

Cisco Intersight VIC Configuration Guide

First Published: 2025-08-12

Last Modified: 2026-05-14

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

© 2025–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **Overview of Cisco Virtual Interface Card (VIC) Configuration Guide** 1

Overview 1

RDMA Over Converged Ethernet (RoCE) Version 2 1

NDIS Poll Mode for Windows 1

NFS Over RDMA for Linux 2

Single Root I/O Virtualization (SR-IOV) 2

CHAPTER 2 **Guidelines, Limitations, and Requirements** 3

RoCEv2 for Windows 3

Guidelines for Using SMB Direct support on Windows using RDMA over converged Ethernet (RoCE) v2 3

Windows Requirements 5

RoCEv2 for Linux 5

Guidelines for using NVMe over Fabrics (NVMeoF) with RoCE v2 on Linux 5

Linux Requirements 6

RoCEv2 for ESXi 7

Guidelines for using NVMeoF with RoCE v2 on ESXi 7

ESXi Requirements 8

SR-IOV for ESXi 8

Guidelines and Limitations 8

SR-IOV ESXi Requirements 9

SR-IOV for Linux 10

Guidelines and Limitations 10

SR-IOV Linux Requirements 10

NDIS Poll Mode support for Windows 11

Guidelines for using NDIS Poll Mode for Windows 11

NFS over RDMA for RHEL 11
Guidelines for Using NFS over RDMA 11

CHAPTER 3 **Configuring RDMA Over Converged Ethernet (RoCE) v2** 13
Configuring SMB Direct with RoCE v2 in Windows 13
Configuring Mode 1 on Cisco Intersight 13
Enabling RoCE Settings in LAN Connectivity Policy 14
Configuring SMB Direct Mode 1 on the Host System 17
Configuring Mode 2 on Cisco Intersight 20
Configuring Mode 2 on the Host System 24
Deleting the RoCE v2 Interface in Cisco Intersight 27
Configuring NVMe over Fabrics (NVMeoF) with RoCE v2 in Linux 28
Configuring RoCE v2 for NVMeoF on Cisco Intersight 28
Enabling RoCE Settings in LAN Connectivity Policy 29
Enabling an IOMMU BIOS Settings 32
Configuring RoCE v2 for NVMeoF on the Host System 34
Installing Cisco enic and enic_rdma Drivers 34
Discovering the NVMe Target 35
Setting Up Device Mapper Multipath 37
Deleting the RoCE v2 Interface in Cisco Intersight 38
Configuring NVMe with RoCEv2 in ESXi 38
Configuring RoCE v2 for NVMeoF on Cisco Intersight 38
Enabling RoCE Settings in LAN Connectivity Policy 40
NENIC Driver Installation 42
ESXi NVMe RDMA Host Side Configuration 42
NENIC RDMA Functionality 42
Create Network Connectivity Switches 43
Create VMVHBA Ports in ESXi 45
Displaying vmnic and vmrdma Interfaces 46
NVMe Fabrics and Namespace Discovery 47
Deleting the RoCE v2 Interface in Cisco Intersight 48
Known Issues 49
Windows 49
Linux 50

ESXi 50

CHAPTER 4 Configuring NFS Over RDMA Using ROCEv2 51

Configuring NFS over RDMA On Cisco Intersight 51

Removing NFS over RDMA Using Cisco Intersight 53

Installing Cisco eNIC and enic_rdma Drivers 54

Configuring NFS over RDMA on the Linux Host (Intel) 55

Configuring NFS over RDMA on the Linux Host (AMD) 56

Mounting and Verifying NFS over RDMA Volumes 57

CHAPTER 5 Configuring Single Root I/O Virtualization (SR-IOV) 59

Configuring BIOS and SR-IOV VFs 59

Enabling BIOS Parameters 59

Create Ethernet Adapter Policy for SR-IOV 60

Enabling SR-IOV VFs using Cisco Intersight GUI 60

Disabling SR-IOV VFs Using Cisco Intersight GUI 72

Configuring SR-IOV VFs on the ESXi Host Server 73

Installing Cisco eNIC Driver 73

Verifying the SR-IOV VFs Per Ports on the Host 73

Creating SR-IOV VFs on the Host 74

Configuring the Switch 75

Creating a Virtual Port 77

Creating a New Virtual Machine (VM) 77

Adding SR-IOV VF on the Virtual Machine 78

Installing OS on Guest VM on ESXi 78

Configuring SR-IOV VFs on the Linux Host Server 79

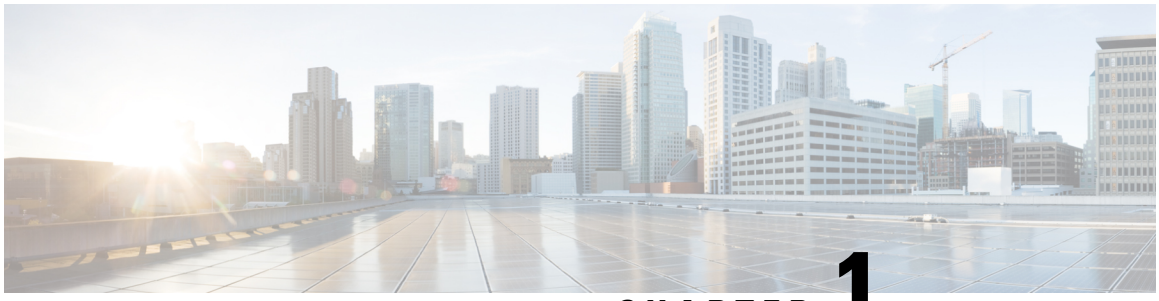
Installing Cisco eNIC Driver and Enabling IOMMU in Linux Kernel 79

Verifying the Total number of SR-IOV VFs per Port on the Host 80

Creating SR-IOV VFs on the Host 81

Creating a New Virtual Machine (VM) 82

Adding SR-IOV VF on the Virtual Machine 83



CHAPTER 1

Overview of Cisco Virtual Interface Card (VIC) Configuration Guide

- Overview, on page 1
- RDMA Over Converged Ethernet (RoCE) Version 2, on page 1
- NDIS Poll Mode for Windows, on page 1
- NFS Over RDMA for Linux, on page 2
- Single Root I/O Virtualization (SR-IOV), on page 2

Overview

A Cisco UCS network adapter can be installed to provide options for I/O consolidation and virtualization support. This guide contains configuration details on RDMA over Converged Ethernet version 2 (RoCEv2) and Single Root I/O Virtualization (SR-IOV).

RDMA Over Converged Ethernet (RoCE) Version 2

RDMA over Converged Ethernet version 2 (RoCEv2) is a network protocol that allows for Remote Direct Memory Access (RDMA) over Ethernet networks. It enables low-latency and high-bandwidth communication between servers or storage systems by leveraging the benefits of RDMA technology. RoCEv2 eliminates the need for traditional TCP/IP networking stack overhead, resulting in improved performance and reduced latency. It allows for efficient data transfers and enables applications to directly access remote memory, enhancing overall network efficiency and scalability. RoCEv2 is often used in data centers and high-performance computing environments to optimize network performance and accelerate data-intensive workloads.

RoCE v2 is supported on Windows, Linux, and ESXi platforms.

NDIS Poll Mode for Windows

Network Driver Interface Specification (NDIS) Poll Mode is an advanced traffic processing mechanism designed to optimize how the operating system manages network data. NDIS Poll Mode allows the operating system to control how network traffic is processed, providing improvements over older methods that relied on Deferred Procedure Calls (DPCs). By using the NDIS layer, the operating system manages network

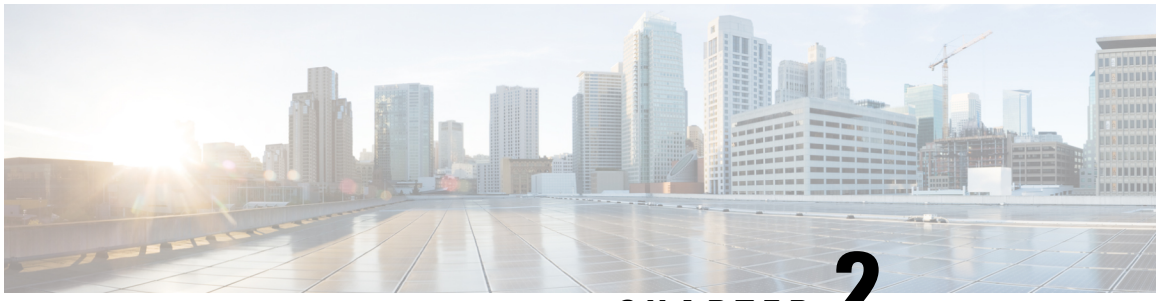
processing directly, which enhances performance, stability, and efficiency, especially during periods of heavy network activity.

NFS Over RDMA for Linux

Network File System (NFS) over Remote Direct Memory Access (RDMA) optimizes network storage performance by allowing the NFS protocol to operate over high-speed, RDMA-capable fabrics. By bypassing traditional network stack processing, this feature significantly reduces Central Processing Unit (CPU) overhead, thereby enhancing overall system efficiency.

Single Root I/O Virtualization (SR-IOV)

Single Root I/O Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead.



CHAPTER 2

Guidelines, Limitations, and Requirements

- RoCEv2 for Windows, on page 3
- RoCEv2 for Linux, on page 5
- RoCEv2 for ESXi, on page 7
- SR-IOV for ESXi, on page 8
- SR-IOV for Linux, on page 10
- NDIS Poll Mode support for Windows, on page 11
- NFS over RDMA for RHEL, on page 11

RoCEv2 for Windows

Guidelines for Using SMB Direct support on Windows using RDMA over converged Ethernet (RoCE) v2

General Guidelines and Limitations:

- Cisco Intersight Managed Mode support Microsoft SMB Direct with RoCE v2 on Microsoft Windows Server 2019 and later. Cisco recommends that you have all KB updates from Microsoft for your Windows Server release.

Note

- RoCE v2 is not supported on Microsoft Windows Server 2016.
- Refer to Windows Requirements for specific supported Operating System(OS).

- Microsoft SMB Direct with RoCE v2 is supported only with Cisco UCS VIC 1400 Series, VIC 14000, and VIC 15000 Series adapters. It is not supported with UCS VIC 1200 Series and VIC 1300 Series adapters. SMB Direct with RoCE v2 is supported on all UCS Fabric Interconnects.

Note RoCE v1 is not supported on Cisco UCS VIC 1400 Series, VIC 14000 Series, and VIC 15000 series adapters.

- RoCE v2 configuration is supported only between Cisco adapters. Interoperability between Cisco adapters and third party adapters is not supported.
- RoCE v2 supports two RoCE v2 enabled vNIC per adapter and four virtual ports per adapter interface, independent of SET switch configuration.
- RoCE v2 enabled vNIC interfaces must have the no-drop QoS system class enabled in Cisco Intersight Managed Mode domain profile.
- The RoCE Properties queue pairs setting must for be a minimum of four queue pairs and maximum number of queue pairs per adapter is 2048.
- The QoS No Drop class configuration must be properly configured on upstream switches such as Cisco Nexus 9000 series switches. QoS configurations will vary between different upstream switches.
- The maximum number of memory regions per rNIC interface is 131072.
- SMB Direct with RoCE v2 is supported on both IPv4 and IPv6.
- RoCE v2 cannot be used on the same vNIC interface as NVGRE, NetFlow, and VMQ features.
- RoCE v2 cannot be used with usNIC.
- RoCE v2 cannot be used with GENEVE offload.
- Set the QoS CoS value to 5 (Platinum). For detailed configuration on upstream Nexus switches, see FlashStack Data Center for Oracle RAC 19c Database on NVMe/RoCE.

MTU Properties:

- In older versions of the VIC driver, the MTU was derived from either a Cisco Intersight server profile or from the Cisco IMC vNIC MTU setting in standalone mode. This behavior varies for Cisco UCS VIC 1400 Series, VIC 14000 Series, and VIC 15000 Series adapters, where MTU is controlled from the Windows OS Jumbo Packet advanced property.
- The RoCE v2 MTU value is always power-of-two and its maximum limit is 4096.
- RoCE v2 MTU is derived from the Ethernet MTU.
- RoCE v2 MTU is the highest power-of-two that is less than the Ethernet MTU. For example:
  - If the Ethernet value is 1500, then the RoCE v2 MTU value is 1024
  - If the Ethernet value is 4096, then the RoCE v2 MTU value is 4096
  - If the Ethernet value is 9000, then the RoCE v2 MTU value is 4096

Windows NDPKI Modes of Operation:

- Cisco's implementation of Network Direct Kernel Provider Interface (NDPKI) supports two modes of operation: Mode 1 and Mode 2. Implementation of Network Direct Kernel Provider Interface (NDKPI) differs in Mode 1 and Mode 2 of operation: Mode 1 is native RDMA, and Mode 2 involves configuration for the virtual port with RDMA. Cisco does not support NDPKI Mode 3 operation.
- The recommended default adapter policy for RoCE v2 Mode 1 is Win-HPN-SMBd.
- The recommended default adapter policy for RoCE v2 Mode 2 is MQ-SMBd.
- RoCE v2 enabled vNICs for Mode 2 operation require the QoS host control policy set to full.

- Mode 2 is inclusive of Mode 1: Mode 1 must be enabled to operate Mode 2.
- On Windows, the RoCE v2 interface supports both MSI & MSIx interrupts mode. Default interrupt mode is MSIx. Cisco recommends you avoid changing interrupt mode when the interface is configured with RoCE v2 properties.

Downgrade Limitations:

- Cisco recommends you remove the RoCE v2 configuration before downgrading to any non-supported firmware release. If the configuration is not removed or disabled, downgrade will fail.

Windows Requirements

Configuration and use of RDMA over Converged Ethernet for RoCE v2 in Windows Server requires the following:

- Windows 2019 or Windows Server 2022 or Windows 2025 with latest Microsoft updates
- VIC Driver version 5.4.0.x or later
- Cisco UCS M5 B-Series or C-Series with Cisco UCS 1400 Series adapters: only Cisco UCS VIC 1400 Series or VIC 15000 series adapters are supported.
- Cisco UCS M6 B-Series, C-Series, or X-Series servers with Cisco UCS VIC 1400, VIC 14000, or VIC 15000 series adapters.
- Cisco UCS M7 C-Series, or X-Series servers with Cisco UCS VIC 1400, VIC 14000, or VIC 15000 series adapters.
- Cisco UCS M8 C-Series, , or X-Series servers with Cisco VIC 15000 series adapters.

Note All Powershell commands or advanced property configurations are common across Windows 2019 and 2022 unless explicitly mentioned.

RoCEv2 for Linux

Guidelines for using NVMe over Fabrics (NVMeoF) with RoCE v2 on Linux

General Guidelines and Limitations:

- Cisco recommends you check UCS Hardware and Software Compatibility to determine support for NVMeoF. NVMeoF is supported on Cisco UCS B-Series, C-Series, and X-Series servers.
- NVMe over RDMA with RoCE v2 is supported with the Cisco UCS VIC 1400, VIC 14000, and VIC 15000 Series adapters.
- When creating RoCE v2 interfaces, use Cisco Intersight provided Linux-NVMe-RoCE adapter policy.

- In the Ethernet Adapter policy, do not change values of Queue Pairs, Memory Regions, Resource Groups, and Priority settings other than to Cisco provided default values. NVMeoF functionality may not be guaranteed with different settings for Queue Pairs, Memory Regions, Resource Groups, and Priority.
- When configuring RoCE v2 interfaces, use both the enic and enic_rdma binary drivers downloaded from Cisco.com and install the matched set of enic and enic_rdma drivers. Attempting to use the binary enic_rdma driver downloaded from Cisco.com with an inbox enic driver will not work.
- RoCE v2 supports maximum two RoCE v2 enabled interfaces per adapter.
- Booting from an NVMeoF namespace is not supported.
- RoCEv2 and GENEVE offload can be configured on different interfaces on the same Cisco UCS VIC adapter.
- RoCEv2 cannot be used with QinQ.
- Layer 3 routing is not supported.
- RoCE v2 does not support bonding.
- Saving a crashdump to an NVMeoF namespace during a system crash is not supported.
- NVMeoF cannot be used with usNIC, VxLAN, VMQ, VMMQ, NVGRE, GENEVE Offload, and DPDK features.
- Set the QoS CoS value to 5 (Platinum). For detailed configuration on upstream Nexus switches, see FlashStack Data Center for Oracle RAC 19c Database on NVMe/RoCE.
- Cisco Intersight does not support fabric failover for vNICs with RoCE v2 enabled.
- The Quality of Service (QoS) no drop class configuration must be properly configured on upstream switches such as Cisco Nexus 9000 series switches. QoS configurations will vary between different upstream switches.
- Spanning Tree Protocol (STP) may cause temporary loss of network connectivity when a failover or failback event occurs. To prevent this issue from occurring, disable STP on uplink switches.

Linux Requirements

Configuration and use of RoCEv2 in Linux requires the following:

- InfiniBand kernel API module ib_core
- nvme-cli package
- Minimum VIC firmware 5.1(2x) or later for IPv6 support
- Cisco UCS B-series, Cisco UCS C-series, and Cisco UCS X-series servers with Cisco UCS VIC 1400 or Cisco UCS VIC 15000 Series adapters
- A storage array that supports NVMeoF connection
- eNIC driver version 4.0.0.10-802.34 or later and enic_rdma driver version 1.0.0.10-802.34 or later

Note Ubuntu 24.04.1 with kernel 6.8.0-51-generic starts supporting RoCE v2 from eNIC driver version 4.8.0.0-1128.4 and enic_rdma driver version 1.8.0.0-1128.4.

- Red Hat Enterprise Linux 8.x, 9.x and 10.x versions

Interrupts

- Linux RoCEv2 interface supports only MSIx interrupt mode. Cisco recommends avoiding changing interrupt mode when the interface is configured with RoCEv2 properties.
- The minimum interrupt count for using RoCEv2 with Linux is 8.

RoCEv2 for ESXi

Guidelines for using NVMeoF with RoCE v2 on ESXi

General Guidelines and Limitations:

- Cisco recommends checking the UCS Hardware and Software Compatibility to determine support for NVMeoF. NVMeoF is supported on Cisco UCS B-Series, C-Series, and X-Series servers.
- Nonvolatile Memory Express (NVMe) over RDMA with RoCE v2 is currently supported only with Cisco VIC 15000 Series adapters.
- When creating RoCE v2 interfaces, use Cisco Intersight provided VMWareNVMeRoCEv2 adapter policy.
- When creating RoCE v2 interfaces, use Cisco recommended Queue Pairs, Memory Regions, Resource Groups, and Class of Service settings. NVMeoF functionality may not be guaranteed with different settings for Queue Pairs, Memory Regions, Resource Groups, and Class of Service.
- RoCE v2 supports maximum two RoCE v2 enabled interfaces per adapter.
- Booting from an NVMeoF namespace is not supported.
- RoCEv2 and GENEVE offload can be configured on different interfaces on the same Cisco UCS VIC adapter.
- RoCEv2 cannot be used with QinQ.
- SR-IOV cannot be configured on the same vNIC with VXLAN, Geneve Offload, QinQ, VMQ/VMMQ, RoCE, or usNIC.
- Layer 3 routing is not supported.
- Saving a crashdump to an NVMeoF namespace during a system crash is not supported.
- NVMeoF with RoCE v2 cannot be used with usNIC, VxLAN, VMQ, VMMQ, NVGRE, GENEVE Offload, ENS, and DPDK features.
- Cisco Intersight does not support fabric failover for vNICs with RoCE v2 enabled.

- The Quality of Service (QoS) no drop class configuration must be properly configured on upstream switches such as Cisco Nexus 9000 series switches. QoS configurations will vary between different upstream switches.
- During the failover or failback event, the Spanning Tree Protocol (STP) can result temporary loss of network connectivity. To prevent this connectivity issue, disable STP on uplink switches.
- Set the QoS CoS value to 5 (Platinum). For detailed configuration on upstream Nexus switches, see FlashStack Data Center for Oracle RAC 19c Database on NVMe/RoCE.

ESXi Requirements

Configuration and use of RoCE v2 in ESXi requires the following:

- VMware ESXi 7.0 U3 and 8.0 or later
- VIC firmware 5.2(3x) or later versions.
- The driver version, *nenic-2.0.4.0-1OEM.700.1.0.15843807.x86_64.vib* that provides both standard eNIC and RDMA support
- A storage array that supports NVMeoF connection.

SR-IOV for ESXi

Guidelines and Limitations

- Cisco recommends checking the UCS Hardware and Software Compatibility to determine support for SR-IOV.
- SR-IOV is supported with Cisco UCS AMD®/Intel® based B-Series, C-Series, and X-Series servers.
  - SR-IOV is not supported in Physical NIC mode.
  - SR-IOV does not support VLAN Access mode.
- Each vNIC supports up to 64 Virtual Functions (VFs). For each VF, the configuration includes: Up to 8 RQs, Up to 8 WQs, Up to 16 CQs, and Up to 16 interrupts
- SR-IOV cannot be configured on the same vNIC with VXLAN, Geneve Offload, QinQ, VMQ/VMMQ, RoCE, or usNIC.
- Cisco IMM does not limit the total number of VFs, Receive Queue Count Per VF, Transmit Queue Count Per VF, Completion Queue Count Per VF, and Interrupt Count Per VF values. However, if any one of the resources exceed the adapter limit, the Server Profile deployment will fail with a resource error. In this case, either reduce the number of VFs or adjust the failed resource value accordingly.
- For ESXi hosts using SR-IOV with VFs on vNICs and VMs, the system may crash with a PSOD during cold or warm reboots. This behavior is related to the handling of VFs in the environment.
- Enabling some of the features concurrently with SR-IOV will lead to a Server Profile Deployment failure. Ensure the following features are disabled when configuring SR-IOV on a vNIC:

- VMQ
- usNIC
- Geneve Offload
- RoCE
- QinQ Tunnelling
- NVGRE
- VXLAN

The following features are not supported on SR-IOV:

- aRFS
- iSCSI Boot
- DPDK when the host has Linux OS
- Precision Time Protocol (PTP)



Note SR-IOV interface supports Message-Signaled Interrupts (MSIs) interrupt mode.

SR-IOV ESXi Requirements

Configuration and use of SR-IOV in ESXi requires the following:

- Cisco VIC firmware version 5.3(2.32) or later
- VMware ESXi 7.0 U3, 8.0, 9.0 or later
- VMs with RHEL 8.7, 9.0, and 10.0 or later
- Cisco VMware nENIC driver version 2.0.10.0 for ESXi 7.0 U3, 2.0.11.0 for ESXi 8.0 U3, and 2.0.18.0 for ESXI 9.0 or later
- Cisco RHEL ENIC driver version 4.4.0.1-930.10 for RHEL 8.7 and 9.0 and later
- Cisco RHEL ENIC driver version 4.9.0.1-1160.11 for RHEL 9.6 and 10.0 and later



Note SR-IOV is not supported on Cisco UCS VIC 1200 and Cisco UCS VIC 1300 series adapters.

