.3(4a)

The following table lists the resolved caveats in Release 4.2.

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

**Table 2: Release 4.2**

Defect ID	Description	First Bundle Affected	Resolved In
CSCwh50478	Microsoft Windows 2022 OS resulted in bugcheck 0x50 when the interrupt count is configured to a value greater than 256  This issue is resolved.	4.3(2c)	4.3(4a)
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 <b>Get-NetAdapterRdma</b> command on the host could show additional vPorts with RDMA Capable Flag as True. Using the <b>Get-smbclientNetworkInterface</b> 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
CSCwm26689	Upgrading the networking adapter nenic driver version to 5.13.24.2 or using this specific driver version on a newly installed Windows Server OS result in a BSOD followed by a host reboot.	If a lower nenic driver version is used instead of the version 5.13.24.2, then the BSOD does not appear.	4.3(2b), 4.3(4a), 4.3(5a)
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

Defect ID	Description	Workaround	First Release Affected
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 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"> <li>• Reboot the ESXi host to get the correct port speed reporting.</li> <li>• 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.</li> </ul>	3.1(1.152)B 2.1(2.56)A
CSCvt66474	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.  If the Windows 2019 neNIC driver is installed on Windows 2016, RDMA is not supported.  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.	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"> <li>1. On Windows Server, go to Disk Management,</li> <li>2. Right click on the disk that has gone offline and click on <b>Online</b>. This should return the LUN to service.</li> </ol>	4.2(1a)A

Defect ID	Description	Workaround	First Release Affected
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

## 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.

### 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.

Physical NIC mode is not supported in trunk mode.

### **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.

INTx interrupt mode is not supported with Windows nenic and fnic drivers.

INTx interrupt mode is not supported with Linux enic and fnic drivers.

**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-nvme\\_fc](https://docs.netapp.com/us-en/ontap-sanhost/nvme_esxi_7.html#validating-nvme_fc)

## 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](#).

## 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.