

Emulex Drivers for Windows Server 2016 Technical Preview 4 Release Notes

Date: December 2015

Purpose and Contact Information

These release notes describe supported features, unsupported features, driver installation, updating firmware, configuration, and known issues for the Emulex® drivers for Windows Server 2016 Technical Preview 4.

Caution: The software is pre-release and is for evaluation and test use only; it is not intended for use in a production environment.

Note:

- Customer support is not available for this pre-release software, but Emulex welcomes your feedback. Email previewfeedback.pdl@avagotech.com if you have issues with the drivers, or if you have suggestions. Emails may not receive a response.
- If you are using an OCe14000-series adapter or an OCm14000-series adapter in conjunction with the Windows Server 2016 Technical Preview driver, upgrade the firmware to the version available on the Windows Server 2016 page at <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>
- The pre-release drivers are compatible with management applications from the version 10.6 software release, which are available on the Management tab of the Windows Server 2016 page at <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>
- Refer to the *Emulex Drivers for Windows version 10.6 User Manual* on the Windows Server 2016 page <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>
- Performance testing was done using Windows 2012 R2.

Supported Hardware and Features

The Emulex driver for Windows Server 2016 Technical Preview 4 supports the following adapters and features.

Table 1 Supported Hardware and Supported Features

Adapter	Features
LPe16202 LPe16000-series OCe15000 LPe12000-series	Nano Server support (FC driver only)

Table 1 Supported Hardware and Supported Features (Continued)

Adapter	Features
OCe11100-series	All drivers and Nano Server support (No advanced feature support)
OCe14000 and OCm14000	All drivers, Nano Server support, and advanced feature support including Host Mode RDMA, VxLAN, and Packet Direct
OCe14000B and OCm14000B	All drivers, Nano Server support, and advanced feature support including Host Mode RDMA, VxLAN, Routable RoCE, and Packet Direct

Note: You must install the OCe14000, OCe14000B, and OCm14000 firmware from <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>

New Features

- Packet Direct is supported in the NIC driver.
- Nano Server support

Unsupported Features

- Multi-channel (UMC, VNIC, Flex10, NPar) is not supported when using Host Mode RDMA, VxLAN, or NVGRE.
- Advanced configurations, such as teaming on Host Mode, have not been tested.
- Simultaneous operation of NVGRE and VxLAN encapsulation over a single port is not supported.

Installing the Driver

For installation instructions, refer to Section 2, Installation, of the *Emulex Drivers Version 10.6 for Windows User Manual*, which is available at <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>

and “Updating Emulex OOB Drivers to a Nano Server VHD” on page 21.

Updating Adapter Firmware

You can use the OneConnect™ Flash Utility, the preferred method, or the OneCommand™ Manager application to update OneConnect adapter firmware.

To update adapter firmware using the OneConnect Flash Utility:

1. Locate the file "OneConnect-Flash-11.0.xxx.yy.iso", usually in the Firmware/Flash-ISO directory in the Emulex package, and burn a CD with this image, or mount it over the network (mounting over the network is easier).
2. Reboot the computer and select the image as the boot device.
After the operating system loads, a prompt appears to confirm the firmware flash.

3. Flash the firmware and reboot the computer.

Refer to the *Using the OneConnect™ Flash Utility to Update OneConnect Adapter Firmware* instructions available on the Emulex website for more information about the utility.

To update adapter firmware using the OneCommand Manager application GUI:

1. Install and start the the OneCommand Manager application.
2. Select **Batch>Download Firmware**.
3. From the "OneCommand Manager Batch Firmware Download dialog box, enter the firmware file path in the "Firmware File:" prompt, or select **Browse...** and navigate to the "oc14-11.0.xxx.yy.ufi" firmware image.
4. Click **Start Download**.
5. Reboot the computer after the download is finished.

Refer to the *OneCommand Manager application version 10.6 User Manual*, or the *OneCommand Manager application Command Line Interface version 10.6 User Manual* available on the Windows Server 2016 page at <http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/> for more information about the application.

Configuration

Enabling the RoCE Profile

You can enable RoCE using the PXESelect utility or the OneCommand Manager application. To enable RoCE using the PXESelect utility:

1. Press **<Ctrl+ P>** at the Emulex PXESelect splash screen as the server boots. A screen displays the global options.
2. Press **<Tab>** to highlight Personality.
3. Select the **NIC+RoCE** personality and the **RoCE-2** profile.
4. Save the settings and follow the instructions to complete the process.

Refer to the *Boot Version 10.6 for NIC, iSCSI, FCoE, and RoCE Protocols User Manual* available on the Emulex website for more information about the PXESelect utility.

To enable RoCE using the OneCommand Manager application GUI:

1. Start the OneCommand Manager application.
2. From the discovery-tree, select the adapter on which you want to enable RoCE.
3. Choose the **Adapter Configuration** tab.
4. Select the **Single personality** option.
5. For Personality, select **NIC+RoCE** from the menu.
6. For NIC+RoCE Configuration Type, select **RoCE-2**.

- Click **Apply** and follow the on screen instructions to complete the process.

Verifying the RoCE Profile is Enabled

Verify that the RoCE profile is enabled by using the Network Interface Property page or a PowerShell script.

Using the Network Interface Property page:

Network Direct is enabled.

