

White Paper

Marvell® FastLinQ® 578xx/3400/8400/34000 Series 1GbE/10GbE/20GbE Adapters

Single Function and Switch Independent NIC Partitioning FAQs

December 2020



- Marvell® FastLinQ® Switch Independent NIC Partitioning can simplify end-to-end networking by allowing division of a physical Ethernet adapter into multiple virtual Ethernet/iSCSI/FCoE adapters, enabling dynamic allocation of bandwidth as needed while reducing the total cost of ownership.

Overview

Administrators using servers enabled with Marvell® FastLinQ® Switch Independent NIC Partitioning can leverage this feature's functionality in powerful ways. These frequently asked questions are designed to provide a better understanding of Switch Independent NIC Partitioning. They provide fundamental information, sometimes about rather complex subjects, all of which can be described in further detail in other collateral available on the Marvell website at: <https://www.marvell.com/support/doc-library>

Frequently Asked Questions (FAQ)

Q1: What are the Key Features of Marvell FastLinQ 578Xx/3400/8400/34000 Series 1Gbe/10Gbe/20Gbe Adapters?

- Ethernet
 - 100BASE-T/1000BASE-T/10GBASE-T (RJ45 interfaced adapters only)
 - 1Gb SFP SX/LX/CX
 - 10Gb SFP+ SR/LR/CR
 - Blade 1Gb KX/10Gb KR/20Gb KR2
- Windows Server 2012/R2/2016 (or later) Receive Segment Coalescing (RSC)
- Data Center Bridging (DCB) (see Table 1)
 - Priority-based Flow Control (PFC), IEEE 802.1Qbb
 - Enhanced Transmission Selection (ETS), IEEE 802.1Qaz
 - Data Center Bridging Exchange (DCBx), IEEE 802.1Qaz
- Fiber Channel over Ethernet (FCoE) Hardware Offload
 - Over 2,500,000 IOPS per dual-ported device
- Internet Small Computer System Interface (iSCSI) Hardware Offload
 - iSCSI Type-Length Value (TLV) for DCB
 - Over 1,500,000 IOPS per dual-ported device
- Switch Independent NIC Partitioning (NPAR)
- Single-Root I/O Virtualization (SR-IOV)
 - Supported on Windows Server 2012/R2 /2016 Hyper-V, Linux RHEL KVM/SLES XEN, Citrix Xen Server, and VMware ESXi
 - Total of 64 SR-IOV Virtual Functions (VFs) per Single Function (SF) port (single and dual port adapters)/ total of 32 VFs per SF port (quad port adapters)
 - Total of 64 SR-IOV VFs per NPAR port on single and dual port adapters/total of 32 VFs per NPAR port on quad port adapters
 - VFs are user divisible (in groups of 8) over an NPAR port's partitions

- No VFs can be assigned to a iSCSI-Offload/FCoE-Offload storage enabled NPAR partition
- FCoE Boot from Storage Area Networking (SAN) on Windows, Linux, and VMware ESXi 5.5 or later
- iSCSI Boot on Windows, Linux, and VMware
- Virtual Extensible LAN (VXLAN) Packet Task Offloads on:
 - VMware ESXi 5.5/6.x (or later) with UDP/TCP RSS Acceleration, NetQueue Acceleration, TX TCP Segmentation Offload (TSO), and TX Checksum Offload (CO)
 - Linux KVM/XEN with UDP/TCP RSS Acceleration, MQ Acceleration, LSO, and TX CO
 - Windows Server 2016 (or later) with UDP/TCP RSS Acceleration, VMQ Acceleration, LSO, and TX CO
- Windows Server 2012 R2 (or later) Network Virtualization using Generic Routing Encapsulation (NVGRE) Packet Task Offloads with RSS Acceleration, VMQ Acceleration, LSO, and TX CO
- GENEVE Network Virtualization Encapsulation (GENEVE) Packet Task Offloads on:
 - VMware ESXi 6.5 (or later) with UDP/TCP RSS Acceleration, NetQueue Acceleration, TSO, and TX CO
 - Linux KVM/XEN with UDP/TCP RSS Acceleration, MQ Acceleration, LSO, and TX CO
- RX/TX multiqueue
 - Windows Virtual Machine Queues (VMQs)
 - Windows Server 2012 (or later) Dynamic VMQs (dVMQs)
 - VMware NetQueues
 - Linux Multi-Queues (MQs)
- Receive-Side Scaling (RSS)/Transmit-Side Scaling (TSS)
- TX/RX TCP/IP/UDP CheckSum Offload (CO)
- TCP Segmentation Offload (TSO)
- Giant Send Offload (GSO)
- Large Send Offload (LSO)/Large Receive Offload (LRO)
- Transparent Packet Aggregation (TPA) hardware coalescing used in conjunction with RSC and LRO
- 256 VM virtual Fibre Channel (vFC) instances per FCoE adapter port in Windows Server 2012 (or later) Hyper-V
- Linux, VMware and Windows FCoE 256 N_Port ID Virtualization (NPIV) instances

Q2: What is Marvell FastLinQ Switch Independent NIC Partitioning?

- Marvell FastLinQ Switch Independent NIC Partitioning allows a single physical adapter to be partitioned into multiple virtual Ethernet/iSCSI - Offload/FCoE - Offload adapters, which present additional PCI functions that appear as networking or storage devices to the operating system.
- Marvell FastLinQ Switch Independent NIC Partitioning simultaneously supports up to:
 - Eight (dual-port and quad-port) or four (single-port) partitions
 - Plus concurrently, two (single-port) or four (dual-port) or eight (quadport) storage offload (iSCSI and/or FCoE) partitions - see Q6 and Q7 for more details.
 - All with user-configurable relative (minimum) weight and maximum bandwidth allocation for traffic shaping and Quality of Service (QoS) control.

Q3: How Many Partitions are Created Per Port?

- A maximum of four partitions are created per port on the single/dual port adapters and two on the quad port adapters - see Figures 1 through 3. For example, there are two ports on the QLE3442/QLE8442 device, so there are a total of eight partitions made available in Switch Independent NIC Partitioning mode. FCoE-Offload/iSCSI-Offload adapters (such as the QLE8442) have two MAC addresses per partition, the second is for iSCSI-Offload and FCoE-Offload FIP and DCBX LLDPs.

- The fixed speed 1GbE ports of the 57800 do not present extra partitions in NPAR mode so there are two partitions created per 10G port with a total of six partitions per device - see Figure 4.