SR-IOV for Linux

Guidelines and Limitations

- Cisco recommends checking the UCS Hardware and Software Compatibility to determine support for SR-IOV.
- SR-IOV is supported with AMD®/Intel® based Cisco UCS C-Series, B-Series, and X-Series servers.
- SR-IOV is not supported in Physical NIC mode.
- SR-IOV does not support VLAN Access mode.
- SR-IOV cannot be configured on the same vNIC with VXLAN, Geneve Offload, QinQ, VMQ/VMMQ, RoCE, or usNIC.
- aRFS is not supported on SR-IOV VF.
- iSCSI boot is not supported on SR-IOV VF.
- DPDK on SRIOV VF is not supported when the host has Linux OS.
- SR-IOV interface supports MSIx interrupt mode.
- Precision Time Protocol (PTP) is not supported on SR-IOV VF.
- The system may experience a PSOD when multiple vNICs are configured with SR-IOV and VMs are enumerated with Virtual Functions (VFs), especially when either cold or warm boots are performed. For Linux operating systems, after a system reboot, the VFs need to be reconfigured because they are not persistent across reboots.

SR-IOV Linux Requirements

Configuration and use of SR-IOV in Linux requires the following:

- Host OS: Red Hat Enterprise Linux 8.10, 9.4 or later, 10.0 or later, and Ubuntu 22.0.4.2 LTS or later
- Guest OS: Red Hat Enterprise Linux 8.10, 9.4 or later, 10.0 or later, and Ubuntu 22.0.4.2 LTS or later
- Virtualization Packages installed on the host
- eNIC driver version 4.7.0.5-1076.6 or later
- Cisco UCS Manager Release 4.3(5a) or later
- Cisco VIC firmware 5.3(4.75) or later

NDIS Poll Mode support for Windows

Guidelines for using NDIS Poll Mode for Windows

Guidelines for using NDIS Poll Mode for Windows

NDIS Poll Mode is enabled by default in supported Cisco VIC driver packages. To modify settings, use the `Set-NetAdapterAdvancedProperty` PowerShell cmdlet or navigate to the **Advanced** tab in the Device Manager properties for the specific Cisco VIC adapter.

Guidelines

- **Supported Operating Systems:** Windows Server 2025 and later
- **Supported Adapters:** Cisco UCS VIC 15000 Series
- **Supported Driver Version:** NENIC Version 5.15.17.4

Note For additional information on NDIS Poll Mode and its underlying mechanisms in Windows drivers, refer to Microsoft's Windows Driver and Network documentation.

NFS over RDMA for RHEL

Guidelines for Using NFS over RDMA

NFS over RDMA uses RDMA-capable network adapters to allow data to be copied directly between storage system memory and host system memory, bypassing CPU interruptions and overhead. This configuration is ideal for latency-sensitive or high-bandwidth workloads.

General Guidelines and Limitations

- Cisco recommends you check Supported Hardware for Intersight Managed Mode specific to your Cisco Intersight release to determine support for NFS.
- NFS over RDMA is supported on Cisco UCS X-Series, B-Series, and C-Series M6 and later servers equipped with Cisco UCS VIC 15000 Series adapters.
- NFS over RDMA is supported for NFS version 4.1 and later.
- Each Cisco VIC adapter supports a maximum of two RDMA-enabled vNICs.
- When creating an RDMA policy for RoCEv2, ensure to configure the following values:
  - **Queue Pairs:** Choose or enter 1024
  - **Memory Regions:** Choose or enter 131072
  - **Resource Groups:** Choose or enter 8

  - **Version:**2

- When configuring NFS over RDMA using RoCEv2 interfaces, install the matched set of enic and enic_rdma binary drivers downloaded from Cisco.com. Attempting to use the binary enic_rdma driver with an inbox enic driver is not supported.
- Set the QoS CoS value to 5 (Platinum). For detailed configuration on upstream Nexus switches, see FlashStack Data Center for Oracle RAC 19c Database on NVMe/RoCE.
- A No-Drop QoS policy should be configured carefully across the setup.
- Set the Maximum Transmission Unit (MTU) size correctly across the VLANs and the QoS policy.
- NFS over RDMA does not support fabric failover, Layer 3 routing, bonding, Netflow monitoring, or GENEVE offload.
- NFS over RDMA does not support usNIC, VxLAN, VMQ, VMMQ, NVGRE, and DPDK features.
- Saving a crashdump to NFS over RDMA mount during a system crash are not supported.

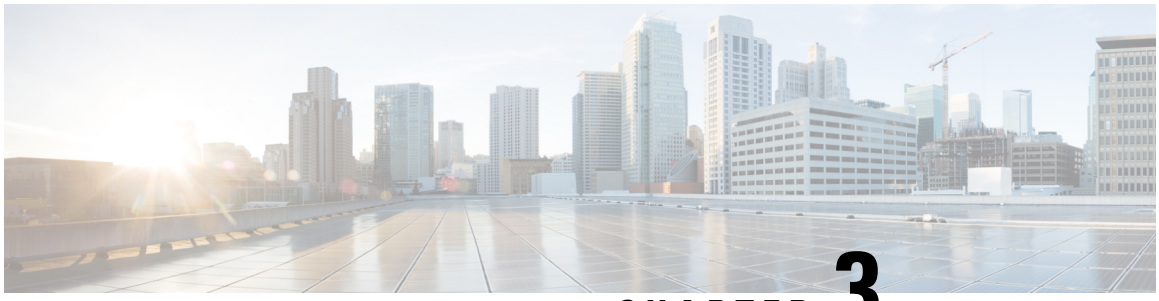
Requirements

- **Supported Adapters**: Cisco UCS VIC 15000 Series adapters.
- **Supported Operating Systems**: Red Hat Enterprise Linux (RHEL) 9.7 or later and 10.1 or later.
- **Storage Support**: A storage array that supports NFS over RDMA (Example, Array AFF-A20; Version 9.17.1P1)
- **QoS Configuration**: A no-drop QoS system class must be enabled in Cisco Intersight.

Interrupts

- NFS Over RDMA supports only MSIx interrupt mode.
- The minimum interrupt count is 256.
- Cisco recommends not to change the interrupt mode once NFS Over RDMA interface is configured.

Downgrade Limitations Before downgrading to a non-supported release version, it is recommended to remove any RDMA-capable vNIC configurations.



CHAPTER 3

Configuring RDMA Over Converged Ethernet (RoCE) v2

- Configuring SMB Direct with RoCE v2 in Windows, on page 13
- Configuring NVMe over Fabrics (NVMeoF) with RoCE v2 in Linux, on page 28
- Configuring NVMe with RoCEv2 in ESXi, on page 38
- Known Issues, on page 49

Configuring SMB Direct with RoCE v2 in Windows

Configuring Mode 1 on Cisco Intersight

Use these steps to configure the RoCE v2 Mode 1 interface on Cisco Interisght.

To avoid possible RDMA packet drops, ensure same no-drop COS is configured across the network. The following steps allows you to configure a no-drop class in System QoS policies and use it for RDMA supported interfaces.

For Cisco UCS M8 C-Series or X-Series servers, the VIC 15000 series is supported, while the Cisco UCS VIC 1400 Series, 14000 Series is not compatible with M8 servers.

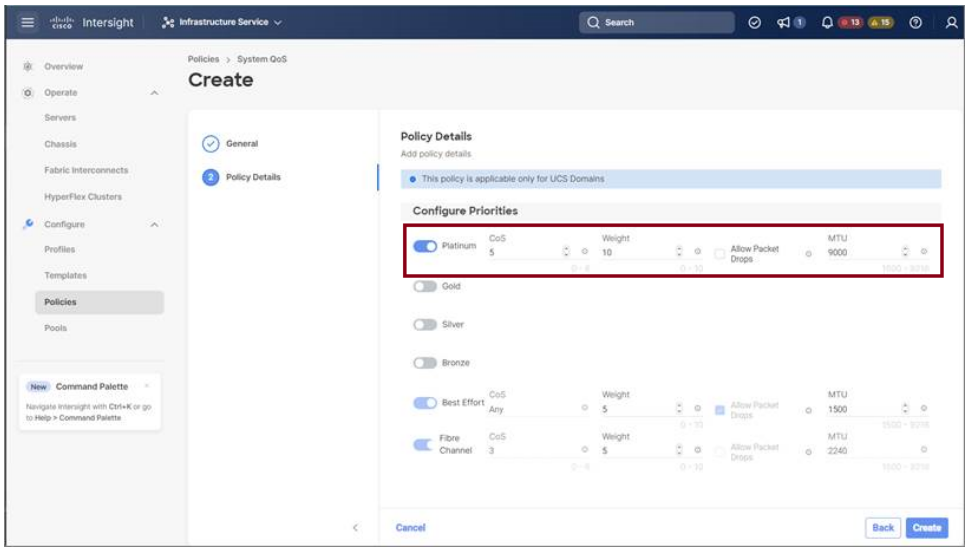
Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Domain** platform type, search or choose **System QoS**, and click **Start**.

Step 2 In the **General** page, enter the policy name and click **Next**, and then in the **Policy Details** page, configure the property setting for System QoS policy as follows:

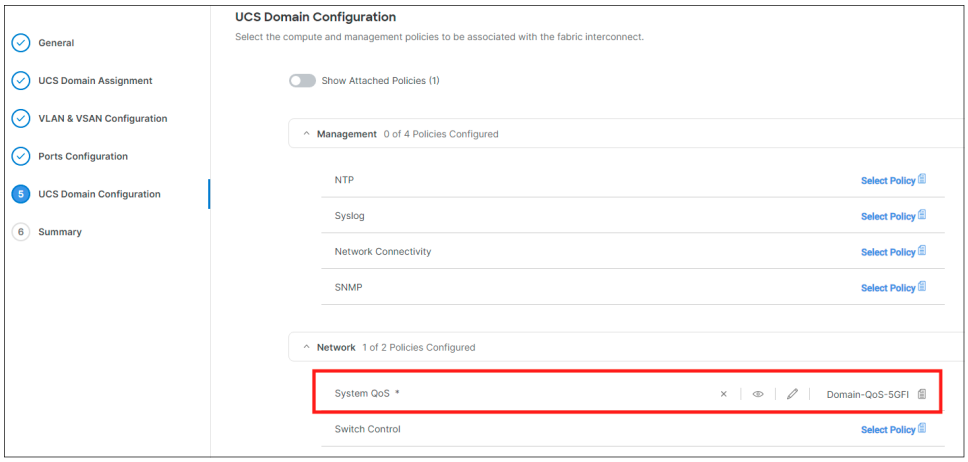
- For **Priority**, choose **Platinum**
- For **Allow Packet Drops**, uncheck the check box.

Note
For more information on MTU field, see *MTU Properties* in Guidelines for Using SMB Direct support on Windows using RDMA over converged Ethernet (RoCE) v2, on page 3



Step 3 Click **Create**

Step 4 Associate the System QoS policy to the Domain Profile and deploy.



Note
For more information, see *Creating System QoS Policy* in Configuring Domain Policies and Configuring Domain Profiles.

The System QoS Policy is successfully created and deployed to the Domain Profile.

What to do next

Configure the server profile with RoCE v2 vNIC settings in LAN Connectivity policy.

Enabling RoCE Settings in LAN Connectivity Policy

Use these steps to configure the RoCE v2 vNIC settings in Mode 1. In Cisco Intersight LAN Connectivity policy, you can enable the RoCE settings on **Ethernet QoS** policy and **Ethernet Adapter** policy for Mode 1 configuration as follows:

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **LAN Connectivity** policy, and click **Start**.

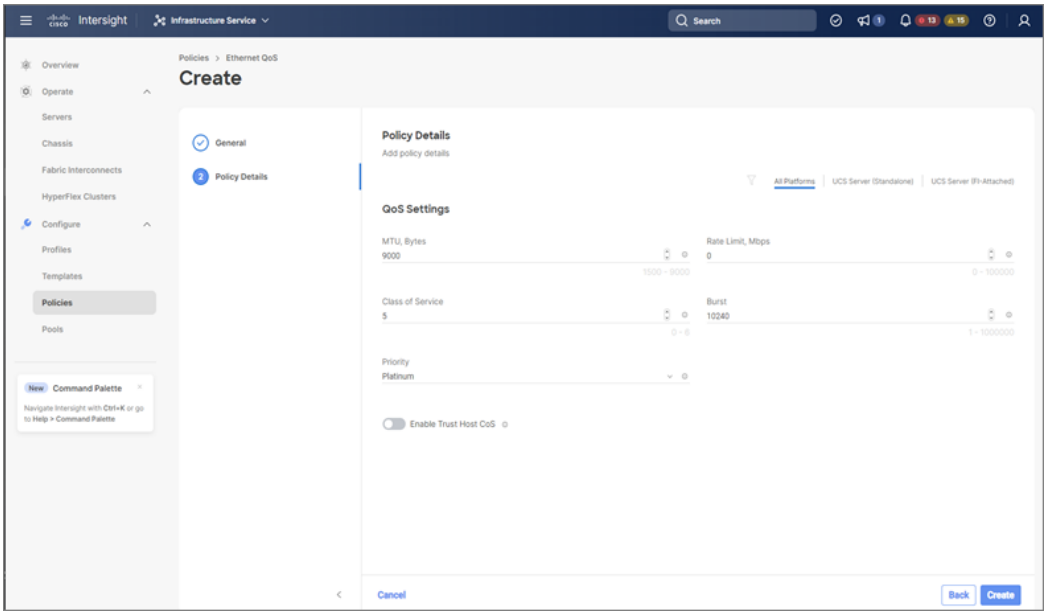
Step 2 In the policy **General** page, enter the policy name, select the Target Platform as **UCS Server (Standalone)** or **UCS Server (FI-Attached)**, and click **Next**.

Step 3 In the **Policy Details** page, click **Add vNIC** to create a new vNIC.

Step 4 In the **Add vNIC** page, follow the configuration parameters to enable the RoCE vNIC settings:

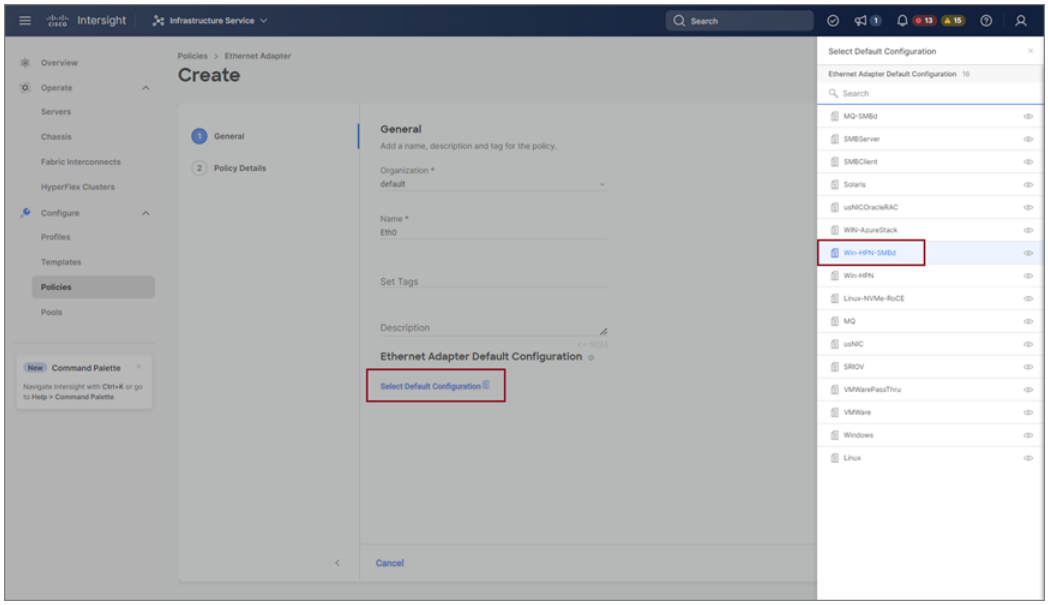
- In the **General** section, provide a name for virtual ethernet interface.
- In the **Consistent Device Naming (CDN)** section of the Standalone server or the **Failover** section of FI-attached server, do the following:
  - Click **Select Policy** link below the **Ethernet QoS**. Use the **Create New** button to create a new Ethernet QoS policy with the following property settings:
    - For **MTU**, choose or enter **1500, 4096, 9000, or 9158**
    - For **Priority**, choose **Platinum** or **any no-drop**
    - For **Class of Service**, choose or enter **5**

      Note
      This property is available only on Standalone servers.

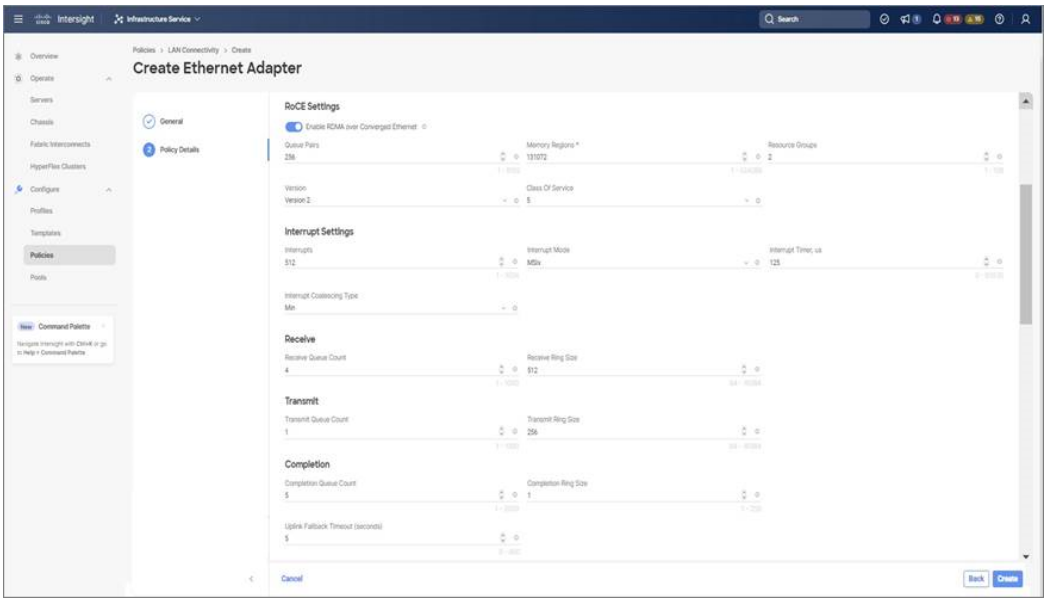


- Click **Select Policy** link below the **Ethernet Adapter**. Follow on to click Create an Ethernet Adapter Policy:
  - **Use the Default Configuration**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy and under Ethernet Adapter Default Configuration click **Select Default**

Configuration to search and select **Win-HPN-SMBd**, the pre-defined Ethernet Adapter Default Configuration. Click **Next** and then **Create**.



- **Configure RoCE Settings in the policy**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy. Under Policy Details page on right pane, use the following property settings, then click **Next**, and then **Create**..
  - For **Enable RDMA over Converged Ethernet**, slide to enable.
  - For **Queue Pairs**, choose or enter **256**
  - For **Memory Regions**, choose or enter **131072**
  - For **Resource Groups**, choose or enter **2**
  - For **Version**, select **Version 2**



- Click **Add** to add and save the new vNIC settings.

Note
All the fields with * are mandatory for creating LAN Connectivity Policy. Ensure they are filled out or selected with appropriate policies.

Step 5 Click **Create** to complete the LAN Connectivity policy with RoCE v2 property settings.

Step 6 Associate the LAN Connectivity policy to the server profile and deploy.

Note
For more information, see *Creating a LAN Connectivity Policy*, *Creating an Ethernet QoS Policy*, and *Creating an Ethernet Adapter Policy* in Configuring UCS Server Policies and Configuring UCS Server Profiles.

The LAN Connectivity policy with the Ethernet QoS policy and Ethernet Adapter policy vNIC setting is successfully created and the server profile is deployed to enable RoCE v2 configuration.

What to do next

Once the policy configuration for RoCE v2 is complete, proceed to enable IOMMU in the BIOS policy.

Configuring SMB Direct Mode 1 on the Host System

You will configure connection between smb-client and smb-server on two host interfaces. For each of these servers, smb-client and smb-server, configure the RoCE v2-enabled vNIC as described below.

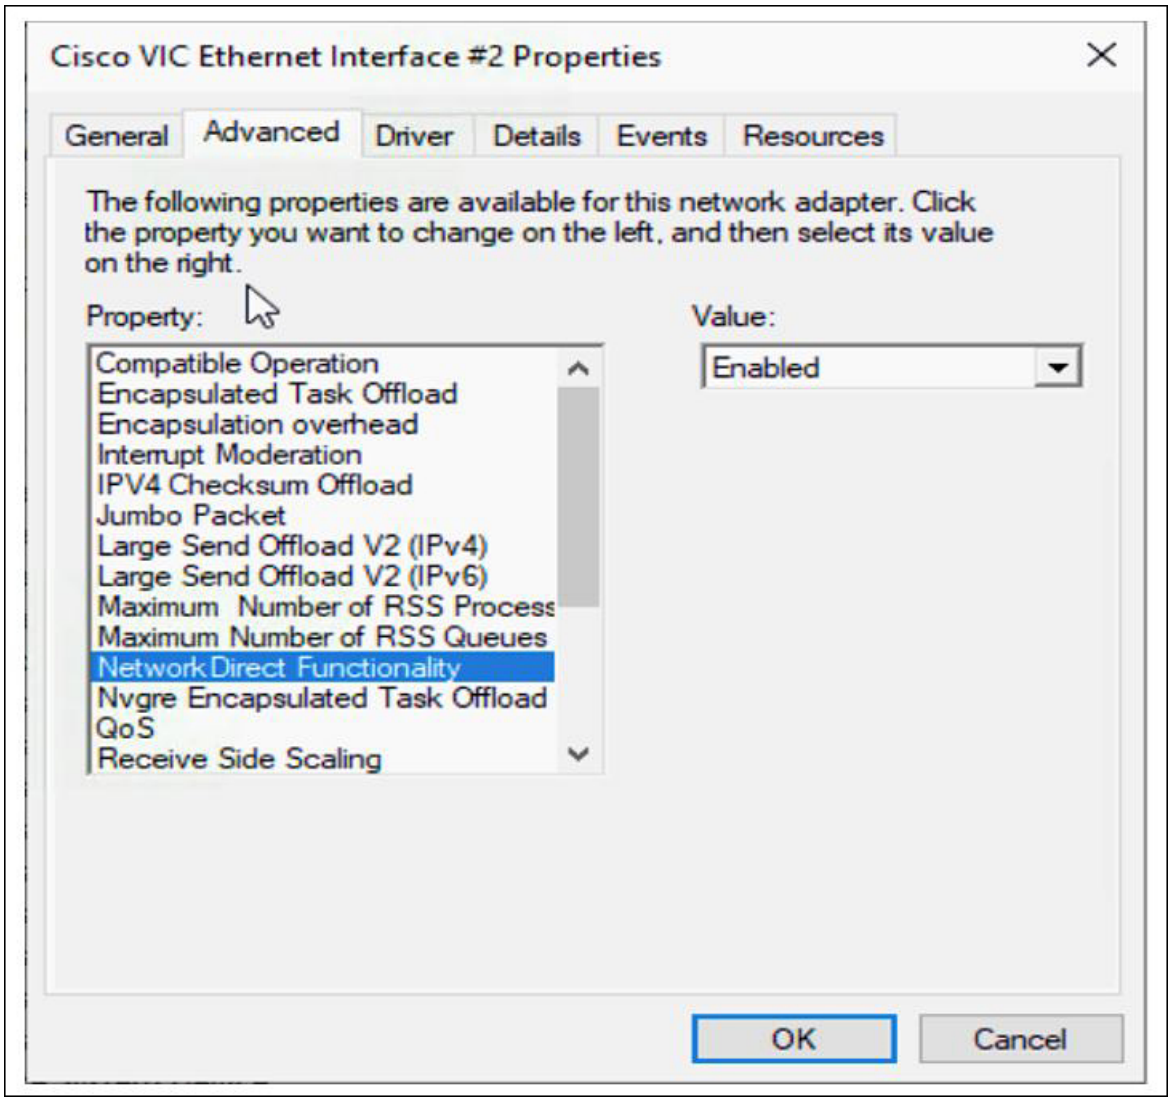
Before you begin

Configure RoCE v2 for Mode 1 in Cisco Intersight.