Using a PowerShell script:

Get-NetAdapterRDMA

Example:

```
C:\Users\Administrator> Get-NetAdapterRDMA
```

Name	InterfaceDescription	Enabled
15-analyzer-88	Emulex OneConnect OCe14102-UX-D 2-por...	True16
Emulex OneConnect	OCe14102-UX-D 2-por...	True

If the profile is correct and NetworkDirect is enabled, you will see active NetworkDirect listeners on IP addresses (port 445) assigned to the NICs.

Configuring Host Mode RDMA

To configure Host Mode RDMA:

- Load the driver.

Note: Do not add VLAN to the network adapter advanced settings prior to creating the vSwitch.

- Using the Hyper-V Manager Virtual Switch Manager, create a new external virtual switch and attach it to the Emulex adapter.
- Using either PowerShell or Virtual Switch Manager, assign any required VLAN IDs to the management operating system.

Note: You cannot configure or change a VLAN ID when the system is running.

- From the Device Manager of the host operating system, select the **Advanced** page of the Hyper-V Virtual Ethernet Adapter and enable Network Direct (RDMA).

Note: Network Direct (RDMA) is disabled by default.

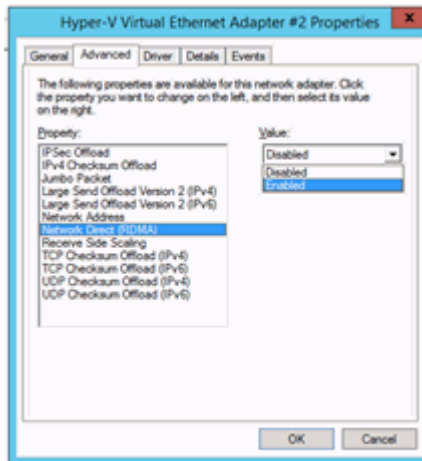


Figure 1 Network Direct (RDMA) Enabled

Configuring Routable RoCE

Routable RoCE is enabled by default. (Only supported on OCe14000B and OCm14000B adapters.)

To configure routable RoCE:

1. From the OneConnect Advanced tab, choose **RoCE Mode**.

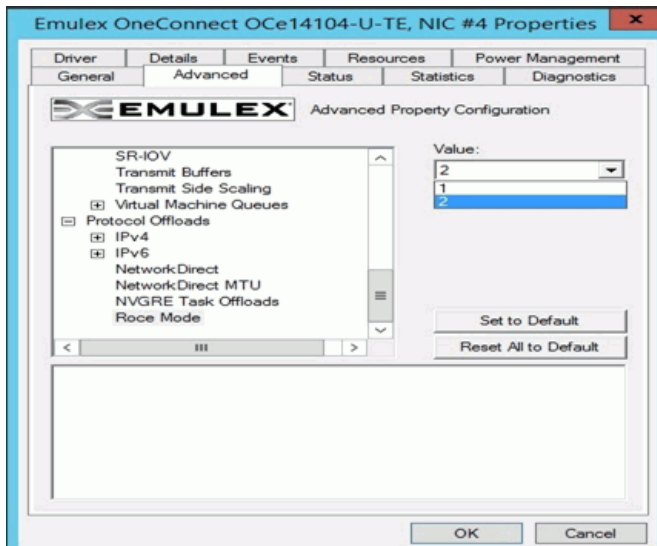


Figure 2 Routable RoCE Enabled (default)

2. From the **Value** menu, choose **2** for Routable RoCE (default setting) or **1** for Native RoCE.
3. Click **OK**.

Testing Routable RoCE (RoCE over UDP)

Emulex tested routable RoCE over the dedicated RoCE NICs (40G adapters) with two machines connected to one switch. Configure the switch port for routing and run traffic.

Check for UDP packets and make sure the tiny log reports routable requests and responses without error.

Enabling or Disabling Encapsulated Task Offload and VxLAN UDP

You can enable or disable Encapsulated Task Offload and modify the VxLAN UDP destination port number using the OneConnect Advanced tab or by using Powershell commands.

When both VxLAN and NVGRE encapsulation are enabled, NVGRE encapsulation is used. To enable VxLAN you must disable NVGRE Encapsulated Task Offload in the adapter property.

Using the OneConnect Advanced Tab

To enable or disable Encapsulated Task Offload (the default is enabled) or modify the VxLAN UDP destination port number using the OneConnect Advanced tab:

1. From the OneConnect Advanced tab, select the **Encapsulated Task Offload** or **VxLAN** parameter.

- Set the value to **Enabled** or **Disabled**, or enter a value for VxLAN UDP.