DUAL Port Device	PCIs Location / Bus / Device / Function	Ethernet and iSCSI-Offload MAC Address	Example IP Address
1st port			
Port 0 Partition 1	X/Y/Z/0	00:10:18:88:E7:A8 00:10:18:88:E7:A9	1.1.1.1
Port 0 Partition 2	X/Y/Z/2	00:10:18:88:E7:AC 00:10:18:88:E7:AD	2.2.2.1
Port 0 Partition 3	X/Y/Z/4	00:10:18:88:E7:B0 00:10:18:88:E7:B1	3.3.3.1
Port 0 Partition 4	X/Y/Z/6	00:10:18:88:E7:B4 00:10:18:88:E7:B5	4.4.4.1
2nd port			
Port 1 Partition 1	X/Y/Z/1	00:10:18:88:E7:AB	5.5.5.1
Port 1 Partition 2	X/Y/Z/3	00:10:18:88:E7:AE 00:10:18:88:E7:AF	6.6.6.1
Port 1 Partition 3	X/Y/Z/5	00:10:18:88:E7:C0 00:10:18:88:E7:C3	7.7.7.1
Port 1 Partition 4	X/Y/Z/7	00:10:18:88:E7:C6 00:10:18:88:E7:C7	8.8.8.1

Figure 1. DUAL-Port PCIe to NPAR with iSCSI Device Diagram

Single Port Device	PCIs Location / Bus / Device / Function	Ethernet and iSCSI-Offload MAC Address	Example IP Address
10GbE port			
Port 0 Partition 1	X/Y/Z/0	00:10:18:88:E7:A8 00:10:18:88:E7:A9	1.1.1.1
Port 0 Partition 2	X/Y/Z/2	00:10:18:88:E7:AC 00:10:18:88:E7:AD	2.2.2.1
Port 0 Partition 3	X/Y/Z/4	00:10:18:88:E7:B0 00:10:18:88:E7:B1	3.3.3.1
Port 0 Partition 4	X/Y/Z/6	00:10:18:88:E7:B4 00:10:18:88:E7:B5	4.4.4.1

Figure 2. Single-Port PCIe to NPAR with iSCSI Device Diagram

57800 QUAD 10GbE Port Device	PCIs Location / Bus / Device / Function	Ethernet and iSCSI-Offload MAC Address	Example IP Address
1st 10GbE port			
Port 0 Partition 1	X/Y/Z/0	00:10:18:88:E7:A8 00:10:18:88:E7:A9	1.1.1.1
Port 0 Partition 2	X/Y/Z/4	00:10:18:88:E7:B0 00:10:18:88:E7:B1	2.2.2.1
2nd 10GbE port			
Port 1 Partition 1	X/Y/Z/1	00:10:18:88:E7:AA 00:10:18:88:E7:AB	3.3.3.1
Port 1 Partition 2	X/Y/Z/5	00:10:18:88:E7:C0 00:10:18:88:E7:C3	4.4.4.1
3rd 10GbE port			
Port 2 Partition 1	X/Y/Z/2	00:10:18:88:E7:AC 00:10:18:88:E7:AD	5.5.5.1
Port 2 Partition 2	X/Y/Z/6	00:10:18:88:E7:B4 00:10:18:88:E7:B5	6.6.6.1
4th 10GbE port			
Port 3 Partition 1	X/Y/Z/3	00:10:18:88:E7:AE 00:10:18:88:E7:AF	7.7.7.1
Port 3 Partition 2	X/Y/Z/7	00:10:18:88:E7:C6 00:10:18:88:E7:C7	8.8.8.1

Figure 3. Quad-Port PCIe to NPAR with iSCSI Device Diagram

57800 DUAL 10GbE and DUAL 1GbE Device	PCIs Location / Bus / Device / Function	Ethernet and iSCSI-Offload MAC Address	Example IP Address
1st 10GbE port			
Port 0 Partition 1	X/Y/Z/0	00:10:18:88:E7:A8 00:10:18:88:E7:A9	1.1.1.1
Port 0 Partition 2	X/Y/Z/4	00:10:18:88:E7:B0 00:10:18:88:E7:B1	2.2.2.1
2nd 10GbE port			
Port 1 Partition 1	X/Y/Z/1	00:10:18:88:E7:AA 00:10:18:88:E7:AB	3.3.3.1
Port 1 Partition 2	X/Y/Z/5	00:10:18:88:E7:C0 00:10:18:88:E7:C3	4.4.4.1
1st 1GbE port			
Port 2	X/Y/Z/2	00:10:18:88:E7:AC	5.5.5.1
2nd 1GbE port			
Port 3	X/Y/Z/3	00:10:18:88:E7:AE	6.6.6.1

Figure 4. 57800 PCIe to NPAR with iSCSI Device Diagram

Q4: What are the Supported Operating Systems?

- The Switch Independent NIC Partitioning mode supported OSs are Windows Server 2012 (or later), Windows Server 2012 (or later) Hyper-V, RHEL (v6.8 or later and v7.4 or later with or without KVM hypervisor), SUSE Linux Enterprise Server 11 (SLES11) SP 3/12 (SLES12) SP2 or later (with or without XEN hypervisor), Citrix® XenServer v6.5/7.0 or later, and VMware® ESXi 5.5/6.x or later as shown in Table 1 - an earlier software (driver/firmware) version combination maybe need for some of the older Linux/VMware/Windows OS versions. The fully featured drivers might not be included in box.

Q5: How Many Ethernet Protocol Device Partitions are Allowed?

- Multiple Ethernet protocols can be enabled per port for a maximum of four on the single-port, six on the 57800, and eight on the dual-port / quad-port devices.
- For Windows OSs, the Ethernet protocol can be enabled or disabled with or without any other offload protocols being enabled on the same partition. The system virtual bus device driver instance is always present. On some 578xx series devices, Windows Server 2012 and Windows Server 2012 R2 TCP Offload Engine (TOE) or TCP Chimney Offload is available on Ethernet protocol enabled partitions (in both SF and NPAR modes). TOE is not enable-able on the QLE3400 or QLE8400 or QL34000 series adapters.
- For Linux®, Citrix XenServer, and VMware ESXi OSs, the Ethernet protocol for all partitions is always enabled, ignoring what might have been set by the various management utilities.

Table 1. Protocols Available in Operating Systems vs. SF and Switch Independent NIC Partitioning Mode

