ESG Lab Review

Accelerating Network Virtualization Overlays with QLogic Intelligent Ethernet Adapters

Date: July 2016 Author: Jack Poller, Senior Lab Analyst

Executive Summary

Enterprise IT organizations are transitioning to the modern data center to address performance and service provision challenges, leveraging lessons learned from hyper-scale and public cloud deployments. More specifically, they are implementing highly virtualized, multi-tenant infrastructures that span geographic locations using overlay network tunneling technologies. VMware, the top hypervisor for these deployments, uses VXLAN encapsulation for overlay network tunneling. Benefits of VXLAN include:

- Scalability supporting up to 16 million virtual LAN IDs
- Secured shared infrastructure enabling separation and isolation for multi-tenant deployments
- Improved resource utilization implementing equal cost multi-path routing in L3 doubles the amount of network bandwidth available compared to the L2 spanning tree protocol

With IT spending budgets limited and data growing at rapid rates, it's important for organizations to address the performance impacts of network virtualization overlays introduced in complex virtualization environments. QLogic looks to alleviate these concerns by offloading tunneling (VXLAN) encapsulation from the host processor with QLogic FastLinQ 45000 Series 25Gb Ethernet (25GbE) Intelligent Ethernet Adapters, which frees up the host. Similar performance improvements are enabled with the QLogic FastLinQ 3400 Series 10GbE Intelligent Ethernet Adapters which supports tunneling while maintaining stateless offloads.

ESG Lab tested the performance benefits of offloading tunneling for network virtualization overlays using IxChariot, an application performance and network assessment tool from Ixia. Performing tunneling with stateless offloads enabled the FastLinQ 3400 Series 10GbE Adapter to deliver an average 19% improvement in throughput. Stateless offloads in conjunction with tunneling encapsulation and decapsulation enabled the host to operate 6% more efficiently. Likewise, the FastLinQ 45000 Series 25GbE Adapter delivered an average 18% increase in throughput. More impressively, by offloading tunneling encapsulation from the host, applications could transfer 55% more data per unit of host CPU work. These gains in efficiency enable QLogic to deliver on the IT mission of doing more with less.

QLogic FastLinQ 45000 Series 25Gb Ethernet



18% Tunneling Throughput Increase

55% Tunneling Throughput Increase per Unit of Host CPU Work

The Challenges

In a recent ESG survey, respondents were asked to identify their top IT priorities for 2016. As they have been since 2010, increased use of server virtualization (cited by 20% of respondents for 2016), increased use of desktop virtualization (20%), improved backup and recovery (20%), and business continuity and disaster recovery (18%) are all in the top ten most frequently cited IT priorities.¹ These priorities place significant performance demands on the network infrastructure. As a result, more than one quarter of those surveyed (27%) cited network performance, and 21% cited difficulties and/or delays

¹ Source: ESG Research Report, <u>2016 IT Spending Intentions Survey</u>, February 2016.

This ESG Lab Review was commissioned by QLogic and is distributed under license from ESG.

^{© 2016} by The Enterprise Strategy Group, Inc. All Rights Reserved.

in provisioning network services for new or updated applications and/or VM mobility as their biggest network challenges (see Figure 1).²

Network performance and VM mobility become more important as organizations transition to the modern data center, incorporating lessons learned from hyper-scale and public cloud environments, including highly virtualized multi-tenant infrastructures that span geographic locations. With more and more VMs and more and more tenants, network infrastructures quickly scale beyond the limits of traditional network overlay technologies such as VLANs.

Figure 1. Top-Ten Biggest Networking Challenges

In your opinion, what are the biggest challenges facing your organization's networking team? (Percent of respondents, N=306, five responses accepted)



Source: Enterprise Strategy Group, 2016

The Solution: Overlay Tunneling with QLogic Intelligent Ethernet Adapters

VMs can be migrated among any servers on the same subnet, and VLANs are used to provide separation and isolation between different groups of VMs and servers. However, there are only 4,096 possible VLAN IDs, which limits the utility of VLANs to just the smallest environments.

Tunneling—the transmission of packets that are encapsulated within an outer-header—also provides separation and isolation. As with VLANs, tunneling can be used to create a virtual L2 network. NVGRE encapsulates an L2 packet inside an IP packet, and is used by Microsoft Windows Server 2012 R2 Hyper-V. VXLAN encapsulates an L2 packet inside a UDP packet, and is used by VMware as well as Linux-based hypervisors such as KVM and Xen.

The benefits of using tunneling instead of VLANs include:

- Scalability—Tunneling can support up to 16 million separate IDs compared with only 4,069 for VLANs.
- Shared infrastructure—Different tenants may use the same MAC addresses and VLANs while maintaining separation and isolation with tunneling.

² Source: ESG Research Report, *Data Center Networking Trends*, February 2016.