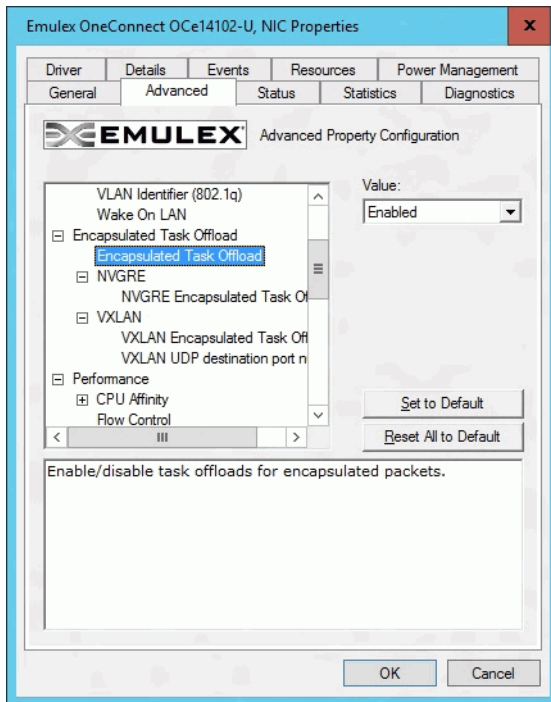


Figure 3 Encapsulated Task Offload (Enabled)




Figure 4 VXLAN Encapsulated Task Offload (Enabled)




Figure 5 VXLAN UDP destination port number

3. Click **OK**.
4. From the OneConnect Advanced tab, choose **NVGRE Encapsulated Task Offload**.

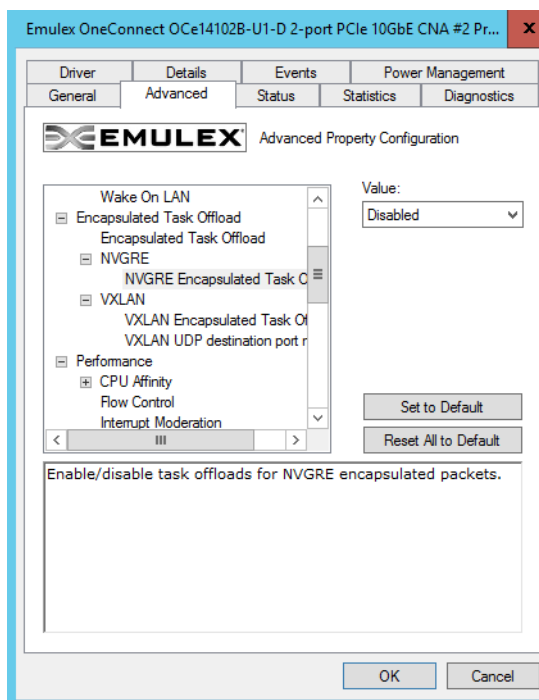


Figure 6 NVGRE Encapsulated Task Offload Disabled

5. From the **Value** menu, choose **Disabled** and click **OK**.

Using Powershell Commands

You can use PowerShell commands to set the values for EncapsulatedPacketTaskOffload, EncapsulatedPacketTaskOffloadVxlan, and VxlanUDPPortNumber subkeys.

For EncapsulatedPacketTaskOffload and EncapsulatedPacketTaskOffloadVxlan subkeys, a value of 0 disables the feature and a value of 1 enables the feature.

The VxlanUDPPortNumber subkey has a default value of 4789 and a value range of 1-65535.

Disable-NetAdapterEncapsulatedPacketTaskOffload,

Enable-NetAdapterEncapsulatedPacketTaskOffload,

Get-NetAdapterEncapsulatedPacketTaskOffload, and

Set-NetAdapterEncapsulatedPacketTaskOffload commands are also available.

Enabling Packet Direct

By default, the advanced driver parameter Packet Direct is enabled.

NIC Driver Options

Use the Get-NetAdapter PowerShell command to list all available adapters in the system.

The Get-Help <cmdl> -full command returns descriptions and help for the cmdlets.

Use the following PowerShell commands to get and set driver parameter values.

To get the driver parameter value:

```
Get-NetAdapterAdvancedProperty -Name <adapter name> -AllProperties  
-RegistryKeyword <registry keyword>
```

Example:

```
Get-NetAdapterAdvancedProperty -Name ""SLOT 6 Port 1"" -AllProperties  
-RegistryKeyword *RSS
```

To set the driver parameter value:

```
Set-NetAdapterAdvancedProperty -Name <adapter name> -AllProperties  
-RegistryKeyword <registry keyword> -RegistryValue <valid registry value>
```

Example:

```
Set-NetAdapterAdvancedProperty -Name ""SLOT 6 Port 1"" -AllProperties  
-RegistryKeyword *RSS -RegistryValue 1
```

Note: Select the <registry Keyword> and the <valid registry value> from Table 2 on page 11.

The following options are available for the NIC driver.

