

Benefits of Offloading I/O Processing to the Adapter

FCoE and iSCSI Protocol Offload Delivers Enterprise-class Performance, Reliability, and Scalability



Key Benefits

Full Offload Capability

- Lowers CPU utilization for I/O processing
- Increases application performance
- Reduce Latency with Universal RDMA
- Allows for more virtual machines (VMs) per server

Advanced Virtualization Techniques

- Enables hardware consolidation
- Reduces I/O emulation overhead
- Dedicates bandwidth quality of service (QoS) for improved scalability
- Lowers adapter, cabling, and management costs

Simplified Management

- Concurrent Ethernet, FCoE, and iSCSI protocols simplify server provisioning
- Minimizes deployment disruptions
- Agnostic to network switches
- Proven and broadest interoperability with OSes, switches, storage, ISV, and IHV fail-over solutions
- Reduces operating overhead—single SKU across OSes/protocols

Hewlett Packard Enterprise (HPE) and Marvell have partnered to deliver to their customers the highest performing, most reliable, and scalable I/O connectivity for the HPE ProLiant server portfolio. HPE and Marvell have a decades long track record of delivering connectivity solutions to simplify HPE ProLiant server connectivity to networks and storage.

Executive Summary

HPE® and Marvell provide converged networking for the HPE ProLiant® rack and tower, HPE Apollo and HPE Synergy platforms. HPE customers with HPE servers can leverage HPE-branded 10Gb, 25Gb and 50Gb Converged Network Adapters (CNA), based on Marvell FastLinQ technology, to simplify server network connectivity for these server platforms. This technology brief will help you make an informed selection regarding these Marvell FastLinQ adapters for HPE.

Benefits Of Hardware Offload

Server virtualization driven consolidation accelerates the move to 10/20/25/50GbE converged networking technology in data center networks. Virtualized environments require higher performance and greater bandwidth. IT administrators can extend the benefits of virtualization by consolidating servers, lowering CPU utilization, accommodating multiple protocols [Ethernet, iSCSI, and Fibre Channel over Ethernet (FCoE)], simplifying bandwidth control (QoS), and reducing cable management complexity as they deploy 10/20/25/50GbE technology. With support for iSCSI and FCoE storage protocols, storage and network traffic can share a single connection, reducing the number of server connections and switch ports consumed. Performance, scalability, and reliability should be understood when specifying converged networking connectivity for data center.

Performance

Although it drives consolidation, virtualization imposes additional overhead on the CPU, network and I/O. With limited I/O bandwidth and CPU resources, it is not logical to further burden the CPU with processing I/O storage requests by using software FCoE or iSCSI initiators. In addition, mobility of virtual servers, load balancing, and failover require high throughput, especially when shared by multiple VMs. For enterprise-class applications, software FCoE and iSCSI solutions not only fail to meet enterprise requirements but can also be detrimental to data center virtualization goals. Deploying an I/O adapter with full hardware offload capability allows you to run more applications per server and extend the useful life of your infrastructure. Initial capital cost of software FCoE or iSCSI may be low, but it is important to consider the operational and intangible costs of a software initiator.

Latency is also a key parameter that impacts performance. To reduce latency, select CNAs from HPE support Remote Direct Memory Access (RDMA). This includes the HPE Ethernet 10/25GbE FLR-SFP28 QL41401-A2G CNA, HPE Synergy 4820C/6820C, HPE CN1200R BASE-T CNA and HPE CN1300R 10/25Gb CNA. All of these adapters are based on Marvell FastLinQ 41000 series technology and support Universal RDMA. This is the ability to concurrently run either iWARP or RDMA over Converged Ethernet (RoCE) protocols. RDMA bypasses the kernel and allows I/O transactions directly between the adapter and memory, reducing transaction time and lowering latency. By supporting both iWARP and RoCE RDMA protocols, these CNAs provide more flexibility for customers in terms of protocol choice and system design.

To further enhance performance, all HPE 10/20/25/50GbE adapters based on technology from Marvell support small packed acceleration with Data Plane Developer Kit (DPDK) libraries. DPDK is supported in the Linux OS to speed processing of small packet IP traffic within the adapter by bypassing the Linux kernel..