Table 1. Protocols Available in Operating Systems vs. SF and Switch Independent NIC Partitioning Mode						
Operating System	SF Mode			Switch Independent NIC Partitioning Mode		
	L2 Ethernet ^b	iSCSI Offload ^b	FCoE Offload	L2 Ethernet ^b	iSCSI Offload ^b	FCoE Offload
Windows 2012 ^a	Yes	Yes ^c	Yes ^g	Yes	Yes ^c	Yes ^g
Windows 2012 Hyper-V ^a	Yes ^d	Yes ^c	Yes ^{g,h}	Yes ^d	Yes ^c	Yes ^{g,h}
Windows 2012 R2 ^a	Yes	Yes ^c	Yes ^g	Yes	Yes ^c	Yes ^g
Windows 2012 R2 Hyper-V ^a	Yes ^{d,e}	Yes ^c	Yes ^{g,h}	Yes ^{d,e}	Yes ^c	Yes ^{g,h}
Windows 2016 ^a	Yes	Yes ^c	Yes ^g	Yes	Yes ^c	Yes ^g
Windows 2016 Hyper-V ^a	Yes ^{d,e,f}	Yes ^c	Yes ^{g,h}	Yes ^{d,e,f}	Yes ^c	Yes ^{g,h}
Citrix XenServer 6.5+(or later)	Yes ^d	No	No	Yes	No	No
CentOS 7.4+ ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^{d,f}	Yes ^c	Yes ^g
RHEL 6.8+ ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^{d,f}	Yes ^c	Yes ^g
RHEL 7.3+ ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^{d,f}	Yes ^c	Yes ^g
SLES 11 SP3+ ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^{d,f}	Yes ^c	Yes ^g
SLES 12 SP2+ ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^{d,f}	Yes ^c	Yes ^g
Ubuntu 14.04.5 LTS ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^f	Yes ^c	Yes ^g
Ubuntu 16.04.3 LTS ^a	Yes ^{d,f}	Yes ^c	Yes ^g	Yes ^f	Yes ^c	Yes ^g
VMWare ESXi 5.5 ^a	Yes ^{d,f}	Yes ^c	Yes	Yes ^{d,f}	Yes ^c	Yes
VMware ESXi 6.0 ^a	Yes ^{d,f}	Yes ^c	Yes	Yes ^{d,f}	Yes ^c	Yes
VMware ESXi 6.5+ ^a	Yes ^{d,f,i}	Yes ^c	Yes ^j	Yes ^{d,f,i}	Yes ^c	Yes ^j

a. DCB (DCBX/PFC/ETS Supported)

b. Jumbo Frame Supported

c. Lossless iSCSI-offload over DCB (iSCSI-TLV) Supported

d. SR-IOV Supported

e. NVGRE Packet Task Offloads Supported

f. VXLAN Packet Task Offloads Supported

g. NPIV Supported

h. vFC Supported

i. GENEVE Packet Task Offloads Supported

j. NPIV Supported with native VMware qlf3f driver

Q6: How Many iSCSI Offload Protocol HBA Device Partitions are Allowed in NPAR Mode?

- Two iSCSI offload protocol instances can be enabled per port (if FCoE-Offload is not enabled) on any two of the partitions for a total maximum of two on the single-port devices, a total of four on the dual-port devices, and a total of eight on the quad port 57840 devices. The 57800 supports up to a total of two on the 10GbE ports only - the two fixed 1GbE ports do not support iSCSI offload.
- In Windows OSs, the iSCSI offload instances can be enabled with or without the Ethernet protocol being enabled on the same partition.
- In Linux and VMware OSs, the iSCSI offload protocol is enabled with the Ethernet protocol being enabled on the same partition.

Q7: How Many FCoE Offload Protocol HBA Device Partitions are Allowed in NPAR Mode?

- One FCoE offload protocol instances can be enabled per port on any one of the partitions for a total maximum of one on the single-port devices, two on the dual-port devices, and a total of four on the quad port 57840 devices. The 57800 supports up to a total of two on the 10GbE ports only - the two fixed 1GbE ports do not support FCoE offload.
- In Windows OSs, the FCoE offload protocol instances can be enabled with or without the Ethernet protocol being enabled on the same partition.
- In Linux and VMware ESXi OSs, the FCoE offload protocol instances is enabled with the Ethernet protocol also being enabled on the same partition.
- Additionally, one iSCSI offload protocol instances can be enabled per port on any of the remaining partitions.
- A maximum of two adapter offload protocols can be enabled per port on the Marvell FastLinQ Adapter.

Q8: How Should the Available Traffic Type Protocols be Configured for the Various Modes?

- For Switch Independent NIC Partitioning mode of operation using the iSCSI non-offload pathway initiator mode on Windows, enable the Ethernet protocol for the specific port's partition. On Linux, Citrix XenServer, and VMware ESXi the Ethernet (i.e., L2) protocol is always enabled.
- For Switch Independent NIC Partitioning mode of operation with iSCSI hardware offload adapter initiator mode on Linux/ Windows/ VMware, enable the iSCSI offload protocol for the specific port's partition.
- For Switch Independent NIC Partitioning mode of operation with FCoE hardware offload adapter initiator mode on Windows/Linux/ VMware, enable the FCoE offload protocol for the specific port's partition.
- For Single Function (SF) mode of operation, all three protocols (Ethernet, iSCSI, and FCoE offload) are simultaneously available by default on the 10GbE /20GbE ports for Windows, VMware and Linux OSs.

Note: For Windows, the offload protocol resource (in the Marvell FastLinQ management utilities) might need to be enabled for the specific offload device to be detected/enumerated when the device is first installed. Refer to the User's Guide for details.

Q9: How Many iSCSI Offload Protocol HBA Device Partitions are Allowed in SF Mode?

- One iSCSI offload protocol instance can be enabled per port on the single function (partition) for a total maximum of one on the single-port devices, two on the dual-port devices, and a total of four on the quad port 57840 devices. The 57800 supports up to a total of two on the two 10GbE ports only - the two fixed 1GbE ports do not support iSCSI offload.
- In SF mode, both FCoE Offload and iSCSI Offload can be enabled on the same single function (partition) of a single port, that means a maximum of two adapter offload protocols (one FCoE Offload and one iSCSI Offload in addition to Ethernet) can be enabled per port on the Marvell FastLinQ Adapter.

Q10: How Many FCoE Offload Protocol HBA Device Partitions are Allowed in SF Mode?

- One FCoE offload protocol instance can be enabled per port on the single function (partition) for a total maximum of one on the single-port devices, two on the dual-port devices, and a total of four on the quad port 57840 devices. The 57800 supports up to a total of two on the two 10GbE ports only - the two fixed 1GbE ports do not support FCoE offload.
- In SF mode, both FCoE Offload and iSCSI Offload can be enabled on the same single function (partition) of a single port, that means a maximum of two adapter offload protocols (one FCoE Offload and one iSCSI Offload in addition to Ethernet) can be enabled per port on the Marvell FastLinQ Adapter.

Q11: What Does a Partition's Relative Bandwidth Weight Value Mean?