Table 2 Windows Server 2016 NIC Driver Options

Option Name	Registry Keyword	Registry Values	Definition
Class of Service (802.1p)	Class of service mode	1 Automatic Priority (default) 2 Filtered Priority 3 User Priority 4 Disable Priority	<p>The following modes are supported for selecting 802.1p priority tags:</p> <p>Automatic Priority – The DCBX standard allows the network adapter to negotiate priority class usage with DCBX aware endpoints such as switches or network cards. If the peer indicates that priority pause is supported for a non-zero priority, the NIC automatically inserts the default priority in all transmitted packets. This is the default mode, allowing priority pause to operate for both storage and network traffic. If the peer indicates a zero default priority (such as when the peer does not support priority pause), the device uses the “Non-Storage Priority” mode discussed below.</p> <p>Filtered Priority – This mode coerces the user priorities in each packet to avoid sending packets on the network function that may disrupt the adapter's storage traffic. The network device uses the next lower priority if a conflict exists. This mode is useful if multiple network priorities are necessary. Only a limited number of classes are supported for priority pause, so typically it does not function optimally in this mode.</p> <p>User Priority – This mode allows any user specified priority value and must be limited to cases where storage functions are not used.</p> <p>Disable Priority – The adapter always transmits either untagged packets, or VLAN ID (802.1q) tagged packets with a priority value (802.1p) of zero.</p>
Enhanced Transmission Selection	ETS	0 Disabled (default) 1 Enabled	<p>If ETS is enabled, the driver filters transmit packets based on the 802.1p priority tag into multiple separate transmit rings. The network switch must be configured for ETS to group priorities into a priority group (or traffic class). Each priority group can be assigned a QoS bandwidth limit. For example, one network priority can support priority flow control to achieve loss-less network traffic. Using separate hardware interfaces in the driver allows each priority to progress at a different rate, or pause temporarily without affecting the other priorities.</p> <p>If ETS is enabled, all configurations regarding bandwidth and priority flow control must be performed on the network switch. The adapter will learn the configuration using the DCBX protocol.</p> <p>NOTE For OCe11102 CNAs only.</p> <p>NOTE ETS is not supported in conjunction with VMQ technology. ETS is not available if SR-IOV is enabled.</p>
Flow Control	*FlowControl	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & TX Enabled) (Default)	<p>The IEEE 802.3x Ethernet specification defines a control frame between peers that can request a pause in packet transmissions. This allows one system to request a temporary halt of all incoming traffic when receive buffer space is exhausted.</p> <p>The network device may be configured to respond to pause frames (Rx Enable) and/or to send pause frames (Tx Enable). Flow control is almost always advantageous to avoid packet drops on the network. The switch or network peer must also have flow control enabled.</p>

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Interrupt Moderation	InterruptModerationLevel	0 (None) 8 (Static 90k Int/sec) 9 (Static 70k Int/sec) 10 (Static 50k Int/sec) 11 (Static 40k Int/sec) 12 (Static 30k Int/sec) 2 (Static 25k Int/sec) 13 (Static 20k Int/sec) 14 (Static 15k Int/sec) 15 (Static 10k Int/sec) 16 (Static 5k Int/sec) 4 (Adaptive) (default)	<p>The network device uses interrupt moderation algorithms to reduce the total amount of CPU cycles spent processing interrupts which increases efficiency for the system. However, interrupt moderation increases the latency of each send and receive. It should only be disabled when short latencies are more important than efficient CPU utilization.</p> <p>The "No Moderation" setting disables all delays to minimize latency.</p> <p>The "Static Moderation" uses a constant interrupt delay to avoid any spikes in interrupt rate.</p> <p>The Adaptive (default) setting causes the driver to dynamically maintain a target interrupt rate. The Adaptive setting value is controlled by a dynamic algorithm that scales well for various adapter link speeds.</p>
IP Checksum Offload (IPv4)	*IPChecksumOffloadIPv4	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & Tx Enabled) (Default)	<p>This offloads the transmit and the receive IPv4 checksum computation.</p> <p>Offloading checksums increases system efficiency.</p>
Large Send Offload v1 (IPv4)	*LsoV1IPv4	0 (Disabled) 1 (Enabled) (Default)	<p>Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (<= "Packet Size") that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet.</p> <p>The Windows Version 1 LSO supports only IPv4.</p>
Large Send Offload v2 (IPv4)	*LsoV2IPv4	0 (Disabled) 1 (Enabled) (Default)	<p>Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (<= "Packet Size") that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet.</p> <p>The Windows Version 2 LSO supports larger offload sizes.</p>
Large Send Offload v2 (IPv6)	*LsoV2IPv6	0 (Disabled) 1 (Enabled) (Default)	<p>Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (less than the MTU) that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the TCP checksums for each individual packet.</p> <p>IPv6 support requires LSO Version 2.</p>
Maximum Number of RSS Processors	*MaxRssProcessors	Min : 1 Max : 16 Default : None	<p>This property sets the maximum number of processors that can be used for RSS.</p>