Reliability

Processing mechanisms for storage traffic (iSCSI or FCoE) must maintain data integrity while sending storage data over Ethernet. Data integrity checking is a compute-intensive process that is either performed by the CPU in solutions using a software initiator or by the offload engine of a CNA. Full hardware offload maintains the highest level of data integrity. Increasing traffic across an Ethernet network produce dropped and out-of-order data frames. Using software initiators, recovering from both these issues can be a significant burden on the CPU and may cripple overall performance on a 10/20/25GbE network. The HPE CN1200R-T/CN1300R, HPE Ethernet 10/25Gb FLR-SFP28 QL41401-A2G CNA, HPE Synergy 4820C/6820C CNAs, and HPE FlexFabric 10GbE 53X Series Adapters, powered by the Marvell offload engine, reassemble out-of-order frames and complete the process of resending those that are dropped while maintaining low CPU utilization, ensuring the highest levels of performance and reliability. Software FCoE or iSCSI initiators lack maturity for enterprise applications and it will take years of qualification testing to be ready for reliable enterprise data center deployments.

Enterprise data centers that run mission-critical applications require extensive, heterogeneous, and scalable deployments with high CPU efficiency and reliability, availability, and serviceability (RAS) built into the design of the data center functions. The Marvell hardware offload technology and solutions provide key benefits to meet such demands.

Hardware-based boot from SAN offers several benefits for data center deployments. Boot images for the servers can be stored remotely with centralized management, leading to higher density and utilization of storage in each server, improved cooling with reduced storage requirement, and efficient management of host image deployments.

This means lower operational costs while improving RAS by avoiding single point of failure on the local disk in the server. A stable, reliable, and extensively field-proven boot from SAN hardware offload solution for the critical data center operations provides key advantages for the enterprise that a software FCoE solution may not adequately meet at this time. Multipath I/O provides the critical high-availability paths for SAN deployments in the data center. It enables the establishment of multiple paths between the I/O devices with redundant paths that can be leveraged when an I/O path fails. Multipath I/O allows enterprisecritical applications to function, providing access to the critical

data required for maintaining mission-critical functions. Increasing the number of virtualized deployments intensifies the importance of the multipath solution in the data center. When a path to a virtualized server fails, it is no longer a singular OS that is impacted. The impact is exponential due to the high density of the VMs that are running with scaled-out applications. The Marvell full hardware offload solution provides the same level of RAS for multipath I/O and fail-over solutions out-of-box, as seen with the boot from SAN solutions.

Scalability

Software FCoE or iSCSI have lower scalability capacity. One adapter running a software initiator could easily consume up to one-third of the CPU processors' capabilities. Adding multiple adapters into a single system that relies on the CPU to perform their multiple operations only compounds the problem, as does additional VMs. Using a software initiator on a NIC requires that every incoming TCP/IP, FCoE, and iSCSI packet be sent over the PCI® bus in the server. Sending packets back and forth increases the PCI bus' busy state, and can cause bottlenecks with other hardware on the PCI bus. Offload engines are a better alternative to software initiators in enterprise servers because organizations can address emerging and future scalability requirements within enterprise data centers. As enterprise application customers strive to achieve density and resource utilization objectives, high I/O performance will emerge as a requirement. Virtualization requires CPU processing capability to efficiently scale VMs. Furthermore, adding a virtualization layer can add overhead, which will degrade I/O performance.



HPE-Branded Marvell FastlinQ Network Adapter Family

Enterprise Requirements for I/O Adapters		Converged Network/FlexFabric Adapter	Software Initiator
PERFORMANCE	Support for I/O-Intense Applications	YES	NO
	Efficient CPU Utilization	YES	NO
	DPDK small packet acceleration	YES	NO
	RDMA - iWARP, RoCE, RoCEv2	YES	NO
RELIABILITY	Data Integrity Assurance	YES	NO
	Comprehensive Fail-over Support	YES	NO
	Enterprise Reliability	YES	NO
	Investment Protection	YES	YES
SCALABILITY	Network Partitioning (NPAR)	YES	NO
	Overlay Offloads (NVGRE/VXLAN)	YES	NO
	Scalability within Virtual Operating Environments	YES	NO
	IOPS Scalability	YES	YES
	Concurrent I/O Support for Consolidation	YES	NO
	Broad OS Support	YES	NO
	Proven Interoperability	YES	NO
	Boot from SAN Across Configurations and OSes	YES	YES

Network Partitioning (NPAR)

HPE Network Adapters from Marvell are optimized to work in Virtual Server environments. The HPE FlexFabric 10GbE 53X Series and the HPE Synergy 4820C/6820C adapters support HPE Virtual Connect virtualization to improve scalability, increase flexibility, and reduce management. The same adapters for HPE ProLiant, Apollo and HPE Synergy servers support NPAR to provide adapter-based virtualization, reducing the number of physical adapters required in a virtual server environment. All these adapters also support overlay offloading via Network Virtualization using Generic Routing Encapsulation (NVGRE) and Virtual Extensible LAN (VXLAN) to simplify network virtualization and security in highly virtualized environments.