- Each individual partition's Relative Bandwidth Weight value can be from 0% to 100%, but the sum total of the Relative Bandwidth Weights for all partitions on the same port must equal either 0% or exactly 100%. The recommended Relative Bandwidth Weight value range is either all 0% in the case where the total sum of all partitions (on that port) Relative Bandwidth Weight is equal to 0% OR no lower than 10% on a single partition in the case where the total sum of all partitions (on that port) Relative Bandwidth Weight is equal to 100%.
- The Relative Bandwidth Weight value is more than a simple minimum bandwidth setting. It is used to dynamically determine the bandwidth ratio of a partition's send traffic versus the other partition's sending traffic. The send traffic is approximately that partition's Relative Bandwidth Weight value divided by the sum of all the other current actively sending partition's Relative Bandwidth Weight values. This assumes each individual actively sending partition is able to reach that amount over their individual sending applications, and that partition's Maximum Bandwidth value is not restricting the flow below that amount. If an individual partition's send traffic flow cannot fill the level allocated to it, the other actively sending partitions will attempt to use the surplus bandwidth within the constraints of their Relative Bandwidth Weight and Maximum Bandwidth settings.
- Setting the individual partition's Relative Bandwidth Weights to equal a sum other than 0% or 100% will result in undefined behavior.
- Setting the Relative Bandwidth Weight value to all 0% is not exactly the same as setting all of the partition values to 25%. All 0% means each individual traffic flow type (no matter what partition it is on) gets equal weight with respect to the other traffic flows, while setting them all to 25% means each partition's aggregated traffic flow gets equal weight with respect to each of the other partition's aggregated traffic flows. The three types of traffic flows are Ethernet or iSCSI-Offload or FCoE-Offload.
- If the sum of the partitions on the same port is not 0%, and one or more of the partitions have a 0% Relative Bandwidth Weight value, then the Relative Bandwidth Weight value used by the sending bandwidth allocation logic is actually 1%. This means that the send traffic percent allocated to a partition with 0% is never zero. The recommended minimum Relative Bandwidth Weight value is no less than 10% on any individual partition, unless all partitions are set to 0%.

Q12: What is a Partition's Maximum Bandwidth Value Range?

- A partition's Maximum Bandwidth value range is 1 to 100 in increments of 1% of the port's current Link Speed (in 200 Mbps increments for a 20GbE link, 100 Mbps increments for a 10GbE link and in 10 Mbps increments for a 1GbE link)
- The 57800 1GbE fixed speed ports do not support NPAR mode and the associated NPAR QoS settings but a higher speed port could be configured to run at 1GbE and that would still support NPAR mode.
- The sum total of the partitions' Maximum Bandwidth values on the same port does not have to equal 100. The allowable sum-total value range on a port with four partitions is from 4 (undersubscribed – all four partitions set to 1) to 400 (oversubscribed – all four partitions set to 100) while a port with two partitions would be 2 to 200.
- If all partitions were set to undersubscribe the available bandwidth, then all unassigned send/transmit direction bandwidth would be unavailable for use by any of the partitions. For example, if the Maximum Bandwidth sum of the partitions on a port was 4%, then the unassigned 96% of the available bandwidth would be unavailable and could not be used.

Q13: How is the Switch Independent NIC Partitioning Transmit Direction Traffic Flow Rate Affected by Different Mode Settings?

- In non-DCB mode, where the sum of the partition's Relative Bandwidth Weights equals zero (i.e., each partition's Relative Bandwidth Weight is set to zero), each individual traffic flow – in this example, Partition 1's iSCSI (P1i) and Ethernet (P1e) through Partition 4's Ethernet (P4e) and iSCSI (P4i) – is equally scheduled to transmit within the limitations of the partition's Maximum Bandwidth and the overall connection's link speed as shown in Figure 4. This means if the Maximum Bandwidth of a specific partition is set to less than 100%, then the traffic flows sharing that partition will be further restricted to where their combined traffic flow bandwidth will be capped by that per partition setting. The actual inter-partition ratio of the two sharing traffic flows is controlled by the host OS. Think of the Maximum Bandwidth as a fixed-sized funnel with the OS determining how the two sharing traffic types are pouring into that funnel. If all four partitions' individual Maximum Bandwidths are set to 100% (i.e., they are all unrestricted), then each actively

sending traffic flow (without regard to which partition they are on) will equally share the transmit directions total bandwidth (i.e., TX link speed).




Figure 4. Non-DCB Mode Zero Sum Relative Bandwidth Weight TX Traffic Flow

- In non-DCB mode, where the sum of the partition's Relative Bandwidth Weights equals 100, each Partition's combined traffic flow is equally scheduled to transmit within the limitations of the partition's Relative Bandwidth Weight and Maximum Bandwidth settings and the overall connection's link speed as shown in Figure 5. This means a specific partition's Relative Bandwidth Weight value will restrict the traffic flows sharing that partition's bandwidth allocation, as one combined traffic flow with respect to the other actively sending partitions. The partition's send flow rate is based on the ratio of that partition's individual weight versus the aggregated weights of all the other actively sending partitions. Furthermore, each partition's combined traffic flow will be capped by that partition's Maximum Weight setting. The actual inter-partition ratio of the two sharing traffic flows is controlled by the host OS. Think of the dynamic weight ratio as a variable sized funnel that could be further restricted by the Maximum Bandwidth fixed-sized funnel with the OS determining how the sharing traffic types are pouring into the combined funnels. The recommended Relative Bandwidth Weight value range in this case is no lower than 10% on a single partition.

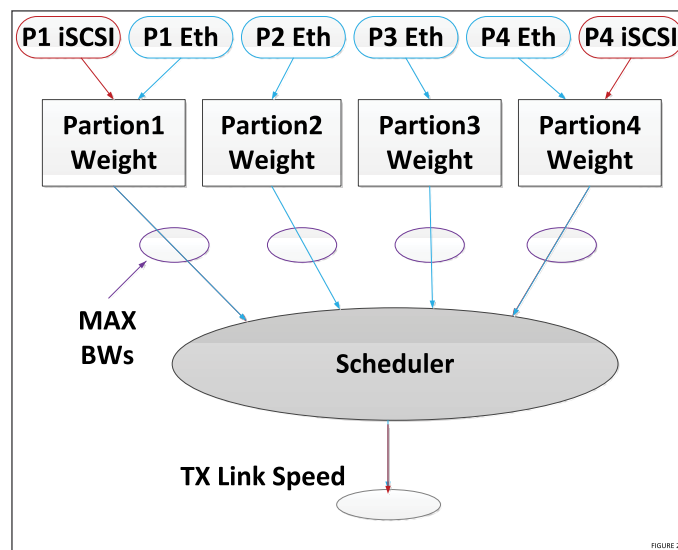


Figure 5. Non-DCB Mode 100% Sum Relative Bandwidth Weight TX Traffic Flow

- In DCB mode, all of the Partition's Relative Bandwidth Weights are disregarded and the individual traffic flows are scheduled to transmit within the limitations of the Priority Group's ETS value (determined by its Traffic Type). This example shows the first Partition has both an FCoE-enabled and Ethernet-enabled protocols transmitting to two different Priority Groups. Each partition's Maximum Bandwidth setting and the overall connection link speed as shown in Figure 6. The FCoE traffic type is in PG1, and all of

the other traffic types (iSCSI offload and Ethernet) are in the PG0. Each Priority Group has its own ETS value. Similar to the other two rate-controlling modes, the host OS determines the actual inter-partition traffic ratio for the cases where two traffic types share the same partition and are in the same Priority Group. Refer to the NPAR user guide's examples for details.