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Maximum Number of RSS Queues	*NumRssQueues	Min : 1 Max : 12 Default : 8	This parameter defines the maximum processor number for the RSS queues on the network adapter within the given processor group. A processor group contains 64 logical processors, so this value ranges from 0 to 63. This value may be modified in conjunction with the "Rss Max Processor Group" to explicitly select the desired RSS processors for the adapter. User input will automatically be adjusted to fit within processor range of the selected group after driver restarts.
Maximum RSS Processor Number	*RssMaxProcNumber	Min : 0 Max : 63 Default : None	This parameter sets the maximum processor number for the RSS CPUs. This is the highest processor number of any processors from the RssMaxProcGroup parameter.
Network Address	Network address	Valid MAC address The default setting is None.	This overrides the permanent MAC address for the interface. The MAC address must follow this format XX:XX:XX:XX:XX:XX, where X is a hexadecimal digit (0-9 or A-F). <ul style="list-style-type: none">The address cannot be a multicast address, which has the lowest bit in the first byte set.The address cannot be all zeros. For example, 01:00:00:00:00:00 is not valid, while 02:00:00:00:00:00 is valid.
NetworkDirect	*NetworkDirect	0 Disabled 1 Enabled (default)	The Network Direct feature enables an offloaded RDMA interface for SMB 3.0 network attached storage traffic using Microsoft's SMB Direct protocol. Emulex supports ROCE. For best performance, priority flow control (PFC) should be configured on the network switch. Emulex defaults to priority 5 for ROCE traffic, although it will still work without PFC enabled. ROCE offload provides a zero copy data transfer path for SMB 3.0 traffic, offering both increased efficiency and lower latency for storage access. Both ends of the connection must support ROCE in the network adapter. ROCE is a non-routable protocol, so both servers must be in the same L2 network subnet. Microsoft currently disables Network Direct support for network adapters bound to a virtual switch.
Network Direct MTU	NdkMtu	256 512 1024 (default) 2048 4096	The MTU or frame size for ROCE traffic can be configured with this parameter.
NVGRE Task Offload (also known as Encapsulated Task Offload)	*EncapsulatedPacketTaskOffload	0 Disabled 1 Enabled (default)	NVGRE Task Offload enables the offloading of network virtualization using GRE tunneling on the NIC adapter. NVGRE offload works in conjunction with VMQ. NOTE For OCe14000-series adapters only.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Encapsulation Overhead	*EncapOverhead	Min : 0 Max : 256 Default : 0 NOTE Valid range is 0 through 256 with step of 32	Encapsulation Overhead defines the amount of overhead required in Ethernet frames due to virtual network overlay encapsulation such as VXLAN and NVGRE. Valid range is 0 through 256 with step of 32. For example 0, 32, 64, 96, 128, etc. are valid values. Effective MTU = *JumboFrame + *EncapOverhead - 14. If the effective MTU is too large for the NIC adapter, the effective MTU = *JumboFrame - 14 and SDN Host Agent will not encap tenant overlay traffic and return an error to the Network Controller.
NVGRE Encapsulated Task Offload	*EncapsulatedPacketTaskOffloadNvgre	0 (Disabled) 1 (Enabled) (Default)	Enable/disable task offloads for NVGRE encapsulated packets.
VXLAN Encapsulated Task Offload	*EncapsulatedPacketTaskOffloadVxlan	0 (Disabled) 1 (Enabled) (Default)	Enable/disable task offloads for NVGRE encapsulated packets.
VXLAN UDP destination port number	*VxlanUDPPortNumber	Min : 1 Max : 65535 Default : 4789	VXLAN UDP destination port number between 1 to 65535 inclusive.
PacketDirect	*PacketDirect	0 (Disabled) 1 (Enabled) (Default)	Enables and disables Packet Direct.
Packet Size	*JumboPacket	1514 (default) 9014 8222 4088	Configures packet size for OneConnect NIC adapters only. This parameter determines the maximum packet size transmitted and received on the interface. A 1514 byte frame size is standard, while larger packets are called jumbo frames. Using a higher frame size is generally more efficient, but it uses more system memory. A larger frame size also requires support on the network switch. Jumbo frames are IPv4-only frames; IPv6 packets will be fragmented by LSO. Switches and the peer must be configured to accept the specified packet size or the size will be negotiated to the common smallest size.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Physical Link Tracking	PLinkTrackEnable	0 (Disabled) 1 (Enabled) (Default)	<p>This parameter enables/disables physical link tracking when SR-IOV is used.</p> <p>When SR-IOV is enabled, a VEB switch is used and the driver link status does not reflect the physical link status. Use this parameter to force the driver link status to reflect the physical link status.</p> <p>By default physical link status tracking is enabled.</p> <p>Disable physical link tracking to allow PF and VF to communicate via VEB switch regardless of physical link status. When physical link status tracking is disabled, the driver always reports link as UP.</p> <p>When physical link tracking is disabled, teaming failover will not work.</p> <p>When SR-IOV is disabled, the driver always reports physical link status.</p> <p>NOTE For OCe11102, LPe16202, and OCe14000-series adapters only.</p>
Performance Tuning	FairnessMode	0 (Maximum Performance) (Default) 1 (Dynamically Balanced) 2 (Statically Balanced)	<p>This parameter selects the driver algorithm for performance tuning, allowing you to balance raw networking throughput with overall system fairness among multiple devices and applications.</p> <p>Maximum Performance - This mode maximizes the network performance for this adapter. This mode supports most cases. However, in systems with a large number of network or storage adapters, this mode may limit the performance of other devices.</p> <p>Statically Balanced - This mode configures the network adapter to throttles CPU usage in all cases, allowing more balance among hardware devices and applications. If system responsiveness is poor, this mode may improve the overall system behavior.</p> <p>Dynamically Balanced - Dynamic balancing adjusts the network adapter's performance based on system metrics, such as CPU usage. This mode can aggressively limit performance for the most stressful networking applications to ensure that all network cards can share limited computer resources, yet it can maintain maximum performance when the system has resources available.</p>
Preferred NUMA Node	*NumaNodeId	Min : 0 Max : 65535 Default : None	<p>Most modern multi-socket servers have separate memory controllers for each CPU socket. These systems have NUMA latencies for a given CPU core to access the local versus remote memory node.</p> <p>By setting this property, the driver attempts to use both memory and CPU cores from the given NUMA node.</p> <p>If the Preferred NUMA node is not set, the driver uses the preferred NUMA node as specified by the computer's BIOS.</p> <p>For best performance, the network applications must use memory and CPU affinity from the same NUMA node. This level of tuning is primarily noticeable when multiple adapters are running.</p>