Procedure

Step 1 In the Windows host, go to the Device Manager and select the appropriate Cisco VIC Internet Interface.

Step 2 Go to **Tools** > **Computer Management** > **Device Manager** > **Network Adapter** > click on **VIC Network Adapter** > **Properties** > **Advanced** > **Network Direct Functionality**. Perform this operation for both the smb-server and smb-client vNICs.



Step 3 Verify that RoCE is enabled on the host operating system using PowerShell.

The `Get-NetOffloadGlobalSetting` command shows NetworkDirect is enabled.

```
PS C:\Users\Administrator> Get-NetOffloadGlobalSetting

ReceiveSideScaling          : Enabled
ReceiveSegmentCoalescing    : Enabled
Chimney                     : Disabled
```

```
TaskOffload                   : Enabled
NetworkDirect                 : Enabled
NetworkDirectAcrossIPSubnets  : Blocked
PacketCoalescingFilter        : Disabled
```

Note
If the NetworkDirect setting is showing as disabled, enable it using the command: `Set-NetOffloadGlobalSetting -NetworkDirect enabled`

Step 4 Bring up Powershell and enter the command:

```
get-SmbClientNetworkInterface
```

```
PS C:\Users\Administrator>
PS C:\Users\Administrator> Get-SmbClientNetworkInterface

Interface Index RSS Capable RDMA Capable Speed    IpAddresses    Friendly Name
--------------- ----------- ------------ -----    -----------    -------------
14              True        False        40 Gbps  {10.37.60.162} vEthernet (vswitch)
26              True        True         40 Gbps  {10.37.60.158} vEthernet (vp1)
9               True        True         40 Gbps  {50.37.61.23}  Ethernet 2
5               False       False        40 Gbps  {169.254.10.5} Ethernet (Kernel Debugger)
8               True        False        40 Gbps  {169.254.4.26} Ethernet 3


PS C:\Users\Administrator>
```

Step 5 Enter **enable - netadapterrdma [-name] ["Ethernetname"]**

Step 6 Verify the overall RoCE v2 Mode 1 configuration at the Host as follows:

a) Use the Powershell command **netstat -xan** to verify the listeners in both the smb-client and smb-server Windows host; listeners will be shown in the command output.

```
PS C:\Users\Administrator>
PS C:\Users\Administrator> netstat -xan

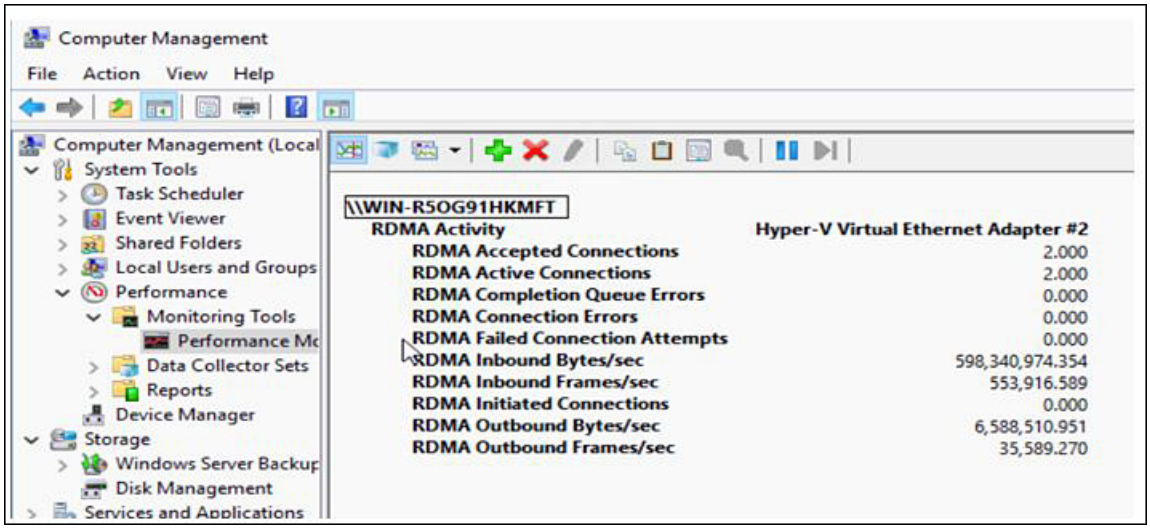
Active NetworkDirect Connections, Listeners, SharedEndpoints

  Mode   IfIndex Type            Local Address          Foreign Address        PID

  Kernel       9 Listener        50.37.61.23:445        NA                     0
  Kernel      26 Listener        10.37.60.158:445       NA                     0
PS C:\Users\Administrator>
```

b) Go to the smb-client server fileshare and start an I/O operation.

c) Go to the performance monitor and check that it displays the RDMA activity.



Step 7 In the Powershell command window, check the connection entries with the **netstat -xan** output command to make sure they are displayed. You can also run **netstat -xan** from the command prompt. If the connection entry shows up in netstat-xan output, the RoCE v2 mode1 connections are correctly established between client and server.

```
PS C:\Users\Administrator> netstat -xan

Active NetworkDirect Connections, Listeners, SharedEndpoints

  Mode   IfIndex Type            Local Address          Foreign Address        PID

  Kernel       4 Connection      50.37.61.22:445        50.37.61.71:2240       0
  Kernel       4 Connection      50.37.61.22:445        50.37.61.71:2496       0
  Kernel      11 Connection      50.37.61.122:445       50.37.61.71:2752       0
  Kernel      11 Connection      50.37.61.122:445       50.37.61.71:3008       0
  Kernel      32 Connection      10.37.60.155:445       50.37.60.61:49092      0
  Kernel      32 Connection      10.37.60.155:445       50.37.60.61:49348      0
  Kernel      26 Connection      50.37.60.32:445        50.37.60.61:48580      0
  Kernel      26 Connection      50.37.60.32:445        50.37.60.61:48836      0
  Kernel       4 Listener        50.37.61.22:445        NA                     0
  Kernel      11 Listener        50.37.61.122:445       NA                     0
  Kernel      32 Listener        10.37.60.155:445       NA                     0
  Kernel      26 Listener        50.37.60.32:445        NA                     0
```

Note
IP values are representative only.

Step 8 By default, Microsoft's SMB Direct establishes two RDMA connections per RDMA interface. You can change the number of RDMA connections per RDMA interface to one or any number of connections.

For example, to increase the number of RDMA connections to 4, type the following command in PowerShell:

```
PS C:\Users\Administrator> Set-ItemProperty -Path `
"HKLM:\SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters"
ConnectionCountPerRdmaNetworkInterface -Type DWORD -Value 4 -Force
```

Configuring Mode 2 on Cisco Intersight

Use these steps to configure the RoCE v2 policies in Mode 2. In Cisco Intersight LAN Connectivity Policy, you can enable the RoCE settings on **Ethernet QoS** policy and **Ethernet Adapter** policy, and **VMMQ Adapter** policy for Mode 2 configuration as follows:

You will apply the VMQ Connection Policy as vmmq.

Before you begin

Configure RoCE v2 Policies in Mode 1.

Use the pre-defined default adapter policy “MQ-SMBd”, or configure a user-defined Ethernet adapter policy with the following recommended RoCE-specific parameters:

- RoCE: Enabled
- Version 1: disabled
- Version 2: enabled
- Queue Pairs: 256
- Memory Regions: 65536
- Resource Groups: 2

- Priority: Platinum

Create a VMQ connection policy with the following values:

- Multi queue: Enabled
- Number of sub-vNIC: 16
- VMMQ adapter policy: MQ-SMBd

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **LAN Connectivity** policy, and click **Start**.

Step 2 In the policy **General** page, enter the policy name, select the Target Platform as **UCS Server (Standalone)** or **UCS Server (FI-Attached)**, and click **Next**.

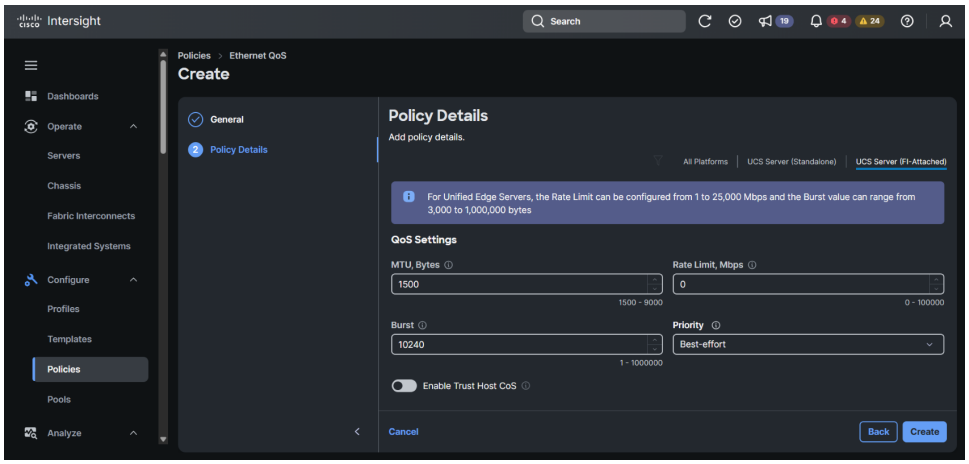
Step 3 In the **Policy Details** page, click **Add vNIC** to create a new vNIC.

Step 4 In the **Add vNIC** page, follow the configuration parameters to enable the RoCE vNIC settings:

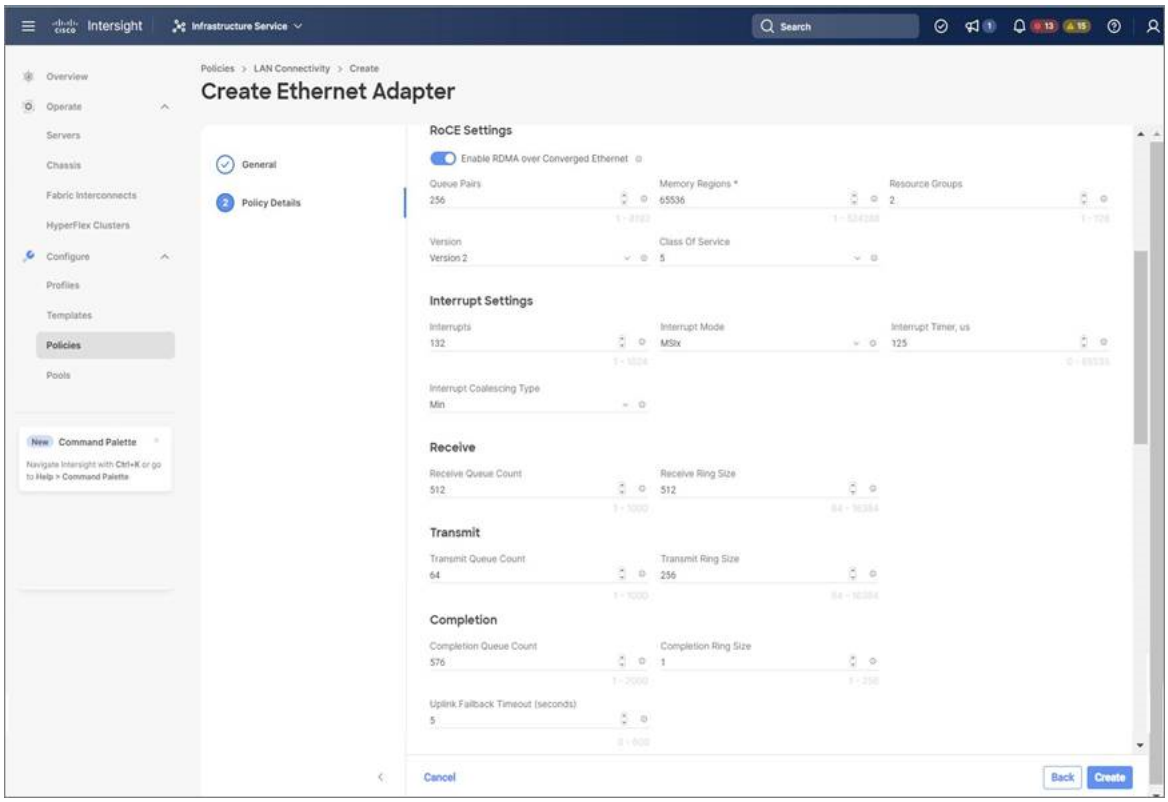
a) In the **General** section, provide a name for virtual ethernet interface.

b) In the **Consistent Device Naming (CDN)** section of the Standalone server or the **Failover** section of FI-attached server, do the following:

- Click **Select Policy** link below the **Ethernet QoS**. Use the **Create New** button to create a new Ethernet QoS policy with the following property settings:
  - **MTU**—The Maximum Transmission Unit (MTU) or packet size that the virtual interface accepts. For **MTU**, choose or enter **1500, 4096, 9000, or 9158.**
  - **Rate Limit, Mbps**—The value in Mbps (0-10G/40G/100G depending on Adapter Model) to use for limiting the data rate on the virtual interface.
  - **Class of Service**—Class of Service to be associated to the traffic on the virtual interface.
  - **Burst**—The burst traffic, in bytes, allowed on the vNIC.
  - For **Priority**, choose or enter **Best-effort**
  - **Enable Trust Host CoS**, slide to enable



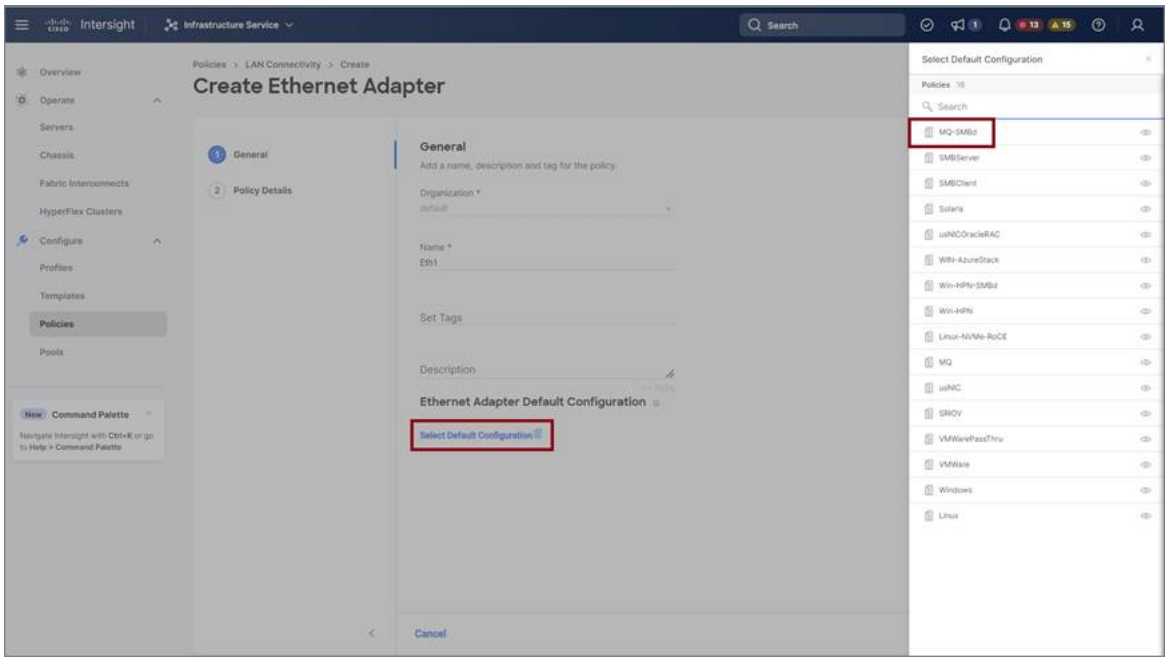
- Click **Select Policy** link below the **Ethernet Adapter**. Use **Create New** button to create a new Ethernet Adapter policy with the following property settings:
  - For **Enable RDMA over Converged Ethernet**, slide to enable.
  - For **Queue Pairs**, select or enter **256**
  - For **Memory Regions**, select or enter **65536**
  - For **Resource Groups**, select or enter **2**
  - For **Version**, choose **Version 2**
  - For **Class of Service**, choose or enter **5**



- In the **Connection** section, use the following property setting for VMQ Connection and to create VMMQ Adapter policy:
  - For connection, select **VMQ**.
  - **Enable Virtual Machine Multi-Queue** using slider button.
  - For **Number of Sub vNICs**, select or enter **4**
  - For **VMMQ Adapter Policy**, click **Select Policy** link below the VMMQ Adapter Policy and do the following:
    - Click **Create New** to create a new policy. In the **General** page, enter the name of the policy and click **Select Default Configuration** to search and select **MQ-SMBd**, the pre-defined VMMQ Adapter default configuration.

      Attention
      Do not modify the pre-defined parameters under Policy Details page, retain the default settings.

    - Click **Next** and then **Create**.



- Click **Add** to add and save the new vNIC settings.

Note
All the fields with * are mandatory and ensure it is filled out or selected with appropriate policies.

Step 5 Click **Create** to complete the LAN Connectivity policy with RoCE v2 property settings.

Step 6 Associate the LAN Connectivity policy to the server profile.

Note
For more information on *Creating an Ethernet QoS*, *Ethernet Adapter Policy*, and *VMMQ Adapter Policy*, see Configuring UCS Server Policies and Configuring UCS Server Profiles.

The LAN Connectivity Policy with Ethernet QoS Policy, Ethernet Adapter Policy, and VMMQ Adapter Policy are successfully created and deployed to enable RoCE v2 configuration.

What to do next

Once the policy configuration for RoCE v2 is complete, reboot the server and proceed with the RoCE v2 Mode 2 configuration on the host operating system.

Configuring Mode 2 on the Host System

This task uses Hyper-V virtualization software that is compatible with Windows Server 2019 and Windows Server 2022.

Follow the below procedure for the host operating system configuration for RoCEv2 Mode 2.

Before you begin

- Configure and confirm the connection for Mode 1 for both Cisco Intersight and Host.
- Configure Mode 2 in Cisco Intersight.

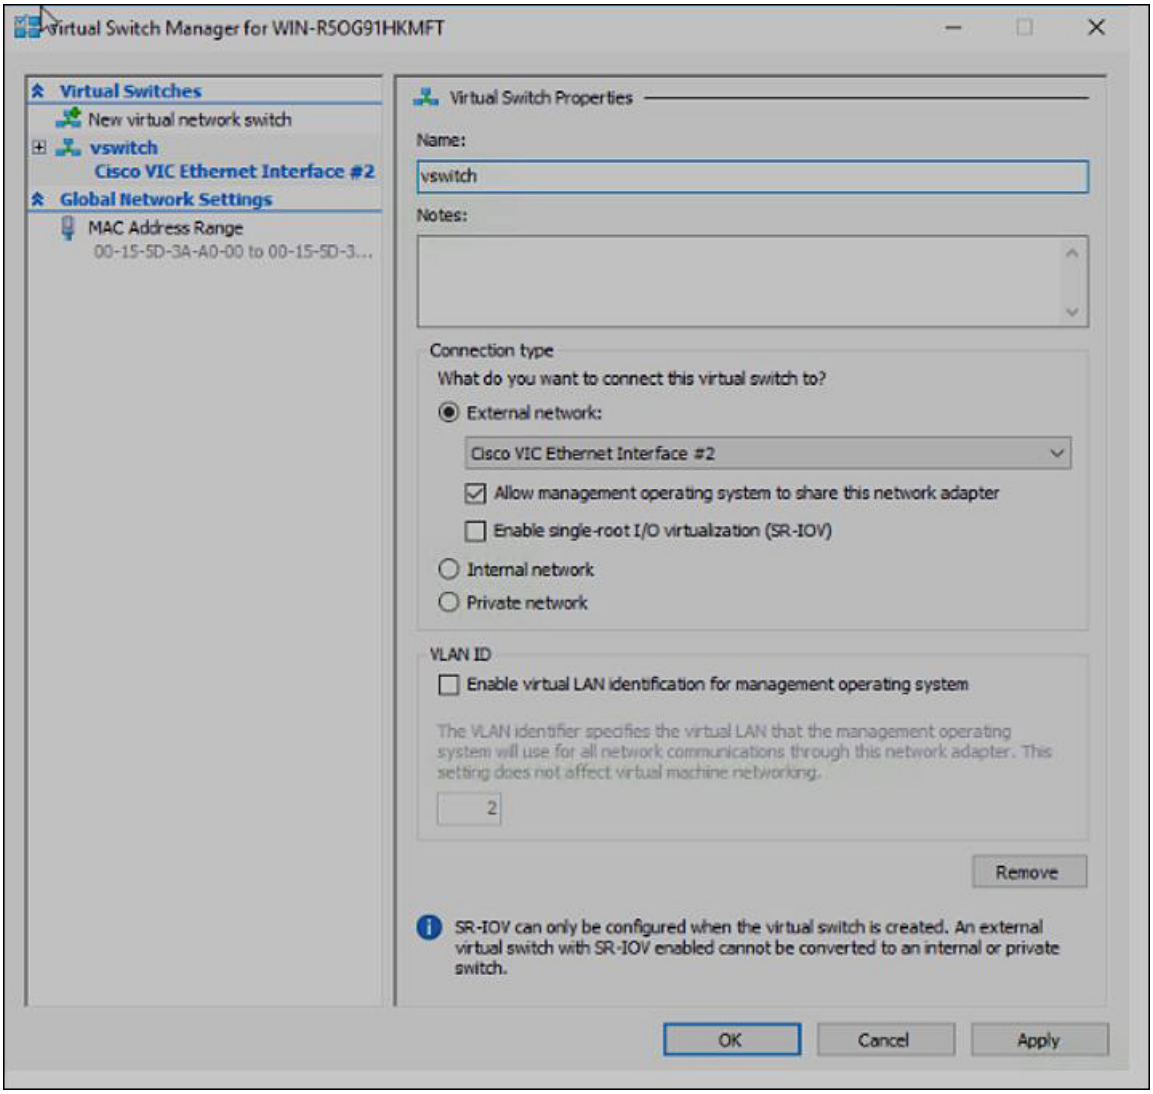
Procedure

Step 1 Go the Hyper-V switch manager.