Figure 6. DCB Mode TX Traffic Flow with Lossless FCoE, Lossy iSCSI, and Lossy Ethernet

- This second example of DCB mode shows the first and fourth Partitions having both an iSCSI-enabled and Ethernet-enabled protocol transmitting to two different Priority Groups and each partition's Maximum Bandwidth setting and the overall connections link speed as shown in Figure 7. In this example, the lossless iSCSI Offload over DCB traffic (also known as iSCSI-TLV) type is in PG1, and all of the other traffic types (specifically Ethernet) are in the default PG0. Each Priority Group has its own ETS value. The two partition's lossless iSCSI offload go through its assigned Priority Group while the Lossy Ethernet traffic goes through the other.




Figure 7. DCB Mode TX Traffic Flow with Lossless iSCSI and Lossy Ethernet

- This third example of DCB mode shows the first Partition again having both an FCoE-enabled and Ethernet-enabled protocol transmitting to two different Priority Groups while the fourth Partition has the iSCSI enabled and Ethernet-enabled protocols also transmitting to two different Priority Groups and each partition's Maximum Bandwidth setting and the overall connections link speed as shown in Figure 8. In this example, the lossless FCoE offload over DCB traffic type (from Partition 1) is in PG1, the lossless

iSCSI offload over DCB traffic (also known as iSCSI-TLV) type (from Partition 4) is in PG2, and all of the other traffic types (specifically the Ethernet from Partitions 1 through 4) are in the default PG0. Each Priority Group has its own ETS value. The two different lossless offloads (FCoE and iSCSI) go through their assigned Priority Groups while the Lossy Ethernet traffic goes through another.




Figure 8. DCB Mode TX Traffic Flow with Lossless iSCSI, Lossless FCoE, and Lossy Ethernet

Note: An adapter with only two partitions per port (such as a quad port device) would act similarly, but would only have two partitions instead of the four shown in the preceding examples.

Q14: How is the Receive Direction Traffic Flow Rate Affected by Different Switch Independent NIC Partitioning Settings?

- A partition’s Relative Bandwidth Weight and Maximum Bandwidth values affect only its send/transmit/outgoing/egress traffic.
- The receive direction traffic of the partitions of a single-port adapter are not weighted or restricted by the Marvell FastLinQ 578xx/3400/8400/34000 device Switch Independent NIC Partitioning QoS settings. The receive direction acts in a similar way when in SF mode: the incoming traffic on the port is controlled by the switch port to which it is connected.

Q15: Does the Marvell FastLinQ Series Support DCB Enhanced iSCSI (iSCSI Offload TLV Over DCB)?

- Yes, the Marvell FastLinQ Series Adapter can assign iSCSI offload traffic to be Lossless and to its own dedicated DCB traffic class priority group independent of the other TCP/IP networking or FCoE offload traffic, in both SF and Switch Independent NIC Partitioning modes.

Q16: How Does a Partition’s Maximum Bandwidth Value affect the Displayed Link Speed of the Partition Adapter?

- A partition’s Maximum Bandwidth value is used by some OS applications to display the partition device/adapter link speed. This value represents the specific partition’s send/transmit/ outgoing maximum rate and not the receive/incoming link speed rate, which is always the port’s current link speed.

Q17: What are the Similarities and Differences Between SF and Switch Independent NIC Partitioning Modes for Advanced Features Such as Stateless Offload, Teaming, VLAN, Etc.?

Refer to Table 2.

Q18: How are PXE/iSCSI/FCOE Remote Boots affected by Switch Independent NIC Partitioning Mode?

- Remote boots normally use the first partition of each port (as if not in partitioned mode) as that boot port's MAC address. This means the applicable offload should be enabled on the booting first partition of a port.
- Windows, VMware ESXi and Linux all support PXE/iSCSI/FCoE remote boots.
- In SF or Switch Independent NIC Partitioning modes, PXE booting of Windows Server 2012/R2 Hyper-V VMs over a Marvell FastLinQ QLASP NIC teaming driver Smart Load Balanced and Failover (SLB) Team is supported. The HyperV Mode must be enabled when creating the Smart Load Balancing™ team. This is not applicable when using the Windows Server 2012 (or later) in-OS NIC teaming load-balancing and failover service. The Marvell FastLinQ QLASP NIC teaming driver is not supported on Windows Server 2016 (or later).

Q19: How is Wake On Lan (WOL) Out Of Band (OOB) Affected by Switch Independent NIC Partitioning Mode?

- Windows/Linux/VMware ESXi WOL – OOB Advanced Configuration and Power Interface (ACPI) S1/S3/S4/S5 states are supported on the first partition of each port and for 1 Gbps link speeds only.
- Not all adapters or adapter ports support WOL.

Q20: How is NON-OOB WOL Affected by Switch Independent NIC Partitioning Mode?

- Windows WOL ACPI S4/S5 states are supported for all partitions of both ports at 1 Gbps link speed only.
- Linux/VMware ESXi WOL ACPI S5 state is supported for all partitions on both ports at 1 Gbps link speed only.
- Not all adapters or adapter ports support WOL.

Q21: How are the Switch Independent NIC Partitioning Settings Modified in the Operating System?

- In Linux/Windows/VMware OSs, you can change and enable/ disable the partitioning mode or settings in the operating system using the Marvell FastLinQ management utilities. These can make changes to Maximum Bandwidth and Relative Bandwidth Weight settings without needing a reboot, but traffic maybe interrupted on the affected ports when the changes are applied. Enabling/ disabling NPAR mode or changing which protocol is enabled on a partition require a reboot to apply.
- In all OS's you can change and enable/disable Switch Independent NIC Partitioning settings in pre-boot Marvell FastLinQ CCM / preboot UEFI HII system BIOS device configuration pages (if supported on your server).
- In newer versions of EthTool, the "ethtool -s ethX speed NNNNN" command supports setting the Marvell FastLinQ 578xx/3400/8400/34000 Series Adapter partition's Maximum Bandwidth value without a system reboot, where ethX is the partition and NNNNN is the speed which is in 1 Mbps increments (i.e., it is NOT a percentage):
 - For a 10GbE link connection you could set it from 100 (which is equivalent to 100 Mbps or 1% of the 10GbE link speed) to 10000 (which is equivalent to 10 Gbps or 100% of the 10GbE link speed) –the last two zero positions are not used.
 - For a 1GbE link connection you could set it from 10 (which is equivalent to 10 Mbps or 1% of the 1GbE link speed) to 1000 (which is equivalent to 1 Gbps or 100% of the 1GbE link speed) – the last zero position is not used.
- The Windows and VMware in-OS utilities can also be used to change a partition Maximum Bandwidth QoS settings.
- In some cases, a system reboot is required before any changes take effect.

Q22: How are Windows QLASP Teaming, Windows Server 2012 (or Later) In-OS NIC Teaming, and Linux/VMware Bonding Affected by Switch Independent NIC Partitioning Mode?