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Receive Buffers	*ReceiveBuffers	Min : 64 Max : 32768 Default : 896.	This determines the number of Ethernet receive buffers allocated per receive queue. If RSS is enabled, 4 or more receive queues are used. Otherwise, a single queue is enabled. Decreasing this value will reduce the required system memory, but performance may decrease. Each buffer is the size of the "Packet Size" parameter.
Receive CPU	RxCpuPolicy	Min : 0 Max : 255 Default : None	The non-RSS receive packets are processed on this logical CPU. By default, the driver will intelligently choose a CPU in the system, so this parameter should only be set for advanced performance tuning. RSS packets will be processed by the set of RSS CPUs provided by Windows operating system. The valid values are 0 to (Number of CPUs on the System - 1).
Receive Side Scaling	*RSS	0 (Disabled) 1 (Enabled) (Default)	Receive Side Scaling (RSS) scales receive processing over multiple CPUs in parallel. This scaling typically improves application performance; however, it tends to increase CPU usage on low end machines. RSS is only supported on two primary adapters per device. It will appear disabled for additional PCI functions in blade server configurations. RSS requires Windows 2003 SP2 and later.
Recv Segment Coalescing (IPv4)	*RSCIPv4	0 (Disabled) 1 (Enabled) (Default)	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This reduces the per-packet receive processing overhead and CPU usage if standard 1514 byte sized frames are in use. NOTE If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. NOTE Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.
Recv Segment Coalescing (IPv6)	*RSCIPv6	0 (Disabled) 1 (Enabled) (Default)	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This reduces the per-packet receive processing overhead and CPU usage if standard 1514 byte sized frames are in use. NOTE If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. NOTE Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.
RoCE Mode	ForceRoutableRoceonSameSubnet	1 2 (Default)	RoCE Mode 2 brings IP based routing feature for RoCE. If there are legacy RoCE adapters in the network that does not support IP based routing, then set RoCE Mode to 1.
RSS Base Processor Group	*RssBaseProcGroup	Min : 0 Max : 32768 Default : None	This defines the base processor group for the RSS queues on the network adapter. A processor group contains 64 logical processors. This value can be modified in conjunction with the "RSS Base Processor Number" to explicitly select the desired RSS processors for the adapter.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
RSS Base Processor Number	*RssBaseProcNumber	Min: 0 Max: 63	Windows will automatically spread the RSS queues for all network cards over the available CPU cores in the computer. This parameter may be set to explicitly define the CPU affinity for the RSS queues of this device. It is the CPU number of the lowest RSS queue for this device. Hyperthreaded systems will only use the lower thread of each core for RSS. A hyperthreaded system with 16 logical processors will only use 8 RSS threads.
RSS Max Processor Group	*RSSMaxProcGroup	Min : 0 Max : 63 Default : None.	RSS Max Processor Group allows you to set the maximum number of processor groups for the RSS CPUs.
RSS Profile	*RSSProfile	1 (Closest Processor) (Default) 2 (Closest Processor Static) 3 (NUMA Scaling) 4 (NUMA Scaling Static) 5 (Conservative Scaling)	The setting determines the RSS load balancing profile implemented by Microsoft for this network adapter. The "Closest Processor" settings will tend to localize the RSS CPUs to one NUMA node, allowing the device driver to allocate memory from the local node. The "NUMA Scaling" settings will use all NUMA nodes on the system, and the memory allocation will not be specific to a particular node. The driver will ignore the Preferred NUMA node setting..
SpeedDuplex		AutoNeg (default) 10GbpsFullDuplex 1GbpsFullDuplex	SpeedDuplex is used for selecting link speed, mainly for 10GBASE-T adapters. If it is set to the default, it auto negotiates 100 Mbps/1 Gbps/10 Gbps with the switch/peer. Link speed can be forced to 1 Gbps, if option 1GbpsFullDuplex is selected. Link speed can be forced to 10 Gbps, if option 10GbpsFullDuplex is selected. 10 Gbps is the maximum supported link speed.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
SR-IOV	*Sriov	0 (Disabled) 1 (Enabled) (Default)	<p>SR-IOV enables the adapter to allocate virtual PCI functions for each virtual machine in Hyper-V. Note that the virtual switch and virtual network adapter must have SR-IOV enabled in the Hyper-V Manager. SR-IOV requires a platform with IOMMU virtualization (VT-d, AMD-Vi).</p> <p>If using SR-IOV, the Emulex NIC driver must be installed on each virtual function within the virtual machine. SR-IOV provides a direct hardware interface from the virtual machine to the networking adapter, which reduces latency and improves performance.</p> <p>The Windows Server 2012 and Windows Server 2012 R2 SR-IOV architecture establishes each Emulex virtual NIC with a corresponding emulated NIC. This allows the virtual machine to seamlessly failover to the emulated NIC if SR-IOV is disabled. It also allows Live Migration to another system, regardless of the installed NIC hardware.</p> <p>NOTE For OCe11102, LPe16202, and OCe14000-series adapters only.</p> <p>NOTE The driver currently supports the following virtual functions for the following adapter families:</p> <ul style="list-style-type: none"> OCe11100-series adapters support a maximum of 24 virtual functions per port. OCe14000-series adapters support a maximum of: <ul style="list-style-type: none"> 2-port 10 Gb: 31 virtual functions/physical function. 4-port 10 Gb: 31 virtual functions/physical function 1-port 40 Gb: 63 virtual functions/physical function
TCP Checksum Offload (IPv4)	*TCPChecksumOffloadIPv4	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & Tx Enabled) (Default)	TCP Checksum Offload (IPv4) offloads the transmit or receive IPv4 TCP checksum computation. Offloading checksums increases system efficiency.
TCP Checksum Offload (IPv6)	*TCPChecksumOffloadIPv6	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & Tx Enabled) (Default)	TCP Checksum Offload (IPv6) offloads the transmit or receive IPv6 TCP checksum computation. Offloading checksums increases system efficiency.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
TCP Connection Offload (IPv4)		Enabled Disabled (default)	<p>NOTE TCP Connection Offload is not supported on 16GFC adapters.</p> <p>If TCP offload is enabled, the device offloads the entire TCP protocol, including acknowledgement processing, retransmits, and timers. Applications that prepost receive buffers (before the data arrives) may avoid data copies in the receive path, which substantially increases the system efficiency and data rates.</p> <p>Windows does not offload TCP connections if any of the following are enabled:</p> <ul style="list-style-type: none"> n Network Load Balancing n IPSEC n Network Address Translation n NDIS 5.1 Intermediate Drivers <p>TCP offload must be enabled in the Windows operating system with the shell command: netsh int tcp set global chimney=enabled</p> <p>This parameter appears disabled if the firmware installed on your device does not support TCP connection offload. Upgrading the firmware may resolve this issue.</p> <p>View the “Statistics” property page to ensure that TCP connection offload is working.</p> <p>NOTE Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.</p>
TCP Offload Optimization		Optimize Latency Optimize Throughput (default)	<p>This parameter only applies to TCP connection offload, which must be enabled in the “Protocol Offloads” section.</p> <p>Most applications perform better with TCP Offload Optimization set to “Optimize Throughput” which handles large data transfers with minimal CPU impact.</p> <p>Setting this parameter to “Optimize Latency” causes receive data to be delivered to the application without waiting for a TCP push flag. This causes additional receive indications that typically decrease total throughput.</p>
Transmit Buffers	*TransmitBuffers	"128 256 512 1024 2048 (Default)	<p>Transmit Buffers sets the number of Ethernet transmits that may be posted to the hardware at any given time.</p> <p>The default value is sufficient to achieve maximum performance. Reducing this value conserves system memory.</p>
Transmit CPU	TxCpuPolicy	Min : 0 Max : 255 Default : None	<p>Transmit packet completion processing will be done on this CPU. By default, the driver will intelligently choose a CPU in the system, so this parameter should only be set for advanced performance tuning.</p> <p>The valid values are 0 to (Num CPUs on the System - 1).</p>
Transmit Side Scaling (TSS)	SendSideScaling	0 (Disabled) 1 (Enabled) (Default)	<p>TSS distributes transmit completions to be processed on multiple CPUs in parallel. It uses the RSS CPU table for distribution and therefore requires RSS to be enabled.</p>