In all HPE CNAs and FlexFabric Adapters from Marvell, the adapter processes the FC driver stack. Offloading FC processing to the adapter enables the CPU resources to support more VMs. This provides greater cost savings and return on investment (ROI) for virtualized environments.




Along with performance, reliability, and scalability benefits, Marvell FastLinQ Converged Network Adapters for HPE offer additional operational efficiencies such as: boot from SAN, multipath I/O, the elimination of PCIe bus bottlenecks, a mature, field-hardened driver, and the largest interoperability base in the industry.

HPE and Marvell Deliver Joint I/O Solution

The HPE FlexFabric and HPE Synergy CNA Series Adapters from Marvell offer a full-featured, high-performance, and scalable solution that includes convergence and virtualization. CNAs provide the flexibility of built-in support for iSCSI and FCoE right off the motherboard without requiring any additional licensing costs. Support for concurrent NIC, iSCSI, and FCoE functions deliver exceptional value and future-proof data center investments.

Flexibility, converged functions, virtualized I/O, and effective system management all replace a single-function I/O for HPE ProLiant DL and ML servers. This powerful new approach to server I/O is a compelling reason to choose the HPE FlexFabric and HPE Synergy Adapters from HPE and Marvell.

Adapter Model Number	Part Number	Details	Adapter Photo	
HPE PROLIANT DL, ML AND SL SERVERS	HPE FlexFabric 10Gb FLR-SFP+ 57810S	700751-B21	10GbE FlexFabric Adapter Form Factor: PCIe 2.0 x8, Flexible LOM Virtualization Support: NPAR, SR-IOV Tunnel Offload: Network Virtualization using NVGRE, VXLAN Connection Type: SFP+ Small Packet Acceleration: DPDK	
	HPE FlexFabric 10Gb 2-port FLR-T 57810S	700759-B21	10GbE FlexFabric Adapter Form Factor: PCIe 2.0 x8, Flexible LOM Virtualization Support: NPAR, SR-IOV Tunnel Offload: NVGRE, VXLAN Connection Type: 10GBASE-T Small Packet Acceleration: DPDK	
	HPE FlexFabric 10Gb 4-port FLR-T 57840S	764302-B21	4-port 10GbE FlexFabric Adapter Form Factor: PCIe 3.0 x8, Flexible LOM Virtualization Support: SR-IOV Tunnel Offload: NVGRE, VXLAN Connection Type: 10GBASE-T Small Packet Acceleration: DPDK	
	HPE Ethernet 10/25GbE FLR-SFP28 QL41401-A2G CNA	867334-B21	10/25GbE Converged Network Adapter Form Factor: PCIe 3.0 x8 Flexible LOM Virtualization Support: SR-IOV Tunnel Offload: NVGRE, VXLAN, GENEVE Universal RDMA: RoCE or iWARP Connection Type: SFP28 Small Packet Acceleration: DPDK	
	HPE CN1300R CNA	Q0F09A	10/25GbE FlexFabric Adapter Form Factor: PCIe 3.0 x8 Virtualization Support: SR-IOV Tunnel Offload: NVGRE, VXLAN, Universal RDMA: RoCE or iWARP GENEVE Connection Type: SFP+ Small Packet Acceleration: DPDK	

	Adapter Model Number	Part Number	Details	Adapter Photo
HPE PROLIANT DL, ML AND SL SERVERS	HPE CN1200R-T CNA	Q0F26A	10GbE Converged Network Adapter Form Factor: PCIe 3.0 x8 Virtualization Support: SR-IOV Tunnel Offload: NVGRE, VXLAN, GENEVE Universal RDMA - RoCE or iWARP Connection Type: 10GBASE-T Small Packet Acceleration: DPDK	
HPE SYNERGY SERVERS	HPE Synergy 4820C 10/20/25Gb CNA	876449-B21	10/20/25GbE Converged Network Adapter Form Factor: PCIe 3.0 x8 Mezzanine Tunnel Offload: NVGRE, VXLAN, GENEVE Virtualization Support: SR-IOV NPAR (8 physical functions/port) Universal RDMA - RoCE or iWARP Small Packet Acceleration: DPDK	
	HPE Synergy 6820C 25/50Gb CNA	P02054-B21	25/50GbE Converged Network Adapter Form Factor: PCIe 3.0 x8 Mezzanine Tunnel Offload: NVGRE, VXLAN, GENEVE Virtualization Support: SR-IOV, NPAR (8 physical functions/port) Universal RDMA - RoCE or iWARP Small Packet Acceleration: DPDK	



To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

Copyright © 2020 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.