- A team should not contain more than one partition from the same port.
- Switch Independent NIC Partitioning adapters can be teamed with other non-Switch Independent NIC Partitioning adapters, such as other Marvell adapters or with other vendor Ethernet adapters.
- The Windows Server 2012 (or later) NIC Teaming, Linux/ VMware Bonding and Marvell FastLinQ QLASP NIC teaming driver on Windows Server 2012/R2 switch independent teaming modes all support Load Balanced and Failover over Switch Independent NIC Partitioned adapters.
- Switch dependent teaming [IEEE 802.3ad LACP and Generic/ Static Link Aggregation (Trunking)] cannot use a Switch Independent NIC Partitioning virtual adapter. This is due to the way the IEEE standards require Switch Dependent Teaming [IEEE 802.3ad LACP and Generic/Static Link Aggregation (Trunking)] mode to work over the entire port instead of MAC address (fraction of a port) granularity.
- FCoE and iSCSI traffic connections use Multipath I/O (MPIO) to provide Load Balancing and Failover protection at the OS's storage level when there is more than one connection to the same storage LUN.
- FCoE and iSCSI offload should not be enabled on Switch Dependent teamed ports, except for FCoE offload on LACP teams on Virtual Port Channel (vPC) configured Cisco switches.
- FCoE and iSCSI offload can be enabled on switch independent NIC teamed ports but is not regulated by the teaming software.

Q23: Is IEEE 802.3X Link-Level Flow Control Supported in SF and Switch Independent NIC Partitioning Mode?

- Yes. The per physical port link-level IEEE Pause flow control is available in both SF and Switch Independent NIC Partitioning modes of operation when DCB's PFC mode is not enabled.
- In non-DCB SF mode, link-level flow control is configured on a per-port basis through the Advanced Properties tab of the NDIS driver of Windows Device Manager, through the Marvell FastLinQ management utilities or through the Linux/VMware in-OS utilities.
- In non-DCB Switch Independent NIC Partitioning mode, link-level flow control is configured on a per-port basis through the Marvell FastLinQ management utilities. It is not configured on a per-partition basis since multiple partitions share the same port, and link-level flow control operates on a per-port basis. This means that the in-OS utilities (such as Windows Advanced Properties or Linux Ethtool utility or VMware in-OS utilities) cannot be used to control link-level flow control in Switch Independent NIC Partitioning mode-this needs to be done using the Marvell FastLinQ management utilities.

Q24: What Criteria is Used to Generate IEEE 802.3X Link-Level Flow Control Pause Frames?

- By default, link-level flow control will only generate pause frames in response to a temporary lack of internal hardware buffering resources.
- Due to the nature of Converged NICs' networking (L2 Ethernet) host buffers and storage (iSCSI/FCoE hardware offloaded) host buffers are maintained separately. If networking host buffers are too few and are allowed to generate link-level pause frame, storage traffic would also be paused. As a result, the Marvell FastLinQ 578xx/3400/8400/34000 Series Adapters do not generate link-level pause frames when the L2 host buffers are limited but instead drop frames destined for the host buffer limited function.
- Windows, VMware, and Linux all provide an option to restore the previous link-level flow control behaviour Setting this to 1 - 4 will set a fixed number of RSS queues to be used by VXLAN tunneled traffic or the HOST. of generating pause frames in response to limited L2 host buffers. In Windows, the option is called "Pause On Exhausted Host Ring" (Windows Device Manager NDIS Advanced Properties) for each port. In Linux, use the module option parameter named "drople3s_fc" (using the "modprobe bnx2x drople3s_fc=1" or equivalent "insmod" commands). In VMware, use the command "esxcfg-module -s drople3s_fc=1 bnx2x" (ESXi 6.5 and earlier) and "esxcfg-module-s drople3s_fc=1 qfle3" (VMware ESXi 6.5 and later). You can also use one of the Marvell FastLinQ management utilities on the specific OS.

- To enable VxLAN with UDP/TCP RSS acceleration: `cfg-module -s RSS=N qfle3 wheesxre` “N” can be 0 – 4. Setting this to 0 will allow the software to automatically control the number of RSS queues used (default setting). Setting this to 1 - 4 will set a fixed number of queues. Both Linux and VMware commands control all partitions of both ports of all adapters in the server simultaneously. When enabled, link-level pause frames will be generated if any host buffers on the designated partition are low, resulting in paused traffic to all functions on all partitions. In non-DCB Switch Independent NIC Partitioning mode, in Windows, all partitions’ Pause On Exhausted Host Ring settings must be enabled for it to be ON or disabled for it to be OFF, whereas in Linux and VMware all Marvell FastLinQ 578xx/3400/8400/34000 Series bnx2x (qfle3 on VMware ESXi 6.5 and later) devices on the host system are enabled (1) or disabled (0) using a single command.
- IEEE 802.3x link-level flow control is disabled when DCB’s PFC mode is enabled.

Q25: Is IEEE 802.1Qbb PFC Supported in SF and Switch Independent NIC Partitioning Modes?

- Yes. Priority Flow Control (PFC) is negotiated as part of DCB which is controlled on a per-port basis. The Data Center Bridging eXchange (DCBx) negotiated lossless traffic type will have PFC support. Lossy traffic types are not flow control protected beyond what is provided by the upper layers (such as with TCP/IP). The Marvell FastLinQ 578xx/3400/8400/34000 Series Adapters DCB support both lossless FCoE offload and lossless iSCSI offload TLV modes of operation.
- If PFC is not enabled during the DCBx exchange, IEEE 802.3x link-level flow controls settings will apply.
- DCBX Link Layer Data Protocol (LLDP) packets uses the iSCSI-Offload MAC address as it’s source address.

Q26: Does Switch Independent NIC Partitioning Support Jumbo Frames?

- Yes. Jumbo frames up to 9600 bytes are supported in both Windows and Linux OSs on all partitions, for both Ethernet and iSCSI offload protocol modes in both SF and Switch Independent NIC Partitioning modes.
- In SF and Switch Independent NIC Partitioning mode, VMware ESXi supports Jumbo frames for both iSCSI hardware offload and L2 Ethernet (which includes iSCSI non-offload initiator mode).
- Jumbo Frames are supported on L2 Ethernet in Oracle® Solaris and Citrix XenServer in both SF and Switch Independent NIC Partitioning modes.
- For Windows, VMware ESXi, and Linux OSs, FCoE protocol’s MTU is always fixed at 2500 bytes. The FCoE MTU size is independent of the Ethernet MTU size setting.

Q27: How Can I View the Current Switch Independent NIC Partitioning Settings in the Operating System?