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
Transmit	VMQTransmit	0 (Disabled) 1 (Enabled) (Default)	If this option is enabled with VMQs, separate transmit queues are created for each VM network interface. Send and receive interrupts for a VM network interface are processed on the same CPU(s). NOTE For OCe11102, LPe16202, and OCe14000-series adapters only.
UDP Checksum Offload (IPv4)	*UDPChecksumOffloadIPv4	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & Tx Enabled) (Default)	UDP offload checksum settings offload the transmit or receive IPv4 UDP checksum computation. Offloading checksums increases system efficiency.
UDP Checksum Offload (IPv6)	*UDPChecksumOffloadIPv6	0 (Disabled) 1 (Tx Enabled) 2 (Rx Enabled) 3 (Rx & Tx Enabled) (Default)	UDP offload checksum settings offload the transmit or receive IPv6 UDP checksum computation. Offloading checksums increases system efficiency.
Virtual Machine Queues	*VMQ	0 (Disabled) 1 (Enabled) (Default)	VMQs are dedicated hardware receive queues for virtual machines that filter receive packets based on the destination MAC address or VLAN. Receive buffers can be allocated for each queue from VM memory. This improves network throughput by distributing processing of network traffic for multiple VMs among multiple processors. It reduces CPU utilization by offloading receive packet filtering to NIC hardware. VMQs prove beneficial when four or more VMs are in use. NOTE For OCe11102, LPe16202, and OCe14000 only.
Virtual Machine Queues Lookahead Split		Enabled (default) Disabled	VMQ enables direct DMA to VM memory. Lookahead improves packet steering performance by PCI prefetching adjacent header buffer into a cache when examining a packet. Header buffers are continuous in physical memory since they belong to one pool. For OCe11102, Lookahead split requires Advanced Mode Support and is enabled in the BIOS controller configuration. NOTE For OCe11102 CNAs only. Not applicable for LPe16202 and OCe14000-series adapters. NOTE Lookahead split is not supported for jumbo frames.
Virtual Machine Queues Transmit		Enabled (default) Disabled	If this option is enabled with VMQs, separate transmit queues are created for each VM network interface. Send and receive interrupts for a VM network interface are processed on the same CPUs. Separate transmit queues increase system overall CPU utilization, but offer greater system scalability. NOTE For OCe11102 CNAs only. Not applicable for LPe16202 and OCe14000-series adapters.