Step 2 Create a new Virtual Network Switch (vswitch) for the RoCE v2-enabled Ethernet interface.

a) Choose **External Network** and select **VIC Ethernet Interface 2** and **Allow management operating system to share this network adapter**.

b) Click **OK** to create the create the virtual switch.



Bring up the Powershell interface.

Step 3 Configure the non-default vport and enable RDMA with the following Powershell commands:

```
add-vmNetworkAdapter -switchname vswitch -name vp1 -managementOS
enable-netAdapterRdma -name "vEthernet (vp1)"
```

```
PS C:\Users\Administrator>
PS C:\Users\Administrator> add-vmNetworkAdapter -switchName vswitch -name vp1 -managementOS
PS C:\Users\Administrator> enable-netAdapterRdma -name "vEthernet (vp1)"
PS C:\Users\Administrator>
```

a) Configure set-switch using the following Powershell command.

```
new-vmswitch -name setswitch -netAdapterName "Ethernet x" -enableEmbeddedTeam $true
```

This creates the switch. Use the following to display the interfaces:

```
get-netadapterrdma
```

```
add-vmNetworkAdapter -switchname setswtch -name svp1
```

You will see the new vport when you again enter

```
get-netadapterrdma
```

b) Add a vport.

```
add-vmNetworkAdapter -switchname setswtch -name svp1
```

You will see the new vport when you again enter

```
get-netadapterrdma
```

c) Enable the RDMA on the vport:

```
enable-netAdapterRdma -name "vEthernet (svp1)"
```

Step 4 Configure the IPV4 addresses on the RDMA enabled vport in both servers.

Step 5 Create a share in smb-server and map the share in the smb-client.

a) For smb-client and smb-server in the host system, configure the RoCE v2-enabled vNIC as described above.

b) Configure the IPV4 addresses of the primary fabric and sub-vNICs in both servers, using the same IP subnet and same unique vlan for both.

c) Create a share in smb-server and map the share in the smb-client.

Step 6 Verify the Mode 2 configuration.

a) Use the Powershell command ***netstat -xan*** to display listeners and their associated IP addresses.

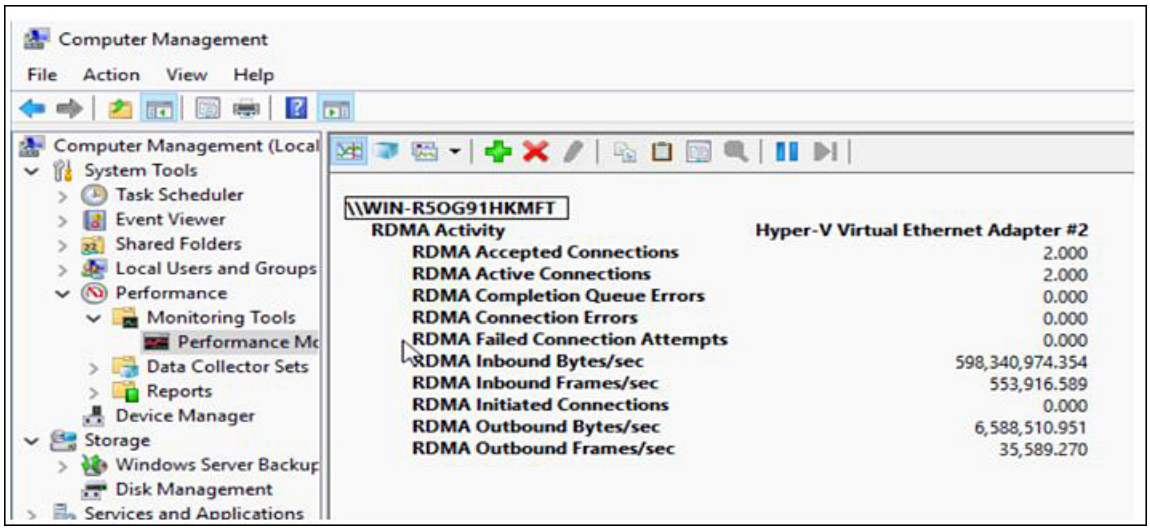
```
PS C:\Users\Administrator>
PS C:\Users\Administrator> netstat -xan

Active NetworkDirect Connections, Listeners, SharedEndpoints

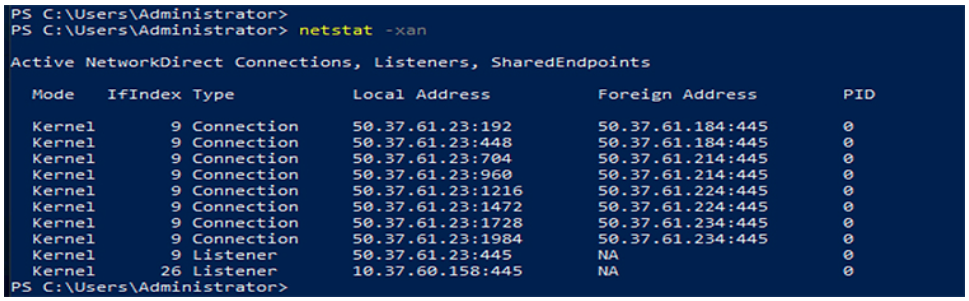
  Mode   IfIndex Type            Local Address          Foreign Address        PID

  Kernel       9 Listener        50.37.61.23:445        NA                     0
  Kernel      26 Listener        10.37.60.158:445       NA                     0
PS C:\Users\Administrator>
```

b) Start any RDMA I/O in the file share in smb-client.



c) Issue the ***netstat -xan*** command again and check for the connection entries to verify they are displayed.



What to do next

Troubleshoot any items if necessary.

Deleting the RoCE v2 Interface in Cisco Intersight

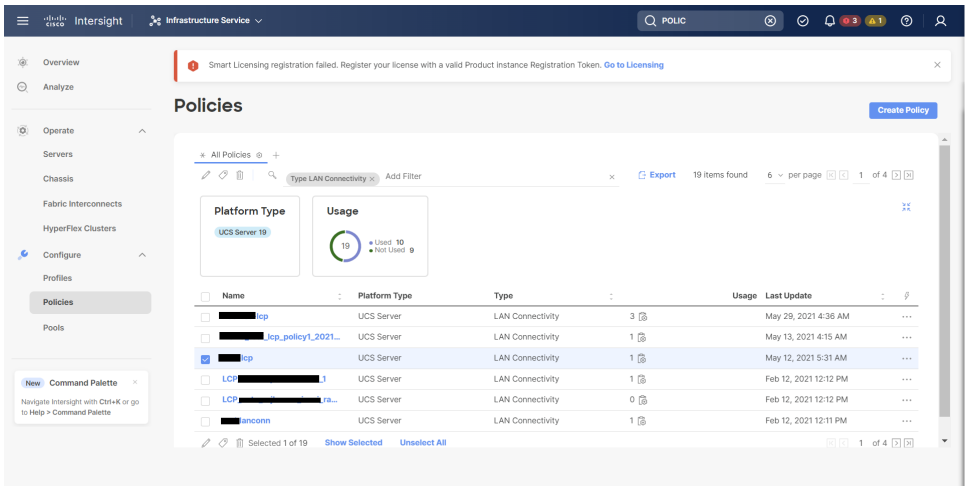
Use these steps to remove the RoCE v2 interface.

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. In the **Add Filter** field, select **Type: LAN Connectivity**.

Step 2 Select the appropriate LAN Connectivity policy created for RoCE V2 configuration and use the delete icon on the top or bottom of the policy list.

Step 3 Click **Delete** to delete the policy.



Step 4 Upon deleting the RoCE v2 configuration, re-deploy the server profile and reboot the server.

Configuring NVMe over Fabrics (NVMeoF) with RoCE v2 in Linux

Configuring RoCE v2 for NVMeoF on Cisco Intersight

Use these steps to configure the RoCE v2 interface on Cisco Intersight.

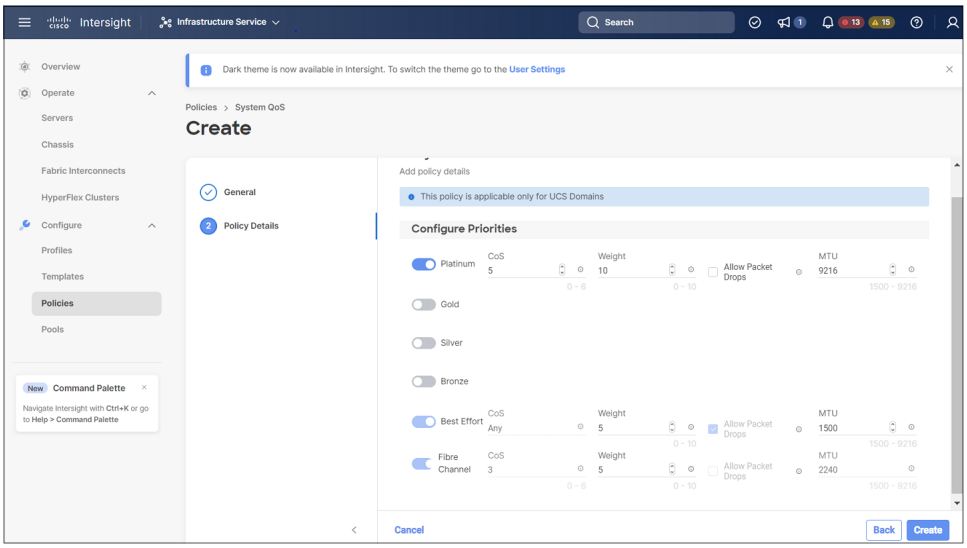
To avoid possible RDMA packet drops, ensure same the no-drop COS is configured across the network. The following steps allow you to configure a no-drop class in System QoS policies and use it for RDMA supported interfaces.

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Domain** platform type, search or choose **System QoS**, and click **Start**.

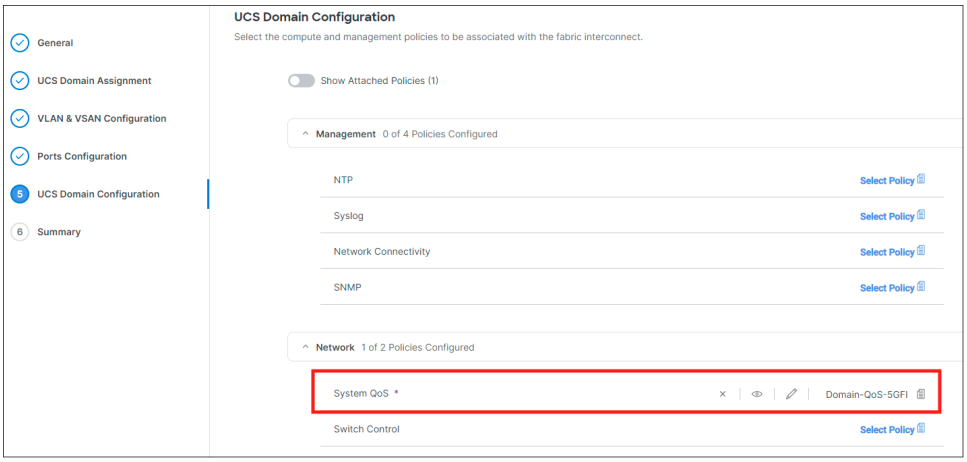
Step 2 In the **General** page, enter the policy name and click **Next**, and then in the **Policy Details** page, configure the property setting for System QoS policy as follows:

- For **Priority**, choose **Platinum**
- For **Allow Packet Drops**, uncheck the check box.
- For **MTU**, set the value as **9216**.



Step 3 Click **Create**.

Step 4 Associate the System QoS policy to the Domain Profile.



Note
For more information, see *Creating System QoS Policy* in Configuring Domain Policies and Configuring Domain Profiles.

The System QoS Policy is successfully created and deployed to the Domain Profile.

What to do next

Configure the server profile with RoCE v2 vNIC settings in LAN Connectivity policy.

Enabling RoCE Settings in LAN Connectivity Policy

Use these steps to configure the RoCE v2 vNIC settings in Mode 1. In Cisco Intersight LAN Connectivity policy, you can enable the RoCE settings on **Ethernet QoS** policy and **Ethernet Adapter** policy for Mode 1 configuration as follows:

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **LAN Connectivity** policy, and click **Start**.

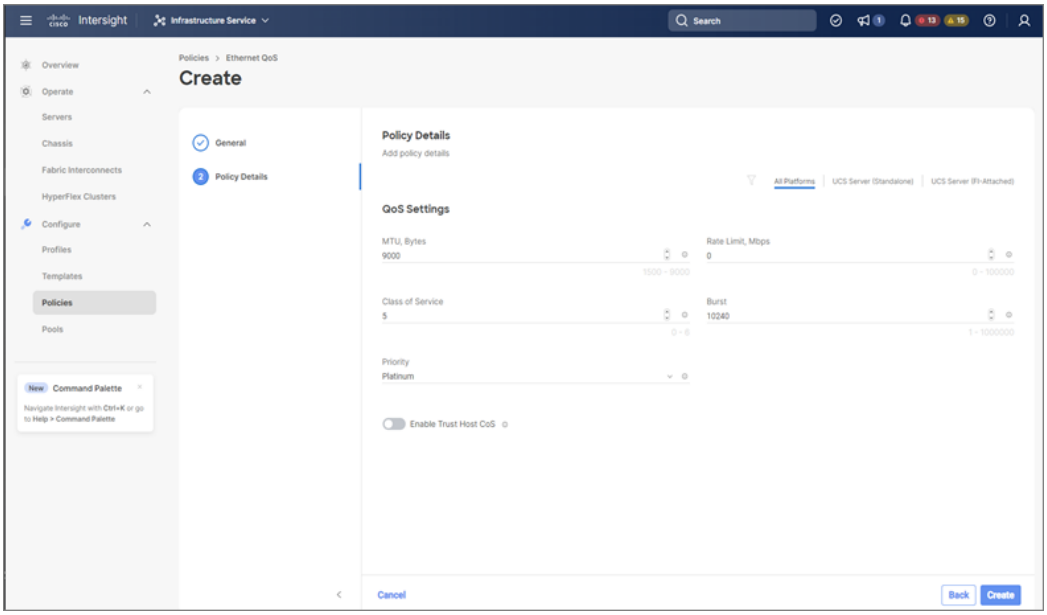
Step 2 In the policy **General** page, enter the policy name, select the Target Platform as **UCS Server (Standalone)** or **UCS Server (FI-Attached)**, and click **Next**.

Step 3 In the **Policy Details** page, click **Add vNIC** to create a new vNIC.

Step 4 In the **Add vNIC** page, follow the configuration parameters to enable the RoCE vNIC settings:

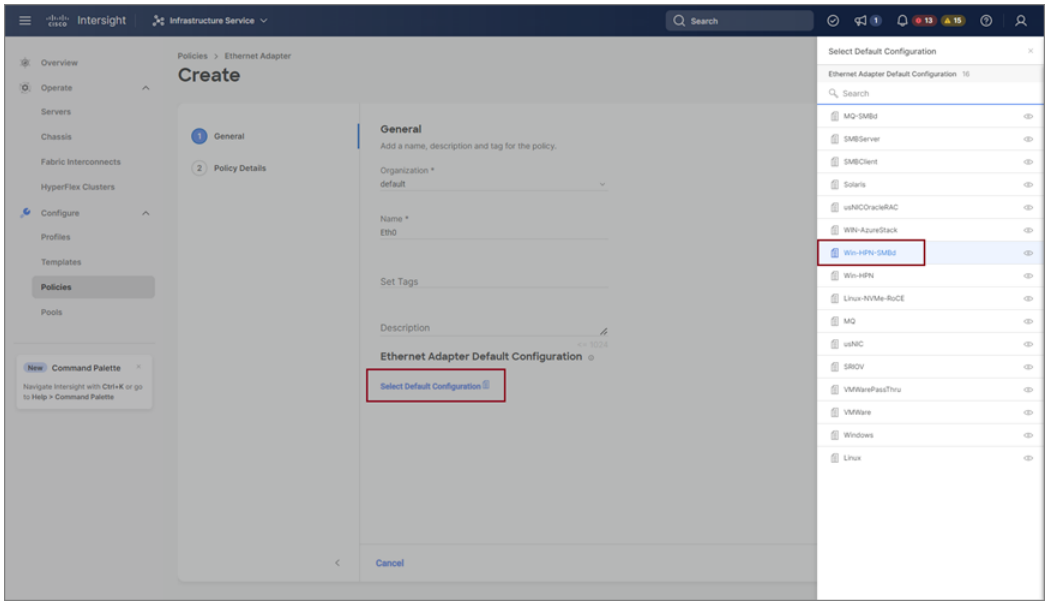
- In the **General** section, provide a name for virtual ethernet interface.
- In the **Consistent Device Naming (CDN)** section of the Standalone server or the **Failover** section of FI-attached server, do the following:
  - Click **Select Policy** link below the **Ethernet QoS**. Use the **Create New** button to create a new Ethernet QoS policy with the following property settings:
    - For **MTU**, choose or enter **1500, 4096, 9000, or 9158**
    - For **Priority**, choose **Platinum** or **any no-drop**
    - For **Class of Service**, choose or enter **5**

      Note
      This property is available only on Standalone servers.

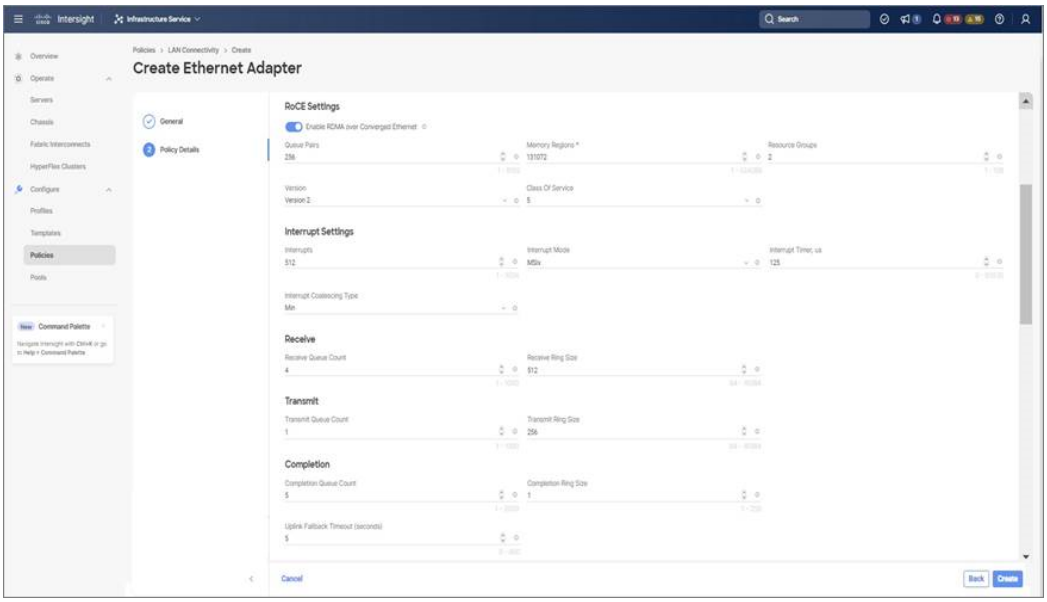


- Click **Select Policy** link below the **Ethernet Adapter**. Follow on to click Create an Ethernet Adapter Policy:
  - **Use the Default Configuration**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy and under Ethernet Adapter Default Configuration click **Select Default**

Configuration to search and select **Win-HPN-SMBd**, the pre-defined Ethernet Adapter Default Configuration. Click **Next** and then **Create**.



- **Configure RoCE Settings in the policy**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy. Under Policy Details page on right pane, use the following property settings, then click **Next**, and then **Create**..
  - For **Enable RDMA over Converged Ethernet**, slide to enable.
  - For **Queue Pairs**, choose or enter **256**
  - For **Memory Regions**, choose or enter **131072**
  - For **Resource Groups**, choose or enter **2**
  - For **Version**, select **Version 2**



- Click **Add** to add and save the new vNIC settings.

Note
All the fields with * are mandatory for creating LAN Connectivity Policy. Ensure they are filled out or selected with appropriate policies.

Step 5 Click **Create** to complete the LAN Connectivity policy with RoCE v2 property settings.

Step 6 Associate the LAN Connectivity policy to the server profile and deploy.

Note
For more information, see *Creating a LAN Connectivity Policy*, *Creating an Ethernet QoS Policy*, and *Creating an Ethernet Adapter Policy* in Configuring UCS Server Policies and Configuring UCS Server Profiles.

The LAN Connectivity policy with the Ethernet QoS policy and Ethernet Adapter policy vNIC setting is successfully created and the server profile is deployed to enable RoCE v2 configuration.

What to do next

Once the policy configuration for RoCE v2 is complete, proceed to enable IOMMU in the BIOS policy.

Enabling an IOMMU BIOS Settings

Use the following steps to configure the server profile with the RoCE v2 vNIC and enable the IOMMU BIOS policy before enabling the IOMMU in the Linux kernel.

Procedure

Step 1 Navigate to **CONFIGURE** > **Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **BIOS**, and click **Start**.

