

Guidelines, Limitations, and Requirements

- RoCEv2 for Windows, on page 1
- RoCEv2 for Linux, on page 3
- RoCEv2 for ESXi, on page 5
- SR-IOV for ESXi, on page 6
- SR-IOV for Linux, on page 8
- NDIS Poll Mode support for Windows, on page 9
- NFS over RDMA for RHEL, on page 9

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.