- In Windows OS, you can see what protocols are enabled for a partition in both Windows Device Manager and Marvell FastLinQ management utilities. You can see a partition’s Relative Bandwidth Weight settings in Marvell FastLinQ management utilities. The Maximum Bandwidth value for an Ethernet protocol-enabled partition can be found in the Network Connections Status window or Marvell FastLinQ management utilities. If a partition has only the iSCSI or FCoE offload protocol enabled, you will only see the device and not its Maximum Bandwidth value in the Windows Device Manager, but everything is viewable and configurable in Marvell FastLinQ management utilities.
- In Linux OS, you can see what protocols are enabled for a partition in Marvell FastLinQ management utilities. You can see the partition’s Relative Bandwidth Weight and Maximum Bandwidth value for all protocol-enabled partitions in Marvell FastLinQ management utilities.
- In Solaris, XenServer, and VMware OSs, you can see the applicable Ethernet, iSCSI, and FCoE protocol-enabled partitions that are loaded and available for an application’s use. You can see a partition’s QoS settings in the Marvell FastLinQ management utilities.

Q28: What Does a Switch Detect When the Device is in Switch Independent NIC Partitioning Mode?

- On the Marvell FastLinQ 578xx/3400/8400/34000 Series Adapters, a switch can see additional MAC addresses (up to four L2 Ethernet MAC addresses), plus two additional iSCSI offload MAC addresses (depending upon the Switch Independent NIC Partitioning configuration), which is slightly more than what a switch would see from the Marvell FastLinQ 578xx/3400/8400/34000 Series Adapters, in SF mode.
- In SF mode, the same port will have one L2 Ethernet MAC address and one additional iSCSI offload MAC address. This is no change from the previous generation's Marvell FastLinQ devices which can also present up to two MAC addresses per port (L2 Ethernet and iSCSI hardware offload).
- A switch port should be set to TRUNK mode if VLANs and/or FCoE protocol mode is enabled on any SF port or Switch Independent Partition.

Q29: Is Switch Independent NIC Partitioning Mode Supported at 1GbE Negotiated Link Speeds on the Marvell FastLinQ 578Xx/3400/8400/34000 Series?

- Yes. The attached device and interface (SFP+ or 10GBase-T) would have to support 1GbE link speeds.
- The overall rate would be 1 Gbps instead of 10 Gbps.
- DCB and FCoE are not supported on lowered link speed 1Gbps connections.
- iSCSI Hardware offload is supported on lowered link speed 1Gbps connections on the applicable OSs.

Q30: How do I Enable DCB to Work with Switch Independent NIC Partitioning Mode?

- DCB is enabled in the various Marvell FastLinQ management utilities - Comprehensive Configuration Management (CCM) preboot utility, UEFI HII system BIOS device configuration pages, FW upgrade utility, QCC GUI, QCS CLI, QCC PowerKit, QCC vSphere GUI plug-in and QCC ESXcli plug-in.
- A reboot is required after enabling or disabling DCB but DCBX changes sent down from the connected DCB switch (via the DCBX LLDPs) are implemented without reboots. Changes in Priority Groups, ETS minimum BW and Losslessness (PFC) are automatically made on the fly.
- DCB works on both Switch Independent NIC Partitioning and Single Function mode ports.
- DCB is configurable on a per port granularity. This means all partitions on a port will have DCB enabled or disabled.
- DCB could be enabled on one port of a device and disabled on another port of the same device.
- DCB requires a Converged Enhanced Ethernet (CEE) or IEEE DCBX capable switch connected to the enabled port.
- DCB can be enabled specifically for the CEE version or for the IEEE version, or it can be set to dynamically detect which DCB version is being used by the link partner (switch).
- DCB is required for FCoE, iSCSI offload TLV, and Windows Server 2012 (or later) DCB-QoS service.

Q31: Does the Marvell FastLinQ 578Xx/3400/8400/34000 Series Support Windows RSS BaseProcessorNumber and MaxProcessorNumber in Windows Server 2012 (or Later)?

- Yes, all RSS capable Marvell FastLinQ devices support changes made to the RSS Indirection Table's BaseProcessorNumber and MaxProcessorNumber values using Windows 2012 (or later) PowerShell version 3.0 (or later).

Q32: How does Setting a Partitions Maximum Bandwidth (or Speed Setting) in VMWare's VClient Affect Those Settings In CCM/UEFI?

- The VMWare VClient speed settings of a VMNic are persistent, and show up in the configuration file /etc/vmware/esx.conf. These settings are stores in the tree as follows:

```
/net/pnic/child[0006]/duplex = "full"  
/net/pnic/child[0006]/mac = "XX:XX:XX:XX:XX:XX"  
/net/pnic/child[0006]/name = "vmnicX"  
/net/pnic/child[0006]/speed = "1000"
```

Note that this sets the speed of that VMNic (which could be a Switch Independent Partition) to 1G. Because this is persistent and applied with every ESXi boot, these speed settings will always override that specific partition's Maximum Bandwidth settings from the pre-boot CCM/UEFI HII system BIOS device configuration page utilities.

Q33: How do you Enable the Windows Server 2012 R2 (or Later) NVGRE/VXLAN Packet Task Offloads (TX LSO, TX CO, RSS and VMQ)?

- To setup Windows Hyper-V NVGRE or VXLAN, see the Microsoft multitenant SDN configuration documentation at <https://docs.microsoft.com/en-us/windows-server/networking/sdn/technologies/hyper-v-network-virtualization/hyper-v-net-work-virtualization> and <https://docs.microsoft.com/en-us/windows-server/networking/sdn/manage/manage-tenant-virtual-networks>.

Q34: How do you Enable the VMWare ESXI 5.5 (or Later) VXLAN Packet Task Offloads (UDP/ TCP RSS, TX TSO and TX CO)?

- To enable VxLAN with UDP/TCP RSS acceleration using the legacy bnx2x driver:
`esxcfg-module -s RSS=N bnx2x`
where "N" can be 2-4. Setting this to 1 - 4 will set a fixed number of RSS queues to be used by VXLAN/GENEVE tunneled traffic or the HOST. Setting it to 1 will essentially disable RSS queueing.
- To enable VxLAN/GENEVE with UDP/TCP RSS acceleration using the native qfle3 driver:
`esxcfg-module -s RSS=N qfle3`
where "N" can be 0-4. Setting this to 0 will allow the software to dynamically control the number of RSS queues used by VXLAN tunneled traffic or the HOST (default setting). Setting this to 1 - 4 will set a fixed number of queues. Setting it to 1 will essentially disable RSS queueing.
- To enable VxLAN with TX TSO and TX CO using the legacy bnx2x driver:
`esxcfg-module -s enable_vxlan_ofld=1 bnx2x`
Setting this to "0" will disable these VxLAN offloads.
- To enable VxLAN/GENEVE with TX TSO and TX CO using the native qfle3 driver: `esxcfg-module -s offload_flags=N qfle3`
where "N" is the Bit-Mask value derived from 1 (TX CSO), 2 (TSO), 4 (for VXLAN), and 8 (for GENEVE). Default is all enabled (1111b or 0xF or 15d).
- There are other qfle3 VXLAN/GENEVE module parameter (such as "DRRS", "rss_engine_nr", and "enable_vxlan_filters"), but these should normally be left at their default settings.
- The current driver module settings can be queried using the command - the driver name can be Legacy (bnx2, bnx2x, bnx2fc, bnx2i, cnic) or Native (qflge, qfle3, qfle3f, qfle3i, qcnic):
`esxcli system module parameters
list -m <driver-Name>
esxcfg-module -i <driver-Name>`

Q35: How do you Enable FCoE-Offload N-Port ID Virtualization (NPIV) on Windows, VMWare and Linux?