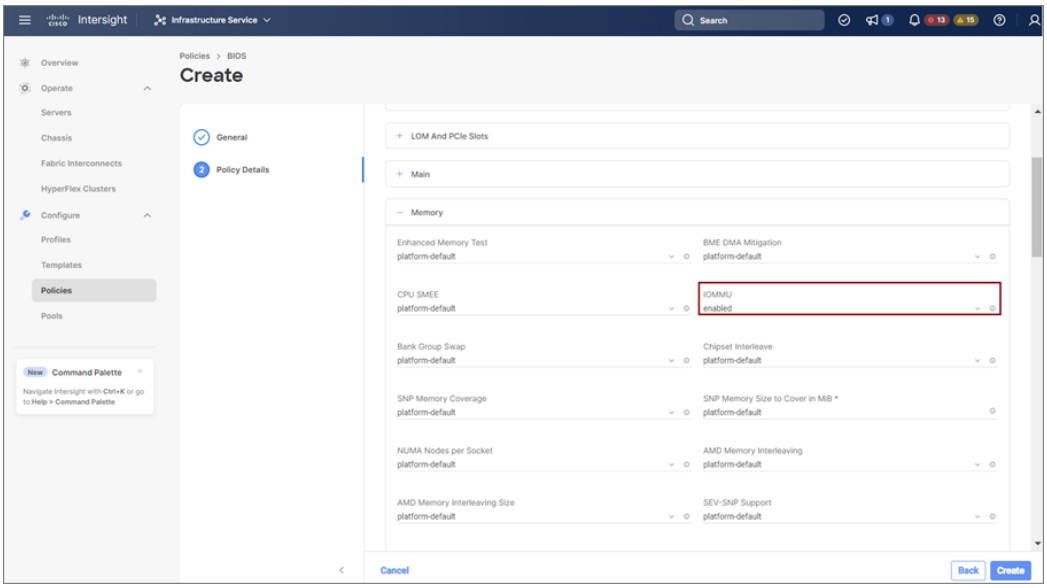
Step 2 On the **General** page, enter the policy name and click **Next**.

Step 3 On the **Policy Details** page, configure the following BIOS:

a) Select **All Platforms**.

b) For a server with an Intel CPU, enable **Intel VT for Directed I/O** under the **Intel Directed I/O** drop-down list and enable **Intel(R) VT** under the **Processor** drop-down list.

c) For a server with an AMD CPU, enable **IOMMU** under the **Memory** drop-down list and enable **SVM Mode** under the **Processor** drop-down list.



Step 4 Click **Create**.

Step 5 Associate the BIOS policy to the server profile and reboot the server.

Note
For more information, see *Creating a BIOS Policy* in Configuring Server Policies and Configuring Server Profile.

The BIOS Policy is successfully created and deployed on the server profile.

What to do next

Configure RoCE v2 for NVMeoF on the Host System.

Configuring RoCE v2 for NVMeoF on the Host System

Before you begin

Configure the Server Profile with RoCE v2 vNIC and the IOMMU enabled BIOS policy.

Procedure

Step 1 Open the `/etc/default/grub` file for editing.

Step 2 Add `intel_iommu=on` to the end of `GRUB_CMDLINE_LINUX`.

```
sample /etc/default/grub configuration file after adding intel_iommu=on:
# cat /etc/default/grub
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap biosdevname=1 rhgb quiet
 intel_iommu=on
GRUB_DISABLE_RECOVERY="true"
```

Step 3 After saving the file, generate a new grub.cfg file.

For Legacy boot:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

For UEFI boot:

```
# grub2-mkconfig -o /boot/grub2/efi/EFI/redhat/grub.cfg
```

Step 4 Reboot the server. You must reboot your server for the changes to take after enabling IOMMU.

Step 5 Verify the server is booted with `intel_iommu=on` option.

```
cat /proc/cmdline | grep iommu
```

Note its inclusion at the end of the output.

```
[root@localhost basic-setup]# cat /proc/cmdline | grep iommu
BOOT_IMAGE=/vmlinuz-3.10.0-957.27.2.el7.x86_64 root=/dev/mapper/rhel-root ro crashkernel=auto
rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet intel_iommu=on LANG=en_US.UTF-8
```

What to do next

Download the enic and enic_rdma drivers.

Installing Cisco enic and enic_rdma Drivers

The enic_rdma driver requires enic driver. When installing enic and enic_rdma drivers, download and use the matched set of enic and enic_rdma drivers on Cisco.com. Attempting to use the binary enic_rdma driver downloaded from Cisco.com with an inbox enic driver, will not work.

Procedure

Step 1 Install the enic and enic_rdma rpm packages:

```
# rpm -ivh kmod-enic-<version>.x86_64.rpm kmod-enic rdma-<version>.x86_64.rpm
```

Note
During enic_rdma installation, the enic_rdmalibnvdimm module may fail to install on RHEL 7.7 because the `nvdimm-security.conf` dracut module needs spaces in the `add_drivers` value. For workaround, please follow the instruction from the following links:

https://access.redhat.com/solutions/4386041

https://bugzilla.redhat.com/show_bug.cgi?id=1740383

Step 2 The enic_rdma driver is now installed but not loaded in the running kernel. Reboot the server to load enic_rdma driver into the running kernel.

Step 3 Verify the installation of enic_rdma driver and RoCE v2 interface:

```
[root@localhost ~]# dmesg | grep enic_rdma
[    3.137083] enic_rdma: Cisco VIC Ethernet NIC RDMA Driver, ver 1.2.0.28-877.2
2 init
[    3.242663] enic 0000:1b:00.1 eno6: enic_rdma: FW v3 RoCEv2 enabled
[    3.284856] enic 0000:1b:00.4 eno9: enic_rdma: FW v3 RoCEv2 enabled
[   16.441662] enic 0000:1b:00.1 eno6: enic_rdma: Link UP on enic_rdma_0
[   16.458754] enic 0000:1b:00.4 eno9: enic_rdma: Link UP on enic_rdma_1
```

Step 4 Load the nvme-rdma kernel module:

```
# modprobe nvme-rdma
```

After server reboot, nvme-rdma kernel module is unloaded. To load nvme-rdma kernel module every server reboot, create nvme_rdma.conf file using:

```
# echo nvme_rdma > /etc/modules-load.d/nvme_rdma.conf
```

Note
For more information about enic_rdma after installation, use the **`rpm -q -l kmod-enic_rdma`** command to extract the README file.

What to do next

Discover targets and connect to NVMe namespaces. If your system needs multipath access to the storage, go to the section for Setting Up Device Mapper Multipath, on page 37.

Discovering the NVMe Target

Use this procedure to discover the NVMe target and connect NVMe namespaces.

Before you begin

Install **nvme-cli** version 1.6 or later if it is not installed already.

Configure the IP address on the RoCE v2 interface and make sure the interface can ping the target IP.

Procedure

Step 1 Create an nvme folder in /etc, then manually generate host nqn.

```
# mkdir /etc/nvme
# nvme gen-hostnqn > /etc/nvme/hostnqn
```

Step 2 Create a settos.sh file and run the script to set priority flow control (PFC) in IB frames.

Note
To avoid failure of sending NVMeoF traffic, you *must* create and run this script after *every* server reboot.

```
# cat settos.sh
#!/bin/bash
for f in `ls /sys/class/infiniband`;
do
        echo "setting TOS for IB interface:" $f
        mkdir -p /sys/kernel/config/rdma_cm/$f/ports/1
        echo 186 > /sys/kernel/config/rdma_cm/$f/ports/1/default_roce_tos
done
```

Step 3 Discover the NVMe target by entering the following command.

```
nvme discover --transport=rdma --traddr=<IP address of transport target port>
```

For example, to discover the target at 50.2.85.200:

```
# nvme discover --transport=rdma --traddr=50.2.85.200

Discovery Log Number of Records 1, Generation counter 2
=====Discovery Log Entry 0======
trtype:  rdma
adrfam:  ipv4
subtype: nvme subsystem
treq:    not required
portid:  3
trsvcid: 4420
subnqn:  nqn.2010-06.com.purestorage:flasharray.9a703295ee2954e
traddr:  50.2.85.200
rdma_prtype: roce-v2
rdma_qptype: connected
rdma_cms:    rdma-cm
rdma_pkey: 0x0000
```

Note
To discover the NVMe target using IPv6, put the IPv6 target address next to the `traddr` option.

Step 4 Connect to the discovered NVMe target by entering the following command.

```
nvme connect --transport=rdma --traddr=<IP address of transport target port>> -n <subnqn value from
 nvme discover>
```

For example, to discover the target at 50.2.85.200 and the subnqn value found above:

```
# nvme connect --transport=rdma --traddr=50.2.85.200 -n
nqn.2010-06.com.purestorage:flasharray.9a703295ee2954e
```

Note
To connect to the discovered NVMe target using IPv6, put the IPv6 target address next to the `traddr` option.

Step 5 Use the **nvme list** command to check mapped namespaces:

```
# nvme list
Node             SN                   Model                                    Namespace Usage
                 Format           FW Rev
---------------- -------------------- ---------------------------------------- ---------
-------------------------- ---------------- --------
/dev/nvme0n1     09A703295EE2954E     Pure Storage FlashArray                  72656       4.29  GB
 /   4.29  GB    512   B +  0 B   99.9.9
/dev/nvme0n2     09A703295EE2954E     Pure Storage FlashArray                  72657       5.37  GB
 /   5.37  GB    512   B +  0 B   99.9.9
```

Setting Up Device Mapper Multipath

If your system is configured with Device Mapper multipathing (DM Multipath), use the following steps to set up Device Mapper multipath.

Procedure

Step 1 Install the `device-mapper-multipath` package if it is not installed already

Step 2 Enable and start multipathd:

```
# mpathconf --enable --with_multipathd y
```

Step 3 Edit the etc/multipath.conf file to use the following values :

```
defaults {
        polling_interval        10
        path_selector           "queue-length 0"
        path_grouping_policy    multibus
        fast_io_fail_tmo        10
        no_path_retry           0
        features                0
        dev_loss_tmo            60
        user_friendly_names     yes
}
```

Step 4 Flush with the updated multipath device maps.

```
# multipath -F
```

Step 5 Restart multipath service:

```
# systemctl restart multipathd.service
```

Step 6 Rescan multipath devices:

```
# multipath -v2
```

Step 7 Check the multipath status:

```
# multipath -ll
```

Deleting the RoCE v2 Interface in Cisco Intersight

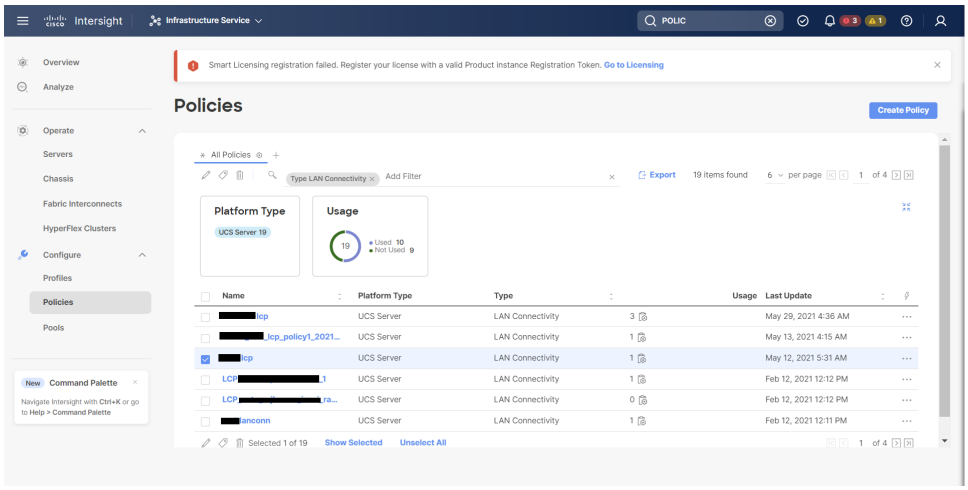
Use these steps to remove the RoCE v2 interface.

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. In the **Add Filter** field, select **Type: LAN Connectivity**.

Step 2 Select the appropriate LAN Connectivity policy created for RoCE V2 configuration and use the delete icon on the top or bottom of the policy list.

Step 3 Click **Delete** to delete the policy.



Step 4 Upon deleting the RoCE v2 configuration, re-deploy the server profile and reboot the server.

Configuring NVMe with RoCEv2 in ESXi

Configuring RoCE v2 for NVMeoF on Cisco Intersight

Use these steps to configure the RoCE v2 interface on Cisco Intersight.

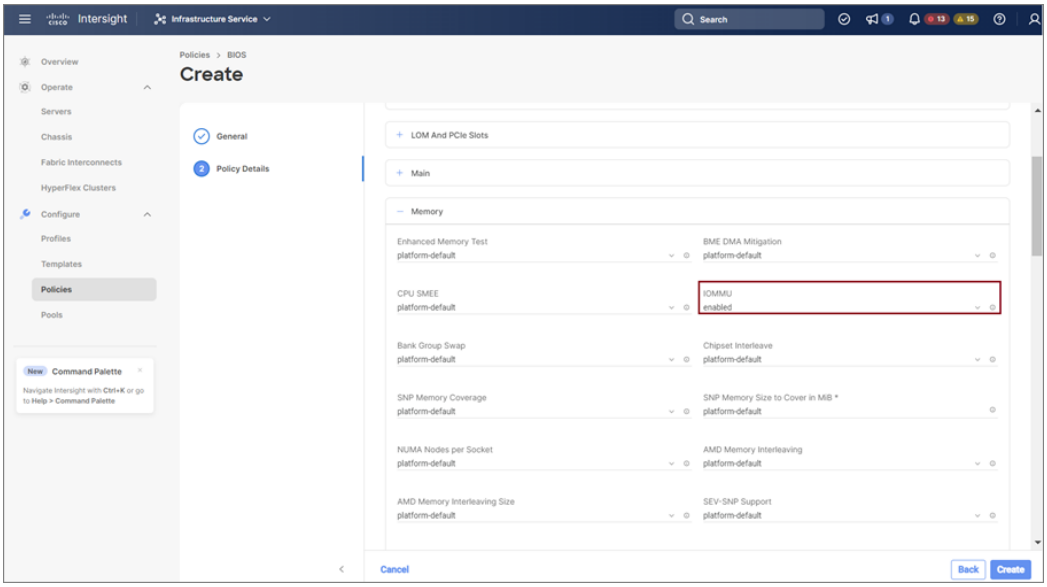
To avoid possible RDMA packet drops, ensure same no-drop COS is configured across the network. The following steps allows you to configure a no-drop class in System QoS policies and use it for RDMA supported interfaces.

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Domain** platform type, search or choose **System QoS**, and click **Start**.

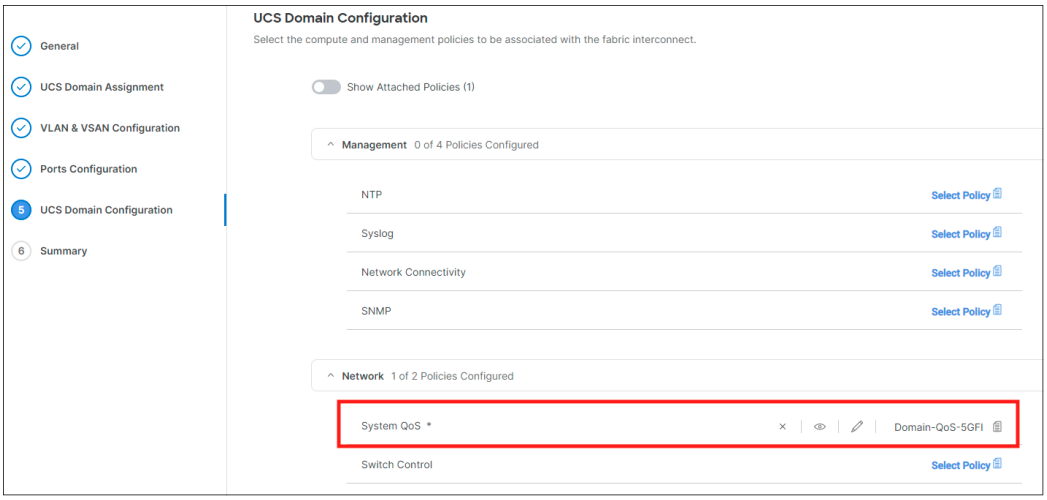
Step 2 In the **General** page, enter the policy name and click **Next**, and then in the **Policy Details** page, configure the property setting for System QoS policy as follows:

- For **Priority**, choose **Platinum**
- For **Allow Packet Drops**, uncheck the check box.
- For **MTU**, set the value as **9216**.



Step 3 Click **Create**.

Step 4 Associate the System QoS policy to the Domain Profile.



Note
For more information, see *Creating System QoS Policy* in Configuring Domain Policies and Configuring Domain Profiles.

The System QoS Policy is successfully created and deployed to the Domain Profile.

What to do next

Configure the server profile with RoCE v2 vNIC settings in LAN Connectivity policy.

Enabling RoCE Settings in LAN Connectivity Policy

Use the following steps to configure the RoCE v2 vNIC. In Cisco Intersight LAN Connectivity policy, you can enable the RoCE settings on **Ethernet Adapter policy** for Linux configuration as follows:

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **LAN Connectivity policy**, and click **Start**.

Step 2 In the policy **General** page, enter the policy name, select the Target Platform as **UCS Server (Standalone)** or **UCS Server (FI-Attached)**, and click **Next**.

Step 3 In the **Policy Details** page, click **Add vNIC** to create a new vNIC.

Step 4 In the **Add vNIC** page, follow the configuration parameters to enable the RoCE v2 vNIC:

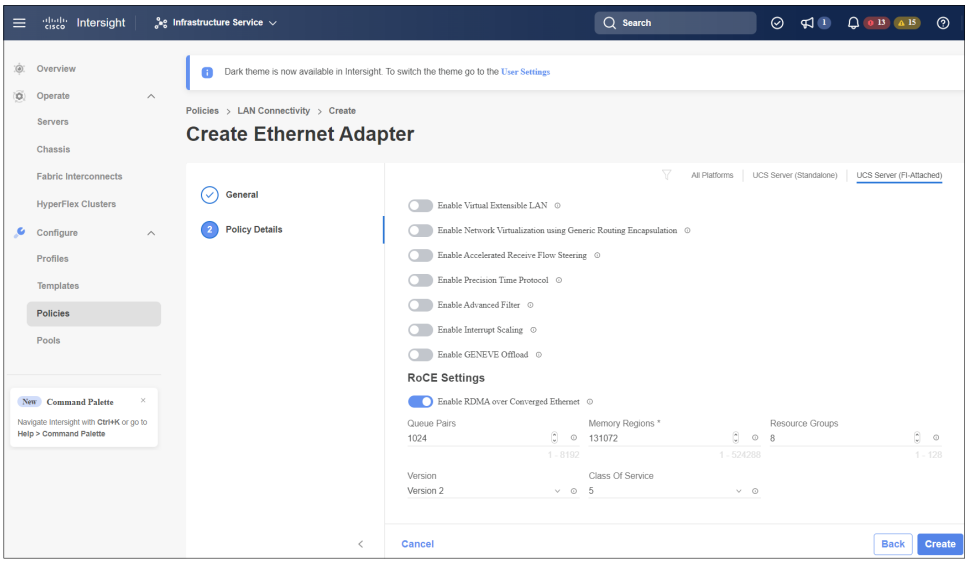
a) In the **General** section, provide a name for virtual ethernet interface.

b) Incase of a Standalone server, click the **Consistent Device Naming (CDN)** or click the **Failover** of a FI-attached server , and do the following:

- Click **Select Policy** under **Ethernet Adapter**.
- In the **Select Policy** window, click **Create New** to create an Ethernet Adapter policy.
- In the **General** page of the Ethernet Adapter Policy, enter the policy name and click **Next**.
- In the **Policy Details** page of the Ethernet Adapter Policy, modify the following property setting:
  - **RoCE Settings**
    - For **Enable RDMA over Converged Ethernet**, slide to enable and set the RoCE on this virtual interface.
    - For **Queue Pairs**, select or enter **1024**
    - For **Memory Regions**, select or enter **131072**
    - For **Resource Groups**, select or enter **8**
    - For **Version**, select **Version 2**
    - For **Class of Service**, select **5**
  - **Interrupt Settings**
    - For **Interrupts**, select or enter **256**.
    - For **Interrupt mode**, select **MSIx**.
    - For **Interrupt Timer, us,** select **125**.
    - For **Interrupt Coalescing Type**, select **Min**.

- **Receive** Settings
  - For **Receive Queue Count**, select or enter **1**.
  - For **Receiving Ring Size**, select or enter **512**.
- **Transmit** Settings
  - For **Transmit Queue Count**, select or enter **1**.
  - For **Transmit Ring Size**, select or enter **256**.
- **Completion** Settings
  - For **Completion Queue Count**, select or enter **2**.
  - For **Completion Ring Size**, select or enter **1**.
  - For **Uplink Failback Timeout(seconds)**, select or enter **5**

- Click **Create** to create an Ethernet Adapter Policy with the above defined settings.



- Click **Add** to save the setting and add the new vNIC.

Note
All the fields with * are mandatory and ensure it is filled out or selected with appropriate policies.

Step 5 Click **Create** to complete the LAN Connectivity policy with RoCE v2 settings.

Step 6 Associate the LAN Connectivity policy to the Server Profile.

Note
For more information, see *Creating a LAN Connectivity Policy* and *Creating an Ethernet Adapter Policy* in Configuring UCS Server Policies and Configuring UCS Server Profiles.

The LAN Connectivity Policy with the Ethernet Adapter policy vNIC setting is successfully created and deployed to enable RoCE v2 configuration.

What to do next

Once the policy configuration for RoCE v2 is complete, configure RoCE v2 for NVMeoF on the Host System.

NENIC Driver Installation

Before you begin

The Ethernet Network Interface Card (eNIC) Remote Direct Memory Access (RDMA) driver requires nenic driver.

Procedure

Step 1 Copy the eNIC vSphere Installation Bundle (VIB) or offline bundle to the ESXi server.

Step 2 Use the command to install nenic driver:

```
esxcli software vib install -v {VIBFILE}
or
esxcli software vib install -d {OFFLINE_BUNDLE}
```

Example:

```
esxcli software vib install -v /tmp/nenic-2.0.4.0-1OEM.700.1.0.15843807.x86_64.vib
```

Note
Depending on the certificate used to sign the VIB, you may need to change the host acceptance level. To do this, use the command:

```
esxcli software acceptance set --level=<level>
```

Depending on the type of VIB installed, you may need to put ESX into maintenance mode. This can be done through the client, or by adding the *--maintenance-mode* option to the above *esxcli*.

What to do next

Configure the Host side for ESXi NVMe RDMA.

ESXi NVMe RDMA Host Side Configuration

NENIC RDMA Functionality

One of the major difference between RDMA on Linux and ESXi is listed below:

- In ESXi, the physical interface (vmnic) MAC is not used for RoCEv2 traffic. Instead, the VMkernel port (vmk) MAC is used.