Table 2 Windows Server 2016 NIC Driver Options (Continued)

Option Name	Registry Keyword	Registry Values	Definition
VLAN Identifier (802.1q)	VlanId	Not Present (default) 0 to 4094	If selected, the adapter adds a VLAN tag to all transmitted packets, and only receives packets with the matching VLAN tag. NOTE This property must not be used if the Emulex Teaming Driver is enabled. In that case, VLAN configuration must be performed in the Teaming Driver application. NOTE This property must not be used with Hyper-V. In that case, the Microsoft Hyper-V Manager must be used to configure VLANs on each virtual machine.
Wake on Magic Packet	WakeonMagic Packet	0 - Disabled 1 (Default) - Enabled	Defines if a network adapter is enabled to wake a computer on the magic packet.
Wake on Pattern Match	WakeonPatternMatch	0 - Disabled 1 (Default) - Enabled	Defines if a network adapter is enabled to wake the computer on pattern matches.

Updating Emulex OOB Drivers to a Nano Server VHD

Emulex recommends that you update the OOB drivers prior to booting to the VHD for the first time. In other words, the OOB drivers must be added to the VHD right after it has been created. This will help prevent the inbox driver from being loaded and linked to an Emulex device that is already present in the system.

To install the drivers:

1. Download the Driver Installer Kits from the following link:
<http://www.emulex.com/downloads/emulex/drivers/windows/windows-server-2016/drivers/>
2. In an elevated command prompt, navigate to the directory where the downloaded driver kit is and run the following commands to unpack the drivers:

```
elxdrvrr-nic-<VERSION>.exe /q2 extract=2
```

```
elxdrvrr-iscsi-<VERSION>.exe /q2 extract=2
```

```
elxdrvrr-fcoe-<VERSION>.exe /q2 extract=2
```

```
elxdrvrr-fc-<VERSION>.exe /q2 extract=2
```

The drivers are extracted to the current user's Documents folder. For example:

```
C:\Users\Administrator\Documents\Emulex\Drivers
```

3. In an elevated PowerShell prompt, navigate to the Nano Server VHD directory and run the following commands:

```
md mountdir
dism /Mount-Image /ImageFile:.\NanoServer.vhd /Index:1 /Mountdir:.\mountdir
dism /Add-Driver /image:.\mountdir
/driver:C:\Users\Administrator\Documents\Emulex\Drivers\NIC-<VERSION>\x64
\winserv10\ocnd65.inf
dism /Unmount-Image /Mountdir:.\mountdir /commit
```

```
md mountdir dism
/Mount-Image /ImageFile:.\NanoServer.vhd /Index:1 /Mountdir:.\mountdir
dism /Add-Driver /image:.\mountdir
/driver:C:\Users\Administrator\Documents\Emulex\Drivers\iscsi-<VERSION>\x64
\winserv10\ocnd65.inf
dism /Unmount-Image /Mountdir:.\mountdir /commit
```

```
md mountdir dism
/Mount-Image /ImageFile:.\NanoServer.vhd /Index:1 /Mountdir:.\mountdir
dism /Add-Driver /image:.\mountdir
/driver:C:\Users\Administrator\Documents\Emulex\Drivers\fcoc-<VERSION>\x64
\winserv10\ocnd65.inf
dism /Unmount-Image /Mountdir:.\mountdir /commit
```

```
md mountdir dism
/Mount-Image /ImageFile:.\NanoServer.vhd /Index:1 /Mountdir:.\mountdir
dism /Add-Driver /image:.\mountdir
/driver:C:\Users\Administrator\Documents\Emulex\Drivers\fc-<VERSION>\x64\w
inserv10\ocnd65.inf
dism /Unmount-Image /Mountdir:.\mountdir /commit
```

Known Issues

Host RDMA and Routable RoCE Known Issues

1. **Some switches strip the VLAN tag from the incoming frame with VLAN ID 0 or VLAN ID 1 values and sends the frame out without the VLAN tag, and therefore without the VLAN priority.**

Workaround

When running NIC+RoCE personality, if PFC is enabled, always configure the interface with a VLAN and make sure the VLAN ID is greater than 1.

2. **Changing the VLAN ID for the management operating system while it is running using the Hyper-V Manager is not supported.**

Workaround

Assign the required VLAN ID to the management operating system when you create the virtual switch.

3. **After a driver reload (ND is disabled on a NIC + ROCE profile or any other non-ROCE profile), throughput via SMB is limited to the highest Link Speed available as shown by the Get-SmbClientNetworkInterface PowerShell command on the client system. This issue is only seen when RDMA is disabled on the adapter and SMB uses TCP.**

Workaround

Use one of the four options below.

- Disable and enable the port of the required interface.
 - Some systems may require additional interfaces, such as Hyper-V hosts. SMB will check the interfaces to determine which can be used to connect the systems. If there are multiple connections, it will use them (RDMA or TCP). Multiple connections must all be configured the same (RDMA or TCP).
 - Reboot the system.
4. **If you are using Packet Direct, set the advanced adapter property “Maximum Number of RSS Processors” to no more than 2.**
 5. **Deleting the switch with Packet Direct enabled could result in a hang.**