- NPIV allows you to assign more than one World Wide Port Name (WWPN) to a single FCoE-Offload port/partition - thus allowing multiple separate FCoE connections.
- In Windows, you can use QCC GUI by right clicking on the FCoE instance and setup the new WWPNs for that adapter instance. You can also use the PowerShell commands “Add-VMFibreChannelHba”, “Get-VMFibreChannelHba”, “Remove-VMFibreChannelHba” and “Set-VMFibreChannelHba” - see <http://technet.microsoft.com/en-us/library/hh848550.aspx>, <http://technet.microsoft.com/en-us/library/hh848588.aspx>, <http://technet.microsoft.com/en-us/library/hh848511.aspx> and <http://technet.microsoft.com/en-us/library/hh848485.aspx> for more information. The Marvell FastLinQ management utilities QCS CLI and QCC PowerKit support similar operations.
- In VMware ESXi 6.5 (or later) using the native qfle3f driver, you can use the QCC vCenter/vSphere GUI plug-in and QCC ESXcli plug-in on the FCoE instance to setup new WWPNs (VNPorts) for that adapter instance - for more information see the Marvell/QLogic Fibre Channel and Converged Network Adapters for VMware ESXi 5.5 and 6.x User’s Guide, QLogic QCC ESXcli plug-in User’s Guides, and the VMware ESXcli documentation. The VMware legacy bnx2fc driver does not support NPIV.
- Linux, you can use the vPort_Create command.
- You can verify the new WWPN are created using a show FLOGI database type command on your connected FCF switch.
- Windows Server 2012 (or later) Hyper-V all support up to four virtual Fibre Channel (vFC) instances in Windows OS VMs. You don’t enable NPIV to use vFC - see <https://technet.microsoft.com/en-us/library/dn551169.aspx>, and <http://technet.microsoft.com/en-us/library/hh831413.aspx> for more information on Win2012 (or later) Hyper-V’s vFC feature.

Q36: Where are the Linux, VMWare and Windows Drivers and Other Related Collateral Located for the Marvell FastLinQ Series Adapters?

- Go to <https://www.marvell.com/support/doc-library>




Table 2. SF Versus Switch Independent NIC Partitioning Mode

Feature	Single Function (SF)	Switch Independent NIC Partitioning
Ethernet Adapters	1 per port	Up to 4 per port; see FAQ entry. See “Q5: How many Ethernet protocol device partitions are allowed?”
iSCSI Hardware Offload Adapters	1 per port	Up to 2 per port; see FAQ entry. See “Q6: How many iSCSI offload protocol HBA device partitions are allowed?”
FCoE Hardware Offload Adapters	1 per port	Up to 1 per port; see FAQ entry. See “Q7: How many FCoE offload protocol HBA device partitions are allowed?”
Receive Segment Coalescing (RSC)	Yes	Yes
Linux/Windows/VMware Virtual Extensible LAN (VXLAN) Packet Task Offloads	Yes	Yes
Windows Network Virtualization using Generic Routing Encapsulation (NVGRE) Packet Task Offloads	Yes	Yes
Linux/VMware GENEVE Packet Task Offloads	Yes	Yes
TCP Segmentation Offload (TSO)	Yes	Yes
Large Send Offload (LSO) and Large Receive Offload (LRO)	Yes	Yes
Giant Send Offload (GSO)	Yes	Yes
IPv4/IPv6 TCP-UDP-IP Checksum Offload (CO)	Yes	Yes
IP Address Multihoming	Yes	Yes
PXE/iSCSI/FCoE Boot	Yes	Yes – first partition of each port; see FAQ entry. See “Q16: How are PXE/iSCSI/FCoE remote boots affected by Switch Independent NIC Partitioning mode?”
Wake-on-LAN (WOL)	Yes	Yes – see FAQ entry. See “Q17: How is Wake On LAN (WOL) Out Of Band (OOB) affected by Switch Independent NIC Partitioning mode?”
PAUSE Flow Control	Yes	Yes – see FAQ entry. See “Q21: Is IEEE 802.3x link-level flow control supported in SF and Switch Independent NIC Partitioning mode?”
Jumbo Frames (up to 9600 Bytes)	Yes	Yes
iSCSI Offload Connections	128 per port	128 per partition – for up to two iSCSI Offload Protocol enabled partitions
RX/TX L2 Ethernet Queues (RSS/NetQueue/VMQ)	16 per port	16 per partition
RX/TX L2 Ethernet Queues (Windows dVMQ)	61 per port ¹	61 per partition ¹
MSI	8	8
MSI-X	1024	1024
OS specific Multipath I/O (MPIO)	Yes	Yes
Teaming – Switch Independent	Yes	Yes – see FAQ entry. See “Q20: How are Windows BASP Teaming, Windows Server 2012 in-OS NIC Teaming, and Linux/VMware Bonding affected by Switch Independent NIC Partitioning mode?”
Teaming – Link Aggregation – Switch Dependent	Yes	No – see FAQ entry. See “Q20: How are Windows BASP Teaming, Windows Server 2012 in-OS NIC Teaming, and Linux/VMware Bonding affected by Switch Independent NIC Partitioning mode?”
Teaming – Generic Trunking – Switch Dependent	Yes	No – see FAQ entry. See “Q20: How are Windows BASP Teaming, Windows Server 2012 in-OS NIC Teaming, and Linux/VMware Bonding affected by Switch Independent NIC Partitioning mode?”
VLAN (insertion/stripping/filtering)	Yes	Yes

1. The current number of VMQs available for use (on that Ethernet physical function) by the OS can be seen in the output of the Microsoft Windows PowerShell® “Get-NetAdapterVmq” command in the Number Of Receive Queues field.

Table 3. FastLinQ 578xx/QLE344x/QLE844x/QL34xxx Driver Names

Feature	Linux	Windows	VMware
Common Core	In L2 Driver	eVBD	In L2 Driver
L2 Ethernet Networking	bnx2x	bxND	bnx2x* or QFLe3**
cNIC (Offload Libraries)	cNIC		cNIC*/cNIC_register* or QcNIC**
iSCSI-Offload	bnx2i	bxOIS	bnx2i* or QFLe3i**
FCoE-Offload	bnx2fc	bxFCoE	bnx2fc* or QFLe3f**

* Legacy Drivers for VMware ESXi 6.5 and earlier

** Native Drivers for VMware ESXi 6.5 and later