Outgoing RoCE packets use the vmk MAC in the Ethernet source MAC field, and incoming RoCE packets use the vmk MAC in the Ethernet destination mac field. The vmk MAC address is a VMware MAC address assigned to the vmk interface when it is created.

- In Linux, the physical interface MAC is used in source MAC address field in the ROCE packets. This Linux MAC is usually a Cisco MAC address configured to the VNIC using UCS Manager.

If you ssh into the host and use the **esxcli network ip interface list** command, you can see the MAC address.

```
vmk0
   Name: vmk0
   MAC Address: 2c:f8:9b:a1:4c:e7
   Enabled: true
   Portset: vSwitch0
   Portgroup: Management Network
   Netstack Instance: defaultTcpipStack
   VDS Name: N/A
   VDS UUID: N/A
   VDS Port: N/A
   VDS Connection: -1
   Opaque Network ID: N/A
   Opaque Network Type: N/A
   External ID: N/A
   MTU: 1500
   TSO MSS: 65535
   RXDispQueue Size: 2
   Port ID: 67108881
```

You must create a vSphere Standard Switch to provide network connectivity for hosts, virtual machines, and to handle VMkernel traffic. Depending on the connection type that you want to create, you can create a new vSphere Standard Switch with a VMkernel adapter, only connect physical network adapters to the new switch, or create the switch with a virtual machine port group.

Create Network Connectivity Switches

Use these steps to create a vSphere Standard Switch to provide network connectivity for hosts, virtual machines, and to handle VMkernel traffic.

Before you begin

Ensure you have nenic drivers. Download and install nenic drivers before proceeding with below steps:

Procedure

Step 1 In the vSphere Client, navigate to the host.

Step 2 On the **Configure** tab, expand **Networking** and select **Virtual Switches.**

Step 3 Click on **Add Networking**.

The available network adapter connection types are:

- **Vmkernel Network Adapter**

  Creates a new VMkernel adapter to handle host management traffic

- **Physical Network Adapter**

  Adds physical network adapters to a new or existing standard switch.

- **Virtual Machine Port Group for a Standard Switch**

  Creates a new port group for virtual machine networking.

Step 4 Select connection type **Vmkernel Network Adapter**.

Step 5 Select **New Standard Switch** and click **Next**.

Step 6 Add physical adapters to the new standard switch.

a) Under **Assigned Adapters**, select **New Adapters**.

b) Select one or more adapters from the list and click **OK**. To promote higher throughput and create redundancy, add two or more physical network adapters to the Active list.

c) (Optional) Use the up and down arrow keys to change the position of the adapter in the Assigned Adapters list.

d) Click **Next**.

Step 7 For the new standard switch you just created for the VMadapter or a port group, enter the connection settings for the adapter or port group.

a) Enter a label that represents the traffic type for the VMkernel adapter.

b) Set a VLAN ID to identify the VLAN the VMkernel uses for routing network traffic.

c) Select IPV4 or IPV6 or both.

d) Select an MTU size from the drop-down menu. Select Custom if you wish to enter a specific MTU size. The maximum MTU size is 9000 bytes.

Note
You can enable Jumbo Frames by setting an MTU greater than 1500.

e) After setting the TCP/IP stack for the VMkernel adapter, select a TCP/IP stack.

To use the default TCP/IP stack, select it from the available services.

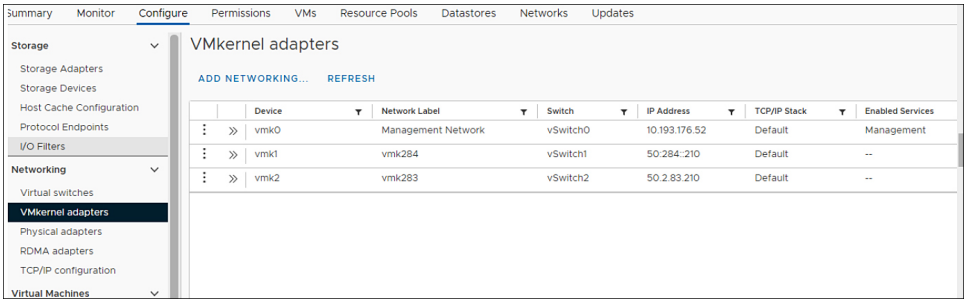
Note
Be aware that the TCP/IP stack for the VMkernel adapter cannot be changed later.

f) Configure IPV4 and/or IPV6 settings.

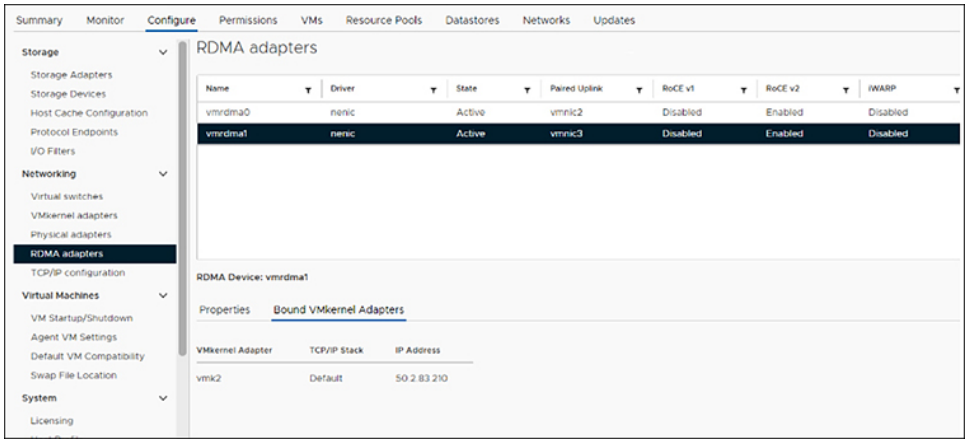
Step 8 On the **Ready to Complete** page, click **Finish**.

Step 9 Check the VMkernel ports for the VM Adapters or port groups with NVMe RDMA in the vSphere client, as shown in the Results below.

The VMkernel ports for the VM Adapters or port groups with NVMe RDMA are shown below.



The VRDMA Port groups created with NVMeRDMA supported vmnic appear as below.



What to do next

Create vmhba ports on top of vmrdma ports.

Create VMVHBA Ports in ESXi

Use the following steps for creating vmhba ports on top of the vmrdma adapter ports.

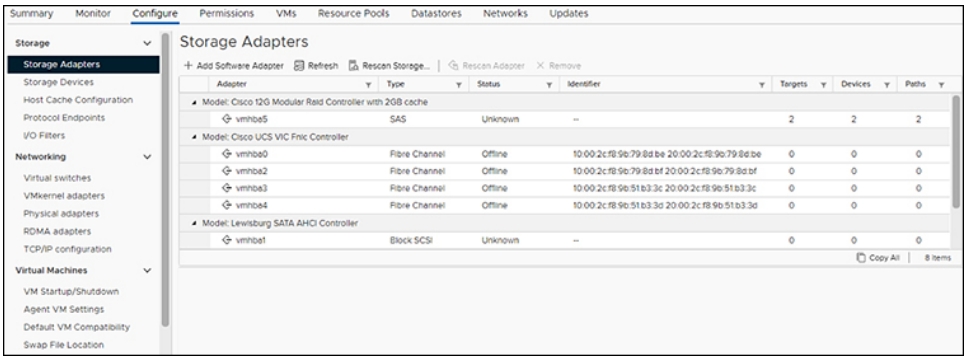
Before you begin

Create the adapter ports for storage connectivity.

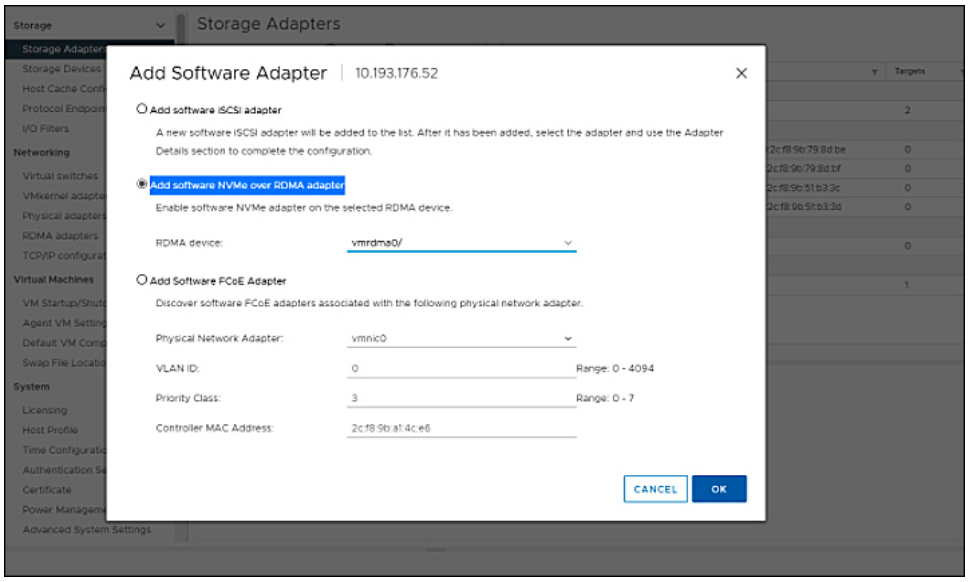
Procedure

Step 1 Go to vCenter where your ESXi host is connected.

Step 2 Click on **Host**>**Configure**>**Storage adapters**.



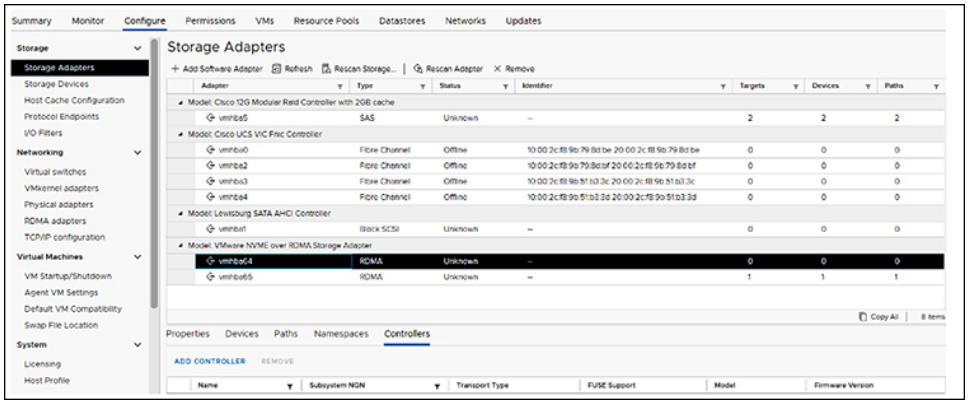
Step 3 Click **+Add Software Adapter**. The following dialog box will appear.



Step 4 Select **Add software NVMe over RDMA adapter** and the vmrdma port you want to use.

Step 5 Click **OK**

The vmhba ports for the VMware NVMe over RDMA storage adapter will be shown as in the example below



Displaying vmnic and vmrdma Interfaces

ESXi creates a vmnic interface for each nenic VNIC configured to the host.

Before you begin

Create Network Adapters and VHBA ports.

Procedure

Step 1 Use **ssh** to access the host system.

Step 2 Enter **esxcfg-nics -l** to list the vmnics on ESXi.

```
Name    PCI          Driver      Link Speed      Duplex MAC Address       MTU    Description
vmnic0  0000:3b:00.0 ixgben      Down 0Mbps      Half   2c:f8:9b:a1:4c:e6 1500   Intel(R) Ethernet Controller X550
vmnic1  0000:3b:00.1 ixgben      Up   1000Mbps   Full   2c:f8:9b:a1:4c:e7 1500   Intel(R) Ethernet Controller X550
vmnic2  0000:1d:00.0 nenic       Up   50000Mbps  Full   2c:f8:9b:79:8d:bc 1500   Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic3  0000:1d:00.1 nenic       Up   50000Mbps  Full   2c:f8:9b:79:8d:bd 1500   Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic4  0000:63:00.0 nenic       Down 0Mbps      Half   2c:f8:9b:51:b3:3a 1500   Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic5  0000:63:00.1 nenic       Down 0Mbps      Half   2c:f8:9b:51:b3:3b 1500   Cisco Systems Inc Cisco VIC Ethernet NIC
```

esxcli network nic list

```
Name    PCI Device    Driver  Admin Status  Link Status  Speed  Duplex  MAC Address         MTU  Description
------  ------------  ------  ------------  -----------  -----  ------  -----------------   ---  -----------
vmnic0  0000:3b:00.0  ixgben  Up            Down             0  Half    2c:f8:9b:a1:4c:e6  1500  Intel(R) Ethernet Controller X550
vmnic1  0000:3b:00.1  ixgben  Up            Up            1000  Full    2c:f8:9b:a1:4c:e7  1500  Intel(R) Ethernet Controller X550
vmnic2  0000:1d:00.0  nenic   Up            Up           50000  Full    2c:f8:9b:79:8d:bc  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic3  0000:1d:00.1  nenic   Up            Up           50000  Full    2c:f8:9b:79:8d:bd  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic4  0000:63:00.0  nenic   Up            Down             0  Half    2c:f8:9b:51:b3:3a  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
vmnic5  0000:63:00.1  nenic   Up            Down             0  Half    2c:f8:9b:51:b3:3b  1500  Cisco Systems Inc Cisco VIC Ethernet NIC
```

Step 3 Use **esxcli rdma device list** to list the vmrdma devices. When the enic driver registers with ESXi the RDMA device for a RDMA capable VNIC, ESXi creates a vmrdma device and links it to the corresponding vmnic.

```
[root@S       RackServer:~] esxcli rdma device list
Name     Driver  State   MTU  Speed    Paired Uplink  Description
-------  ------  ------  ----  -------  -------------  -----------
vmrdma0  nenic   Active  4096  50 Gbps  vmnic1         Cisco UCS VIC 15XXX (A0)
vmrdma1  nenic   Active  4096  50 Gbps  vmnic2         Cisco UCS VIC 15XXX (A0)
[root@S       RackServer:~] esxcli rdma device vmknic list
Device   Vmknic  NetStack
-------  ------  --------
vmrdma0  vmk1    defaultTcpipStack
vmrdma1  vmk2    defaultTcpipStack
```

Step 4 Use **esxcli rdma device protocol list** to check the protocols supported by the vmrdma interface.

For enic, RoCE v2 is the only protocol supported from this list. The output of this command should match the RoCEv2 configuration on the VNIC.

Step 5 Use **esxcli nvme adapter list** to list the NVMe adapters and the vmrdma and vmnic interfaces it is configured on.

```
[root@ESXi7U3      ~] esxcli nvme adapter list
Adapter  Adapter Qualified Name         Transport Type  Driver    Associated Devices
-------  -----------------------------  --------------  --------  ------------------
vmhba64  aqn:nvmerdma:2c-f8-9b-79-8d-bc  RDMA            nvmerdma  vmrdma0, vmnic2
vmhba65  aqn:nvmerdma:2c-f8-9b-79-8d-bd  RDMA            nvmerdma  vmrdma1, vmnic3
[root@ESXi7U3      :~]
```

Step 6 All vmhbas in the system can be listed using **esxcli storage core adapter list**. The vmhba configured over RDMA.

Note
For vmhba64 and vmhba65, you may observe that the driver's Link State displays *link-n/a* instead of *Online*. This is a known issue in ESXi 7.0 Update 3. For more information, see Known Issues - ESXi.

NVMe Fabrics and Namespace Discovery

This process is performed through the ESXi command line interface.

Before you begin

Create and configure the Ethernet Adapter Policy.

Procedure

Step 1 Check and enable Nonvolatile Memory Express (NVMe) on the vmrdma device. If NVMe is enabled, the system returns the following message.

Example:

```
esxcli nvme fabrics enable -p RDMA -d vmrdma0
```

Step 2 Discover the NVMe on the array by entering the following command.

```
esxcli nvme fabrics discover -a vmhba64 -l transport_address figure with esxcli nvme fabrics discover
 -a vmhba64 -l 50.2.84.100
```

The NVMe controller displays the output that include Transport Type, Address Family, Subsystem Type, Controller ID, Admin Queue, Max Size, Transport Address, Transport Service ID, and Subsystem NQN.

You will see output on the NVMe controller.

Step 3 Perform NVMe fabric interconnect.

```
esxcli nvme fabrics discover -a vmhba64 -l transport_address p Transport Service
ID -s Subsystem NQN
```

Step 4 The NVMe controller displays a list of the controllers connected to NVMe. The NVMe namespace list should shows all the NVMe drivers discovered.

```
esxcli nvme fabrics discover -a vmhba64 -l transport_address p Transport Service
ID -s Subsystem NQN
```

The following example shows esxcli discovery commands executed on the server.

Example:

```
[root@ESXiUCSA:~] esxcli nvme fabrics enable -p RDMA -d vmrdma0
NVMe already enabled on vmrdma0 [root@ESXiUCSA:~] esxcli nvme fabrics discover -a vmhba64 -l 50.2.84.100
Transport Type Address Family Subsystem Type Controller ID Admin Queue Max Size Transport
Address Transport Service ID Subsystem NQN
------------- ------------- ------------- ------------- --------------------
----------------- -------------------- -------------
RDMA IPV4 NVM 65535 31 50.2.84.100
4420 nq.210-06.com.purestorage:flasharray:2dp1239anjkl484
[root@ESXiUCSA:~] esxcli nvme fabrics discover -a vmhba64 -l 50.2.84.100 p 4420 -s
nq.210-06.com.purestorage:flasharray:2dp1239anjkl484
Controller already connected
```

Deleting the RoCE v2 Interface in Cisco Intersight

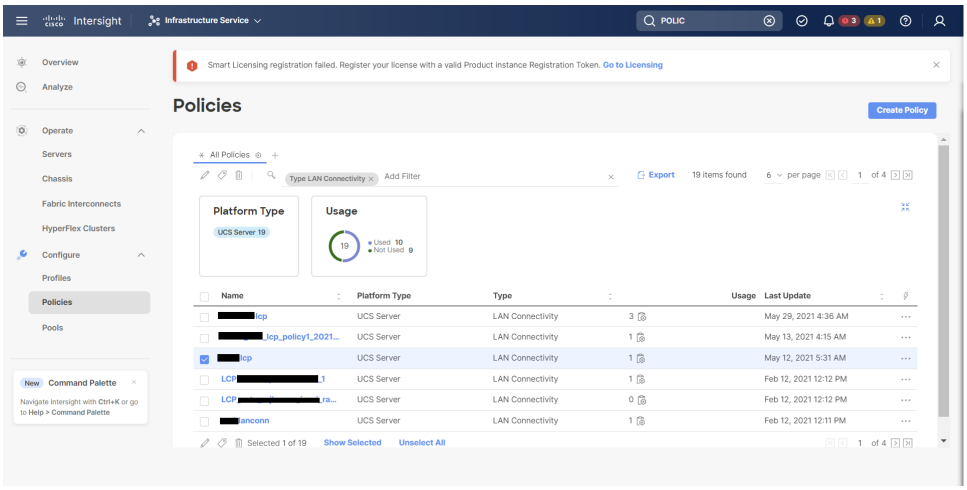
Use these steps to remove the RoCE v2 interface.

Procedure

Step 1 Navigate to **CONFIGURE > Policies**. In the **Add Filter** field, select **Type: LAN Connectivity**.

Step 2 Select the appropriate LAN Connectivity policy created for RoCE V2 configuration and use the delete icon on the top or bottom of the policy list.

Step 3 Click **Delete** to delete the policy.



Step 4 Upon deleting the RoCE v2 configuration, re-deploy the server profile and reboot the server.

Known Issues

Windows

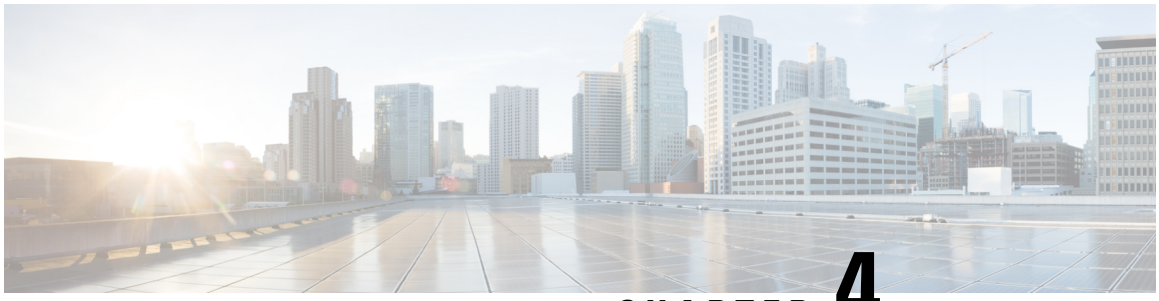
Symptom	Conditions	Workaround
On VIC 1400 Series adapters, the neNIC driver for Windows 2019 can be installed on Windows 2016 and the Windows 2016 driver can be installed on Windows 2019. However, this is an unsupported configuration.	Case 1 : Installing Windows 2019 nenic driver on Windows 2016 succeeds-but on Windows 2016 RDMA is not supported. 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.	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.

Linux

Symptom	Conditions	Workaround
When sending high bandwidth NVMe traffic on some Cisco Nexus 9000 switches, the switch port that connected to the storage sometimes reaches the max PFC peak and does not automatically clear the buffers. In Nexus 9000 switches, the nxos command "**`show hardware internal buffer info pkt-stats input peak`**" shows that the `Peak_cell` or `PeakQos` value for the port reaches more than 1000.	The NVMe traffic will drop.	To recover the switch from this error mode. **1.** Log into the switch. **2.** Locate the port that connected to the storage and shut down the port using "shutdown" command **3.** Execute the following commands one by one: `# clear counters` `# clear counter buffers module 1` `# clear qos statistics` **4.** Run **`no shutdown`** on the port that was shut down.

ESXi

Symptom	Conditions	Workaround
When using the command **esxcli storage core adapter list** to list the vmhba, the Driver's Link State for vmhba64 and vmhba65 rdma ports displays *Link-n/a* instead of *Online*. **Note** VMware Developer Center Partner Network (DCPN) Case ID - 00113157	This is a known issue in ESXi 7.0 Update 3.	None



CHAPTER 4

Configuring NFS Over RDMA Using ROCEv2

- Configuring NFS over RDMA On Cisco Intersight, on page 51
- Removing NFS over RDMA Using Cisco Intersight, on page 53
- Installing Cisco eNIC and enic_rdma Drivers, on page 54
- Configuring NFS over RDMA on the Linux Host (Intel), on page 55
- Configuring NFS over RDMA on the Linux Host (AMD), on page 56
- Mounting and Verifying NFS over RDMA Volumes, on page 57

Configuring NFS over RDMA On Cisco Intersight

Cisco Intersight includes default adapter policies prepopulated with operational parameters. While you do not need to manually create the adapter policy, you must create the vNIC and associate it with the appropriate RoCEv2-enabled policy.

Ensure that you have created an RDMA-compatible service profile.

Procedure

Step 1 Navigate to **Configure** > **Policies**. Click **Create Policy**, select **UCS Server** platform type, search or choose **LAN Connectivity** policy, and click **Start**.

Step 2 On the policy **General** page, enter the following information, and click **Next**.

- Enter the policy name.
- Select the Target Platform as **UCS Server (Standalone)** or **UCS Server (FI-Attached)** or **UCS Server (Unified Edge)**.
- **Description** to help identify the policy.
- **Tag** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.

Step 3 On the **Policy Details** page, configure the following:

- Click **Add vNIC** to create a new vNIC.
- Provide a name for virtual ethernet interface.

  In the **Add vNIC** page, follow the configuration parameters to enable the RoCE vNIC settings:

- Click **Select Policy** link below the **Ethernet Network**. Use the **Create Policy** button to create a new Ethernet Network policy with the following property settings:
  - Enter the policy name and click **Next**.
  - Under VLAN Settings, configure the following:
    - VLAN Mode—Option to determine if the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic
    - Enable QinQ Tunneling—Enable QinQ (802.1Q-in-802.1Q) Tunneling on the vNIC.
    - QinQ VLAN—When activating VIC QinQ (802.1Q-in-802.1Q) Tunneling, a particular VLAN ID is set. In Access VLAN mode, this QinQ VLAN ID is established as the default VLAN.

- Click **Select Policy** link below the **Ethernet QoS**. Use the **Create Policy** button to create a new Ethernet QoS policy with the following property settings:
  - Enter the policy name and click **Next**.
    - For **MTU**, choose or enter **1500, 4096, 9000, or 9158**
    - For **Rate Limit**, enter the value in Mbps (0-10G/40G/100G depending on Adapter Model).
    - For **Class of Service**, choose or enter **5**

      Note
      This property is available only on Standalone servers.
    - For **Burst**, enter the range in between 1024 and 1000000. The default value is 10240.
    - For **Priority**, choose **Platinum** or **any no-drop**
    - Slide to **Enable Trust Host CoS** toggle button.

      Note
      This property is available only on Intersight Managed Mode servers.

- Click **Select Policy** link below the **Ethernet Adapter**, click **Create an Ethernet Adapter Policy**:
  - **Use the Default Configuration**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy and under Ethernet Adapter Default Configuration click **Select Default Configuration** to search, the pre-defined Ethernet Adapter Default Configuration. Click **Next** and then **Create**.
  - **Configure RoCE Settings in the policy**: Click **Create New** to create a new policy. In the **General** page, enter the name of the policy. Under Policy Details page on right pane, use the following property settings, then click **Next**, and then **Create**..
    - For **Enable RDMA over Converged Ethernet**, slide to enable.
    - For **Queue Pairs**, choose or enter **1024**
    - For **Memory Regions**, choose or enter **131072**
    - For **Resource Groups**, choose or enter **8**
    - For **Version**, select **Version 2**

Step 4 Click **Add** to add and save the new vNIC settings.

Note
All the fields with * are mandatory for creating LAN Connectivity Policy. Ensure they are filled out or selected with appropriate policies.

Step 5 Click **Create** to complete the LAN Connectivity policy with RoCE v2 property settings.

Step 6 Associate the LAN Connectivity policy to the server profile and deploy.

Note
For more information, see *Creating a LAN Connectivity Policy*, *Creating an Ethernet QoS Policy*, and *Creating an Ethernet Adapter Policy* in Configuring UCS Server Policies and Configuring UCS Server Profiles.

The Cisco Intersight applies the new adapter policy to the server, and the server initiates a reboot to finalize the configuration.

What to do next

After the server reboots, proceed to configure the host-side NFS over RDMA settings as described in the host configuration procedure.

Removing NFS over RDMA Using Cisco Intersight

Use these steps to remove the Network File System (NFS) over Remote Direct Memory Access (RDMA) configuration from a specific vNIC.

Before you begin

You must be logged in with administrative privileges.

Procedure

Step 1 Navigate to **Configure** > **Profiles** and select your server profile with RoCEv2 LAN Connectivity policy associated.

Step 2 Expand the node for the organization where the service profile is located. (If the system does not include multitenancy, expand the root node).

Step 3 Select the vNIC intended for RDMA traffic.

Step 4 Disable the **Enable RDMA over Converged Ethernet** radio button to remove the RDMA on this virtual interface.

Step 5 Click **Save and Deploy**.

What to do next

- The Cisco Intersight applies the standard adapter policy to the server.
- Reboot the server to ensure the changes take effect and the RoCEv2-based NFS interface is fully removed from the host operating system.

Installing Cisco eNIC and enic_rdma Drivers

Use the steps to install Cisco eNIC and enic_rdma Drivers for NFS over RDMA.

Before you begin

- Ensure you have downloaded the matched set of eNIC and enic_rdma drivers from the Cisco support portal.
- Ensure you have root or sudo privileges to perform driver installation.



Note The driver files are available in the driver ISO under Unified Computing System (UCS) Drivers, typically named in the format ucs-cxxx-drivers-linux-X.Y.Zq.iso.

Procedure

Step 1 Install the eNIC and enic_rdma driver package on the host: `# rpm -ivh kmod-enic-<version>.x86_64.rpm kmod-enic_rdma-<version>.x86_64.rpm`

```
    root@ec2-50-19-41-89:~# rpm -ivh kmod-enic-4.10.366.0-1233.0.rhel10u1_6.12.0_124.8.1.x86_64.rpm
Verifying...                          ################################# [100%]
Preparing...                          ################################# [100%]
Updating / installing...
1:kmod-enic-4.10.366.0-1233.0.rhel1################################# [100%]
root@ec2-50-19-41-89:~# rpm -ivh kmod-enic_rdma-1.10.366.0-1233.0.rhel10u1_6.12.0_124.8.1.x86_64.rpm
Verifying...                          ################################# [100%]
Preparing...                          ################################# [100%]
Updating / installing...
1:kmod-enic_rdma-1.10.366.0-1233.0.################################# [100%]
root@ec2-50-19-41-89:~#
```

Step 2 Reboot the server to load the drivers into the running kernel: `# reboot`

Step 3 Verify that the drivers are loaded and RoCEv2 is enabled by checking the system logs: `# dmesg | grep enic_rdma.`

The system should display the following output:

```
[    1.562913] enic_rdma: Cisco VIC Ethernet NIC RDMA Driver, ver 1.10.366.0-1233.0 init
[    1.687502] enic 0000:75:00.0 enp117s0f0: enic_rdma: FW v4 RoCEv2 enabled
[    1.771968] enic 0000:75:00.1 enp117s0f1: enic_rdma: FW v4 RoCEv2 enabled
```

The eNIC and enic_rdma drivers are successfully installed and loaded. The system logs confirm that the Cisco VIC interfaces are initialized and RoCEv2 functionality is enabled on the respective ports.

What to do next

After verifying the driver installation, proceed to configure the NFS over RDMA mount points.

Configuring NFS over RDMA on the Linux Host (Intel)

Before you begin

- Ensure the server is configured with RoCEv2 vNICs.
- Ensure you have administrative (root or sudo) privileges to the Linux host.
- Verify that the host system is Intel-based.

Use this procedure to configure NFS over RDMA on the Red Hat Enterprise Linux (RHEL) host (Intel-based system).

Procedure

Step 1 Open the `/etc/default/grub` file for editing.

Step 2 Add `intel_iommu=on` to the end of the line for `GRUB_CMDLINE_LINUX` as shown in the sample file below.

```
sample /etc/default/grub configuration file after adding intel_iommu=on:
# cat /etc/default/grub
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap biosdevname=1 rhgb quiet
 intel_iommu=on
GRUB_DISABLE_RECOVERY="true"
```

Step 3 Save the file.

Step 4 After saving the file, run the following command to generate a new grub.cfg file:

- For Legacy boot:

  ```
  # grub2-mkconfig -o /boot/grub2/grub.cfg
  ```

- For UEFI boot:

  ```
  # grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
  ```

Step 5 Reboot the server.

You must reboot your server for the changes to take after enabling IOMMU.

Step 6 Verify that the server booted with the `intel_iommu=on` option: `# cat /proc/cmdline | grep iommu. The output should confirm the inclusion of intel_iommu=on at the end of the line.`

`Sample Output: The output should confirm the inclusion of intel_iommu=on at the end of the line.`

```
[root@localhost basic-setup]# cat /proc/cmdline | grep iommu
BOOT_IMAGE=/vmlinuz-3.10.0-957.27.2.el7.x86_64 root=/dev/mapper/rhel-
root ro crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb
quiet intel_iommu=on LANG=en_US.UTF-8
```

The IOMMU option is enabled, and the system is prepared for the installation of the necessary RDMA drivers.

What to do next

- Proceed to install the downloaded `enic` and `enic_rdma` drivers.
- Configure your NFS mount points to utilize the RDMA interface.

Configuring NFS over RDMA on the Linux Host (AMD)

Before you begin

- Ensure the server is configured with RoCEv2 vNICs.
- Ensure you have administrative (root or sudo) privileges to the Linux host.
- Verify that the host system is AMD-based.

Use this procedure to configure NFS over RDMA on the Red Hat Enterprise Linux (RHEL) host for AMD-based systems.

Procedure

Step 1 Edit the GRUB configuration. Open the `/etc/default/grub` file for editing.

Step 2 Add `amd_iommu=on` to the end of the line for `GRUB_CMDLINE_LINUX` as shown in the following sample:

Sample:

```
# cat /etc/default/grub
                        GRUB_CMDLINE_LINUX="... rhgb quiet amd_iommu=on"
```

Step 3 Save the file.

Step 4 Generate the new GRUB configuration file (grub.cfg) for UEFI systems:

For UEFI boot:

Sample:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

Step 5 Reboot the server to apply the changes.

```
# reboot
```

Note: The server must be rebooted for the IOMMU changes to take effect.

Step 6 Verify that the server booted with the `amd_iommu=on` option.

```
# cat /proc/cmdline | grep iommu
```

The output should confirm the inclusion of `amd_iommu=on` at the end of the line.

Sample Output:

```
BOOT_IMAGE=(hd0,gpt2)/vmlinuz-6.12.0-124.8.1.el10_1.x86_64 ... rhgb quiet amd_iommu=on
```

The IOMMU option is enabled, and the system is prepared for the installation of the necessary RDMA drivers.

What to do next

- Proceed to install the downloaded `enic` and `enic_rdma` drivers.
- Configure your NFS mount points to utilize the RDMA interface as described in the mounting procedure.

Mounting and Verifying NFS over RDMA Volumes

Use these steps to mount NFS volumes using RDMA capable VNICs:

Before you begin

- Confirm Cisco VIC firmware and drivers (enic and enic_rdma) are up to date for RoCEv2.
- Set the MTU size consistently on all VLANs and set the appropriate QoS policies on upstream switches. The MTU must not exceed the maximum supported value of the storage arrays.
- Confirm Cisco Intersight is configured correctly for RoCEv2, including enabling RoCE properties on the vNIC and disabling incompatible failover features.
- Ensure kernel modules `enic`, and `enic_rdma` are loaded on the host OS.
- Validate that RDMA interfaces are up and configured with correct IP addressing matching the storage network.
- You have administrative (root) privileges on the host.

Procedure

Step 1 Create a local directory to serve as the mount point: `# mkdir /<mount_point>`

Sample

```
# mkdir /mnt/nfs_volume1
```

Step 2 Mount the NFS volume using the RDMA protocol:

- **IPv4:**

  ```
  mount -o proto=rdma <storage IPV4>:/<volume path>/<mount_point>
  ```

  Sample `mount -o proto=rdma 192.168.1.100:/mnt/nfs_volume1`

  Note
  The `nconnect=4` option is recommended to provide multiple connections for increased parallelism and performance.

- **IPv6:**

  ```
  mount -o proto=rdma6 '[IPV6 address]:/<volume path>' /<mount point>
  ```

Replace <IPv4_address>/<IPv6_address>, <volume_path>, and <mount_point> with your actual IP addresses, NFS export path, and desired local mount directory.

Step 3 Verify that the NFS over RDMA mount is created and active. Run the following command to filter the mount list for RDMA-specific connections:

```
# mount | grep rdma
```

Sample

```
50.19.40.5:/mnt/nfs_volume1 type nfs4
(rw,relatime,vers=4.1,rsize=65536,wsize=65536,namlen=255,hard,fatal_neterrors=none,
proto=rdma,nconnect=4,port=20049,timeo=600,retrans=2,sec=sys,clientaddr=<interface
address>,local_lock=none,addr=<target address>)
```

Step 4 Verify the volume capacity and mount status using the df command. Use the `df` command to ensure the volume is correctly mapped and reporting the expected storage capacity:

```
# df -h
```

Sample `<target IP>:/nfs_vol 501G 3.6G 497G 1% /mnt/nfs_volume1`

Step 5 Check the RDMA statistics to verify traffic flow:

This command provides NFS statistics on the mount, confirming it is using RDMA and showing incrementing counts during active traffic.

```
# mountstats --xprt <mount_point>
```

Step 6 To unmount the volume, use the following command:

```
# umount /<mount_point>
```

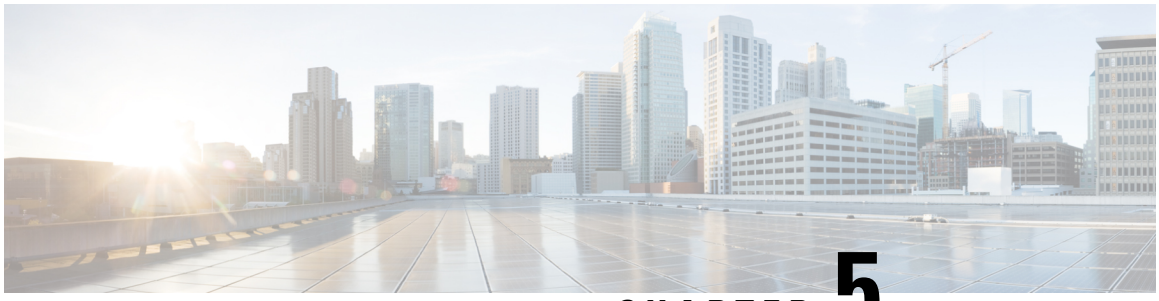
What to do next

After successfully mounting and verifying the NFS over RDMA volume, perform the following tasks:

- Configure Persistent Mounting: To ensure the volume mounts automatically after a system reboot, add an entry to the `/etc/fstab` file. *Example entry:*

  ```
  <storage_ip>:/<volume_path> <mount_point> nfs rdma,nconnect=4,x-systemd.automount 0
          0
  ```

- Configure Application Workloads: Update your application or database configurations to point to the new `<mount_point>` to begin utilizing the high-throughput, low-latency RDMA storage.



CHAPTER 5

Configuring Single Root I/O Virtualization (SR-IOV)

- Configuring BIOS and SR-IOV VFs, on page 59
- Configuring SR-IOV VFs on the ESXi Host Server , on page 73
- Configuring SR-IOV VFs on the Linux Host Server , on page 79

Configuring BIOS and SR-IOV VFs

Enabling BIOS Parameters

Before you begin

- Ensure your BIOS policy is set up with the following options:
  - For Intel based servers, enable **Intel VT for directed IO** under **Intel Directed IO** tab.

Note Intel VT for directed IO is not available for Intel C220 M8 and Intel C240 M8 platforms.

  - For AMD based servers, enable **IOMMU** and **SVM Mode** under **Processor** tab.

  To update BIOS options, see, Cisco UCS Server BIOS Tokens in Intersight Managed Mode.

- You must have a server profile already created for SR-IOV configuration. To create a Server Profile see Creating a UCS Server Profile. Once the Server Profile is created, follow the steps in this procedure to enable the BIOS policy.

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 Navigate to **Configure** > **Policies** > **Create a Policy**

Step 3 On the **Select Policy Type** page, select **BIOS**, click **Start**.

Step 4 At the **General** page, enter the policy name, and click **Next**.

Step 5 On the **Policy Details** page, configure the following BIOS settings:

- Select **All Platforms**.
- For a server with an Intel CPU, configure the BIOS settings as follows:
  - Enable **Intel VT for Directed IO** under the **Intel Directed IO** drop-down list.
  - Enable **Intel(R) VT** under the **Processor** drop-down list.
- For a server with AMD CPU, configure the BIOS settings as follows:
  - Enable **IOMMU** under **Memory** drop-down list.
  - Enable **SVM Mode** under **Processor** drop-down list.

Step 6 Click **Create**.

Step 7 Associate the BIOS policy to the server profile, and reboot the server.

Note
For more information, see Creating a BIOS Policy and Configuring Server Profile.

Create Ethernet Adapter Policy for SR-IOV

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 In the **Navigation** pane, choose **Configure** > **Policies**, and then click **Create Policy**.

Step 3 Select **Ethernet Adapter**, and then click **Start**.

Step 4 At the **General** page, enter the policy name.

Step 5 Click **Select Cisco Provided Configuration**, choose **SRIOV-HPN**, and click **Select**

Step 6 Click **Next**.

Step 7 Click **Create**.

Enabling SR-IOV VFs using Cisco Intersight GUI

To enable SR-IOV from Cisco Intersight, you must

- Create an SRIOV HPN Connection Policy with desired number of VFs.
- Assign the SRIOV HPN Connection Policy to a Server Profile.

Before you begin

- Ensure that the required BIOS options are enabled before performing this procedure.

Procedure

Step 1 Log in to Cisco Intersight.

Step 2 In the **Navigation** pane, choose **Configure** > **Policies**, and then click **Create Policy**.

Step 3 Select **LAN Connectivity**, and then click **Start**.

Step 4 On the **General** page, enter the following information:

- **Name** of your policy.
- **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.

  A LAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a LAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.
- **Set Tags** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- **Description** to help identify the policy.

Step 5 Click **Next**.

Step 6 On the **Policy Details** page, configure the following:

To set up a vNIC without using a template, click **Add vNIC** and configure the following parameters:

Table 1: For Standalone Servers

Property	Description
Add vNIC Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.	
Name	vNIC name.
Placement Placement Settings for the virtual interface. **Simple** When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.	
Slot ID	When automatic slot ID assignment is disabled, the slot ID needs to be entered manually. Supported values are (1-15) and MLOM.
Uplink Port	Adapter port on which the virtual interface will be created.

Property	Description
PCI link The PCI link used as transport for the virtual interface. **Note** The host device order can get impacted when using both the PCI links and while adding or removing vNICs.	
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs.
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Ethernet Network	Relationship to the Ethernet Network Policy. Select the created Ethernet adapter policy from above. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on Standalone servers. Select or create an Ethernet Network policy.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an Ethernet QoS policy.
Ethernet Adapter	Relationship to the Ethernet Adapter Policy. Select or create an Ethernet Adapter for SR-IOV from above.
Connection	
Disabled	Configuration is disabled.
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.

Property	Description
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. select policy
Class of Service	Class of Service to be used for traffic on the usNIC.
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.

Table 2: For FI-Attached Servers

Property	Description
Enable Azure Stack Host QoS	Enabling AzureStack-Host QoS on an adapter allows the user to carve out traffic classes for RDMA traffic which ensures that a desired portion of the bandwidth is allocated to it.
IQN	
None	This option ensures the IQN name is not associated with the policy.
Pool	
IQN Pool	Relationship to the iSCSI Qualified Name Pool. Select or create an IQN pool.
Static If you select this option, enter a static IQN for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain	
IQN Identifier	User provided static iSCSI Qualified Name (IQN) for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain.
vNIC Configuration	
Manual vNICs Placement	If you select this option, you must manually specify the placement for each vNIC. You can also use the Graphic vNICs Editor to create and specify the placement for each vNIC manually by adding vNICs and slots, and defining the connection between them. **Note** For manual placement, PCI Link is not supported on UCS VIC 1400 Series adapters. If a LAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
Auto vNICs Placement	If you select this option, vNIC placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.
Add vNIC Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.	
Name	Name of the virtual ethernet interface.

Property	Description
Pin Group Name	Pingroup name associated to vNIC for static pinning. LCP deploy will resolve pingroup name and fetches the correspoding uplink port/port channel to pin the vNIC traffic.
MAC	
Pool	If you select this option, select the MAC pool that you want to associate with the LAN Connectivity policy.
MAC Pool	The MAC pool that is assigned. Select or create a MAC pool.
Static	Click **Static** and enter a static MAC address for MAC address assignment. This option is available only for Cisco Intersight Managed FI-Attached servers.
Static MAC Address	The MAC address must be in hexadecimal format xx:xx:xx:xx:xx:xx. To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix 00:25:B5:xx:xx:xx.
Placement	
Simple When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0. **Note** Not applicable for Auto vNIC Placement.	
Switch ID	Refers to the Fabric Interconnect that carries the vNIC traffic.
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The order should start from zero with no overlaps. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. All VIC adapters have a single PCI link except VIC 1340, VIC 1380 and VIC 1385 which have two. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs. Not applicable for Auto vNIC Placement.

Property	Description
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover	Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.
Enabled	Enabling failover ensures that traffic from the vNIC automatically fails over to the secondary Fabric Interconnect, in case the specified Fabric Interconnect path goes down. Failover applies only to Cisco VICs that are connected to a Fabric Interconnect cluster.
Ethernet Network Group	Select or create an Ethernet Network Group policy. You can add multiple Ethernet Network Group Policies (ENGPs) on vNICs. The maximum number of ethernet network group policies is restricted to 50 including shared policies. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. You can associate only one ethernet network group policy with a vNIC if QinQ is configured. The native VLAN must be the same across all ethernet network group policies, or must be set in only one ethernet network group policy. Relationship to the Fabric Ethernet Group Policy. Select or create an ethernet network group policy.
Ethernet Network Control	Relationship to the Fabric Ethernet Network Policy. Select or create an ethernet network control policy.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an ethernet QoS policy.
Ethernet Adapter	Relationship to the the Ethernet Adapter Policy. Select or create an ethernet adapter policy.

Property	Description
iSCSI Boot	Relationship to the boot iSCSI Policy. • Not applicable to SR-IOV. • This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. Select or create an iSCSI boot policy.
Connection	
Disabled	Configuration is disabled.
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. Select or create a usNIC adapter policy.
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	

Property	Description
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.
Placement - Advanced	
Automatic Slot ID Assignment	When enabled, slot ID is determined automatically by the system.
Slot ID	When automatic slot ID assignment is disabled, the slot. ID needs to be entered manually. Supported values are (1-15) and MLOM.
Automatic PCI link Assignment	When enabled, PCI link is determined automatically by the system. **Note** If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vNICs are placed on the same PCI link (link 0). If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vNIC will be placed on PCI link 0.
Load Balanced	When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links. If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to specify the PCI order to load balance the vNICs. If both automatic PCI link assignment and automatic Slot ID are disabled, you need to specify the slot and the PCI order to load balance the vNICs. **Note** You cannot change the PCI link mode of two vNICs from Load Balanced mode to Custom mode without deleting and recreating the vNICs. Enter the following fields for Load Balanced option: Switch ID, PCI Order.

Property	Description
Custom	If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order and the PCI link. **Note** You cannot change the PCI link mode of two vNICs from Custom mode to Load Balanced mode without deleting and recreating the vNICs. Enter the following fields for Custom option: PCI Link, Switch ID, PCI Order.
PCI Link	The PCI Link used as transport for the virtual interface. PCI Link is only applicable for select Cisco UCS VIC 1300 models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored. **Note** Not applicable for Auto vNIC Placement.
Switch ID	The fabric port to which the vNICs will be associated.
PCI Order	The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter. **Note** You cannot change the PCI order of two vNICs without deleting and recreating the vNICs. Not applicable for Auto vNIC Placement.
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.	

Property	Description
Enabled	Enabling failover ensures that traffic from the vNIC automatically fails over to the secondary Fabric Interconnect, in case the specified Fabric Interconnect path goes down. Failover applies only to Cisco VICs that are connected to a Fabric Interconnect cluster.
Ethernet Network Group	Select or create an Ethernet Network Group policy. You can add multiple Ethernet Network Group Policies (ENGPs) on vNICs. The maximum number of ethernet network group policies is restricted to 50 including shared policies. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. You can associate only one ethernet network group policy with a vNIC if QinQ is configured. The native VLAN must be the same across all ethernet network group policies, or must be set in only one ethernet network group policy. Relationship to the Fabric Ethernet Group Policy. Select or create an ethernet network group policy.
Ethernet Network Control	Relationship to the Fabric Ethernet Network Control Policy. Select or create an ethernet network control policy. **Note** This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet QoS	Relationship to the Ethernet QoS Policy. Select or create an ethernet QoS policy.
Ethernet Adapter	Relationship to the Ethernet Adapter Policy. Select or create an ethernet adapter policy.
iSCSI Boot	Relationship to the boot iSCSI Policy. • Not applicable to SR-IOV. • This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers. Select or create an iSCSI boot policy.
Connection	
Disabled	Configuration is disabled.

Property	Description
usNIC	
Number of usNICs	Number of usNIC interfaces to be created. When usNIC is enabled, the valid values are from 1 to 225. When usNIC is disabled, the default value is 0.
usNIC Adapter Policy	Ethernet Adapter policy to be associated with the usNICs. select policy
VMQ	
Enable Virtual Machine Multi-Queue	Enables Virtual Machine Multi-Queue feature on the virtual interface. VMMQ allows configuration of multiple I/O queues for a single VM and thus distributes traffic across multiple CPU cores in a VM.
Number of Interrupts	The number of interrupt resources to be allocated. Recommended value is the number of CPU threads or logical processors available in the server. **Note** Number of Interrupts overrides the Interrupts value in the selected Ethernet Adapter Policy. Number of Virtual Machine Queues overrides Receive Queue Count, Transmit Queue Count, and Completion Queue Count values of the selected Ethernet Adapter Policy.
Number of Virtual Machine Queues	The number of hardware Virtual Machine Queues to be allocated. The number of VMQs per adapter must be one more than the maximum number of VM NICs.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead. **Note** SR-IOV setting for Windows Target OS is not supported.	
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.

Property	Description
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.
Template	Setting up a vNIC using a template.
Add vNIC from Template The source vNIC template to apply to the vNIC instance. All configuration settings from the vNIC template will be applied to the vNIC instance except the overridden list of configurations.	
Name	vNIC name.
vNIC Template	The source vNIC template to apply to the vNIC instance. All configuration settings from the vNIC template will be applied to the vNIC instance except the overridden list of configurations. Select or create an vNIC Template.
Graphics vNIC Editor	Displays the graphics vNIC editor details.

Step 7 Click **Add** and then click **Create**.

Disabling SR-IOV VFs Using Cisco Intersight GUI

Procedure

Step 1 In the **Navigation** pane, click **Policies**.
Step 2 On the **Policies** page, click **Search**.
Step 3 Enter the name of created LAN Connectivity policy from above.
Step 4 Click the **policy**.
Step 5 From **Actions**, select **Edit**.
Step 6 Click **Next**.
Step 7 Select vNIC that you want to disable SR-IOV VFs, and click **Edit**.
Step 8 From **Connection**, click **Disabled**, and then **Update**.
Step 9 Click **Save & Proceed**.

Configuring SR-IOV VFs on the ESXi Host Server

Installing Cisco eNIC Driver

Before you begin

Ensure that the required BIOS parameters and SR-IOV VFs configurations are completed.

The inbox driver does not support SR-IOV functionality. To enable SR-IOV, you must install the appropriate driver. For example, Cisco recommends using the enic drivers for SR-IOV functionality.

Procedure

Step 1 Install the enic driver on the host.

The following example shows the installation of eNIC driver on ESXi:

```
[root@localhost:/vmfs/volumes/645c8bdd-c655e553-8ba0-e8d32272f6c0] esxcli software vib install -v
/vmfs/volumes/C240M7-Standalone/CIS_bootbank_nenic_2.0.15.0-1OEM.800.1.0.20613240.vib --no-sig-check

Installation Result

   Message: The update completed successfully, but the system needs to be rebooted for the changes
to be effective.

   VIBs Installed: CIS_bootbank_nenic_2.0.15.0-1OEM.800.1.0.20613240

   VIBs Removed: CIS_bootbank_nenic_2.0.16.0-1OEM.800.1.0.20613240

   VIBs Skipped:

   Reboot Required: true

   DPU Results:

[root@localhost:/vmfs/volumes/645c8bdd-c655e553-8ba0-e8d32272f6c0]
```

Step 2 Reboot the server to load the enic driver into the running kernel.

Step 3 After reboot, execute the command `esxcli software vib list | grep nenic` to check the driver version.

For more information, see Installing Cisco enic and enic_rdma Drivers, on page 34.

Verifying the SR-IOV VFs Per Ports on the Host

You can verify the total number of SR-IOV VFs in the following two ways:

Procedure

Step 1 Verify by logging into the VMware ESXi Host Client.:

- Login to the VMware ESXi Host Client.
- Execute the following command to check the vNIC with SR-IOV capability:

```
root@localhost:~] esxcli network sriovnic list
Name    PCI Device    Driver  Link  Speed  Duplex  MAC Address         MTU  Description
-----------------------------------------------------------------------------------------------------
vmnic0  0000:1b:00.0  nenic   Up    50000  Full    f4:ee:31:30:80:40  1500  Cisco Systems Inc
Cisco VIC Ethernet NIC
```

The following output shows the number of VF configured on vNIC:

```
[root@localhost:~] esxcli network sriovnic vf list -n vmnic0
VF ID  Active  PCI Address     Owner World ID
0   false  00000:027:00.1   -
1   false  00000:027:00.2   -
2   false  00000:027:00.3   -
3   false  00000:027:00.4   -
4   false  00000:027:00.5   -
5   false  00000:027:00.6   -
6   false  00000:027:00.7   -
7   false  00000:027:01.0   -
```

Step 2 Alternatively, you can also access your host from vSphere vCenter Client.

For more information on configuring SR-IOV VFs on the host, see Creating SR-IOV VFs on the Host.

After you reboot the host server, do the following:

- Login to the ESXi Host Client, and choose **Networking** > **Virtual Switches**.
- Click **Add Standard Virtual Switch**.
- Add a switch name in the **vSwitch Name** field, select the vmnic with SR-IOV capability, and click **Add**.

  The maximum number of Virtual Functions (VFs) is set to 10.
- In the **Port Groups** tab, click **Add Port Group**.
- In the **Add Port Group** dialog-box, add a new port group and select the switch from the **Virtual Switch** drop-down.

Creating SR-IOV VFs on the Host

Procedure

Step 1 Login to your VMware ESXi Host Client.

Alternatively, you can also access your host from vSphere vCenter Client and browse to **Configure** > **Networking** > **Physical adapters**.

Step 2 Go to **Host** > **Manage** and select the **Hardware** tab.

Step 3 Select **PCI Devices** from the list.

Step 4 From the drop-down list, select **SR-IOV Capable**.

The list shows all the SR-IOV capable devices.

Step 5 Select the vNIC for which you wish to create the VFs.

Step 6 Click **Configure SR-IOV**.

Configure SR-IOV for Cisco VIC Ethernet NIC window is displayed.

Step 7 Perform the following:

Field	Description
Enabled radio button	Select **Yes** to enable the configuration.
Virtual functions field	Number of VFs as configured on SRIOV connection policy that are available for the configuration. Enter an integer between 1 and 64.

Step 8 Click **Save** and then reboot the host server.

Configuring the Switch

Before you begin

Ensure that the SR-IOV VFs are configured.

Procedure

Step 1 Login to your VMware ESXi Host Client.

Step 2 Navigate to **Host** > **Networking** and select the **Virtual switches** tab.

Step 3 Click **Add Standard Virtual Switch**.

Step 4 Enter the name for the switch.

Step 5 Select a SR-IOV Capable Vmnic from the list.

Step 6 Click **Add**.

Step 7 Complete the following:

Field	Description
vSwitch Name field	Enter a suitable name for the virtual switch.
MTU field	Enter the maximum transmission unit. The default is 1500 bytes. The maximum value is 9000 bytes.

Field	Description
Uplink 1 drop-down list	From the drop-down list, select the PCIe devices for which you created the SR-IOVs.
Link Discovery	From the drop-down list, select the **Mode** and the **Protocol**. **Note** These fields remain as default.
Security	Choose from the following options: • **Promiscuous mode**—**Accept**, **Reject**, or **Inherit from vSwitch**. • **MAC address changes**—**Accept**, **Reject**, or **Inherit from vSwitch**. **Forged trasmits**—**Accept**, **Reject**, or **Inherit from vSwitch**.
NIC teaming	Choose from the following: • **Load balancing**—From the drop-down list choose the Load balancing. Values are: **Inherit from vSwitch**, • **Network failover detection**—From the drop-down list choose the network failover detection. Values are: **Inherit from vSwitch**, • **Notify switches**—Choose the notify switches. Values are **Yes**, **No**, **Inherit from vSwitch**. • **Fallback**—Choose the fallback. Values are **Yes**, **No**, **Inherit from vSwitch**. • **Override failover order**—From the drop-down list choose the override failover order. Values are **Yes** or **No**, • **Failover order**—Choose the failover order.
Traffic Shaping	Perform the following: • **Status**—Choose the status. Values are **Enabled**, **Disabled**, **Inherit from vSwitch**. • **Average bandwidth**—Enter the average bandwidth. • **Peek bandwidth**—Enter the peek bandwidth. • **Burst size**—Enter the burst size. **Note** Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the virtual switch.

What to do next

Creating a Virtual Port, on page 77

Creating a Virtual Port

Before you begin

Ensure that the SR-IOV VFs are configured.

Procedure

Step 1 Login to your VMware ESXi Host Client.

Step 2 Go to **Host** > **Networking** and select the **Port Groups** tab.

Step 3 Click **Add port group**.

Add port group-New port group window is displayed

Step 4 Complete the following:

Field	Description
Name field	Enter a suitable name for the virtual port.
VLAN ID field	Enter the VLAN ID.
Virtual Switch drop-down list	From the drop-down list, select the virtual switch.
Security	Choose from the following options: • **Promiscuous mode**—**Accept**, **Reject**, or **Inherit from vSwitch**. • **MAC address changes**—**Accept**, **Reject**, or **Inherit from vSwitch**. **Forged trasmits**—**Accept**, **Reject**, or **Inherit from vSwitch**.

Step 5 Click **Add**.

Creating a New Virtual Machine (VM)

Before you begin

- Login to vCenter using the credentials
- OS ISO image is copied to the datastore of the host server.

Procedure

Installing OS on Guest VM on ESXi.

Adding SR-IOV VF on the Virtual Machine

Before you begin

Power off the Virtual Machine.

Procedure

Step 1 In the Virtual Machine Manager, right-click on the Virtual Machine and select **Open**.

Step 2 Click the **Show Virtual Hardware Detail** icon next to **Monitor** icon.

Step 3 Click **Add Hardware**.

Step 4 In the **Add New Virtual Hardware** window, select **PCI Host Device**. Under the **PCI Device Details** tab, assign a created SR-IOV VF to the Virtual Machine.

Step 5 Click **Finish**.

Step 6 Power on the Virtual Machine.

What to do next

You can now log into the virtual machine, install Cisco eNIC driver version, reboot the virtual machine, and then use the ip link command to verify the added SR-IOV VF. For more information, see Installing Cisco eNIC Driver.

Installing OS on Guest VM on ESXi

Before you begin

Upload the Linux operating system ISO on the datastore.

Procedure

Step 1 Right-click the host node and navigate to **vCenter** > **New Virtual machine**.

Step 2 Select a **Creation Type** > **Create New Virtual Machine**, and click **Next**.

Step 3 Enter a name for the folder, and click **Next**.

Step 4 Select a compute resource, choose a node and click **Next**.

Step 5 Select Storage and check the datastore radio-button, and click **Next**.

Step 6 Select the compatability ESXi 8.0 or later and click **Next**.

Step 7 Select a guest OS version as **RHEL Linux9 (64-bit)**, and click **Next**.

Step 8 Customize the hardware set **CPU** to 2, and **Memory values** to 4 GB.

Step 9 Expand the **Memory** tab, and check **Reserve all guest memory (All locket)** check box.

Step 10 Select **New CD/DVD Drive** *(Datastore ISO file)*, and check the **Connect At Power On** check box.

Step 11 Under **CD/DVD Media**, browse and select the Linux ISO image and click **Next**.

Step 12 Click **Finish**.

Configuring SR-IOV VFs on the Linux Host Server

Installing Cisco eNIC Driver and Enabling IOMMU in Linux Kernel

Before you begin

Ensure that the required BIOS parameters and SR-IOV VFs configurations are completed.

Procedure

Step 1 Install the enic driver on the host.

Following example shows the installation of eNIC driver on RHEL:

```
[user@rack-111 drivers]# rpm -ivh kmod-enic-4.7.0.5-1076.6.rhel9u4_5.14.0_427.13.1.x86_64.rpm
Verifying...                          ################################# [100%]
Preparing...                          ################################# [100%]
Updating / installing...
   1:kmod-enic-4.7.0.5-1076.6.rhel9u4_################################# [100%]
[user@rack-111 drivers]#
```

Step 2 Enable IOMMU on the host using **grubby** command.

Following example shows how to enable IOMMU on RHEL:

```
[user@rack-111 drivers]# grubby --update-kernel=ALL --args="intel_iommu=on iommu=pt"
```

Step 3 Reboot the server to load the enic driver into the running kernel.

Step 4 Execute **modinfo enic** to check enic driver is loaded.

Following example shows the output of **modinfo enic** command:

```
[user@rack-111 drivers]# modinfo enic
filename:       /lib/modules/5.14.0-427.13.1.el9_4.x86_64/extra/enic/enic.ko
version:        4.7.0.5-1076.6
retpoline:      Y
license:        GPL v2
author:         Scott Feldman scofeldm@cisco.com
description:    Cisco VIC Ethernet NIC Driver
rhelversion:    9.4
srcversion:     3A1B1E81C9641925B34D1B2
alias:          pci:v00001137d000002B7sv*sd*bc*sc*i*
alias:          pci:v00001137d00000071sv*sd*bc*sc*i*
```

```
alias:          pci:v00001137d00000044sv*sd*bc*sc*i*
alias:          pci:v00001137d00000043sv*sd*bc*sc*i*
depends:
retpoline:      Y
name:           enic
vermagic:       5.14.0-427.13.1.el9_4.x86_64 SMP preempt mod_unload modversions
sig_id:         PKCS#7
signer:         Cisco UCS Driver Signing REL Cert
sig_key:        D0:54:9A:88:88:DD:0E:7A
sig_hashalgo:   sha256
signature:      89:9C:DA:53:D1:FF:0A:DA:98:9A:7F:AF:63:29:66:EB:FF:0C:D6:65:
                39:6C:15:40:30:6E:99:4B:2C:F0:54:2E:EB:A4:8A:33:D5:9C:41:7A:
                A4:DB:C8:52:55:74:3A:68:F3:22:36:7B:2A:7C:7C:40:8B:7F:6D:9E:
                A5:CF:06:F1:23:42:E6:60:DB:78:0E:46:C9:0C:BC:06:9B:02:A0:AA:
                5A:FC:36:A3:FB:B0:FE:76:F2:EB:2F:AD:AD:84:89:61:30:7D:E9:2F:
                5D:E1:3E:EA:7C:10:B2:42:94:CD:4F:74:19:A6:16:FE:75:B6:78:49:
                E8:F0:4A:A9:01:BB:92:44:A9:FE:C7:CE:DB:E8:F5:08:AF:36:1E:5F:
                30:D3:B1:5F:70:62:56:6F:C2:38:8E:F2:88:28:0F:44:29:E5:44:66:
                34:B7:5C:A7:5E:21:C3:5D:42:D8:C0:87:CA:40:5E:C4:C0:2C:DA:26:
                D2:25:9B:58:A8:84:C6:A6:41:B3:24:9C:D7:E6:4A:79:42:00:32:82:
                7A:CB:36:D8:79:1D:41:1A:9E:1C:A8:0D:39:6D:C8:F1:0D:44:FA:00:
                93:1E:A3:C9:61:AA:DE:25:4A:38:68:C3:9C:14:55:5B:D3:AC:1C:85:
                00:FE:57:F1:DE:F7:A8:04:64:0E:5D:35:D8:AF:CF:A4
parm:           rxcopybreak:Maximum size of packet that is copied to a new buffer on receive (uint)
[user@rack-111 drivers]#
```

Verifying the Total number of SR-IOV VFs per Port on the Host

Before you begin

Ensure that Cisco eNIC driver is installed.

Procedure

Log into the host server and run the following command and replace *interface_name* with actual interface name on the host.

cat /sys/class/net/*interface_name*/device/sriov_totalvfs

Example

Following example shows the total number for SR-IOV VFs created from SRIOV HPN Connection Policy on p1p1 interface:

```
[user@rack-111 ~]# cat /sys/class/net/p1p1/device/sriov_totalvfs
32
[user@rack-111 ~]#
```

Creating SR-IOV VFs on the Host

Enabling SR-IOV VFs from SRIOV HPN Connection Policy does not create SR-IOV VFs on the host by default. To create SR-IOV VFs on the host, use the following procedure:

Procedure

Step 1 Execute the following command to create SR-IOV VFs on the host:

echo number_of_sriov_devices > /sys/class/net/*sriov interface_name***/device/sriov_numvfs**

Example:

Following example shows the creation of 6 SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# echo 6 > /sys/class/net/p1p1/device/sriov_numvfs
[user@rack-111 ~]#
```

Step 2 Execute the following command to verify the SR-IOV VFs created:

cat /sys/class/net/interface_name/device/sriov_numvfs

Example:

Following example shows the verification of SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# cat /sys/class/net/p1p1/device/sriov_numvfs
6
[user@rack-111 ~]#
```

Step 3 (Optional) Alternatively, IP link command shows created SR-IOV VFs.

ip link show interface_name

Example:

Following example shows created 6 SR-IOV VFs on p1p1 interface.

```
[user@rack-111 ~]# ip link show pipl
2: p1p1:  <BROADCAST, MULTICAST, UP, LOWER_UP>mtu 9000 qdisc mq state UP mode
DEFAULT group default qlen 10 00
link/ether 98: a2:c0:66:32:80 brd ff:ff:ff:ff:ff:ff
vf 0 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 1 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 2 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 3 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 4 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
vf 5 link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff, spoof checking off,
link-state auto, trust off, query_rss off
altname enp9s0
altname eno5
[user@rack-111 ~]#
```

Note

After the host server reboots, the created SR-IOV VFs are removed from the host. By adding the command from Step 1 to rc.local file, the same number of SR-IOV VFs can be created each time the host server boots up.

What to do next

You can create a new virtual machine.

Creating a New Virtual Machine (VM)

Before you begin

- Host with Desktop Environment
- Virtualization packages are installed
- Copy OS ISO image to the sever datastore

Procedure

Step 1 Verify the virtualization is enabled on the host server by using this command.

lscpu | grep Virtualization

Example:

This example shows the Intel's virtualization technology VT-x is enabled.

```
[user@rack-111 ~]$ lscpu | grep Virtualization
Virtualization: VT-x
[user@rack-111 ~]$
```

Step 2 Verify the KVM modules are loaded by using this command.

lsmod | grep kvm

Example:

This example shows KVM modules are loaded in the host server.

```
[user@rack-111 ~]$ lsmod | grep kvm
kvm_intel    409600    8
kvm         1134592    1 kvm_intel
irqbypass      6384    290 vfio_pci_core, kvm
[user@rack-111 ~]$
```

Step 3 Type **virt-manager** command at the terminal to launch Virtual Machine Manager GUI.

Step 4 At the Virtual Machine Manager, click **File** > **New Virtual Machine** to create a new virtual machine.

Step 5 At **New VM window**, select **Local install media (ISO image or CDROM)** option and click **Forward**.

Step 6 At **Choose ISO or CDROM install media**, click **Browse**.

Step 7 At **Locate ISO media volume** window, click **Browser Local**.

Step 8 Go to the folder that has ISO image. Select ISO image and click **Open**.

Step 9 Click **Forward**.

Step 10 Select the desire Memory and CPU settings for the VM and click **Forward**.

Step 11 Choose the VM's disk image size and click **Forward**.

Step 12 Enter a name for the VM in the **Name** field and click **Finish**.

You may monitor the OS installation progress.

Adding SR-IOV VF on the Virtual Machine

Before you begin

Power off the Virtual Machine.

Procedure

Step 1 In the Virtual Machine Manager, right-click on the Virtual Machine and select **Open**.

Step 2 Click the **Show Virtual Hardware Detail** icon next to **Monitor** icon.

Step 3 Click **Add Hardware**.

Step 4 In the **Add New Virtual Hardware** window, select **PCI Host Device**. Under the **PCI Device Details** tab, assign a created SR-IOV VF to the Virtual Machine.

Step 5 Click **Finish**.

Step 6 Power on the Virtual Machine.

What to do next

You can now log into the virtual machine, install Cisco eNIC driver version, reboot the virtual machine, and then use the ip link command to verify the added SR-IOV VF. For more information, see Installing Cisco eNIC Driver.