- **Reduced impact on network nodes**—Network nodes no longer need to be aware of each VM's MAC address, shrinking the size of MAC/IP tables, which can be a critical scaling factor in modern leaf-spine network architectures.
- Simplified networking management and migration—Tunneling decouples logical addresses (used in the inner headers by the VMs) from physical addresses (used in the outer headers by network nodes for switching and routing), providing location-independent addressing.
- Improved resource utilization—VLAN uses spanning tree protocol (STP) to avoid network loops, which can block up to half of the available network bandwidth. VXLAN uses L3 equal cost multi-path routing (ECMP) to use all available network paths.

In a fully virtualized network, traditional network adapters cannot use their stateless offload features due to the added encapsulation of NVGRE or VXLAN. As a result, the hypervisor processes the tunneling traffic, which can exhaust system resources and can significantly affect the performance of the entire system.

The QLogic FastLinQ 3400 Series 10GbE Adapters support tunneling with stateless offload and the FastLinQ 45000 Series 25GbE Adapters provide full tunneling encapsulation offload, accelerating performance and reducing CPU load. Performance gains come from:

- Transmit large segment offload (LSO)/TCP segment offload (TSO)—The Ethernet adapter segments large data blocks into network packets, encapsulates the packets, and transmits the packets without the participation of the host CPU.
- Receive side scaling (RSS)/transmit side scaling (TSS)—The Ethernet adapter hashes the header of the encapsulated data and routes the packet to one of many queues, allowing traffic to spread across multiple CPU cores for load balancing and scaling.
- Large receive offload (LRO)—The Ethernet adapter receives and de-encapsulates multiple packets, reassembling data into large blocks, without the participation of the host CPU (25Gbps Ethernet adapters only).
- **Tunneling offloads**—The Ethernet adapter encapsulates and de-encapsulates the packets for supported tunneling protocols including VXLAN, NVGRE, and GENEVE.

Table 1 provides details for the different offload capabilities supported by the FastLinQ 3400 Series 10GbE Adapters and FastLinQ 45000 Series 25GbE Adapters.

Table 1. QLogic FastLinQ Ethernet Adapter Offload Capabilities

Offload Capabilities	FastLinQ 3400 Series 10GbE Adapters	FastLinQ 45000 Series 25GbE Adapters
Transmit large segment offload (LSO)/ TCP segment offload (TSO)	\checkmark	\checkmark
Receive side scaling (RSS)/ Transmit side scaling (TSS)	\checkmark	\checkmark
Large receive offload (LRO)		\checkmark
Tunneling encapsulation and de- encapsulation		\checkmark

Figure 2. QLogic FastLinQ 3400 Series 10GbE and FastLinQ 45000 Series 25GbE Intelligent Ethernet Adapters³

QLogic FastLinQ 3400 Series 10GbE Intelligent Adapters and QLogic FastLinQ 8400 Series 10GbE Converged Network Adapters QLogic FastLinQ QL45000 Series 25GbE Intelligent Ethernet Adapters

Source: Enterprise Strategy Group, 2016

QLogic 10GbE and 25GbE tunneling offload adapters are available in PCIe network interface card (NIC), rack network daughter card (NDC), blade NDC, and blade mezzanine form factors with one (10GbE adapters only) or two ports to meet most modern data center needs. Using these adapters in the modern fully virtualized data center brings the following benefits:

- Accelerated server performance—Full line-rate 10Gbps and 25Gbps throughput across all ports with full hardware offload for iSCSI and Fibre Channel over Ethernet (FCoE) storage traffic, along with tunneling offload (VxLAN, NVGRE and GENEVE) for enhanced throughput and efficiency in the modern data center.⁴
- Increased virtualization CPU scaling—Full support for network virtualization technologies including switch independent NIC partitioning (NPAR), single root I/O virtualization (SR-IOV), VMware NetQueue, and Microsoft VMQ.
- Enhanced deployment and management—Unifies adapter and storage management with integrated lifecycle controller management and QLogic QConvergeConsole(QCC) / QLogic Control Suite (QCS).

Validating Tunneling Acceleration

ESG Lab performed hands-on testing of FastLinQ 3400 Series 10GbE Adapters and FastLinQ 45000 Series 25GbE Adapters, focusing on the performance gains that result from offloading tunneling. The test bed, shown in **Figure 3**, consisted of two VMware servers, each installed with its own QLogic Ethernet adapter. For 10Gbps testing, we used the QLogic QLE3442-SR Dual-port 10GbE Adapter, and for 25Gbps Ethernet testing, we used the QLogic QL45212 Dual-port 25GbE Adapter. We connected one high-speed port (10Gbps or 25Gbps) on each server to a 10Gbps/25Gbps Ethernet switch.

The test bed environment also included an additional server and a workstation for management and test control. The environment was managed through a separate 1Gbps Ethernet management network with its own 1Gbps Ethernet switch.

³ For a complete list of QLogic Intelligent Ethernet adapters, please refer to

http://www.qlogic.com/Products/adapters/Pages/IntelligentEthernetAdapters.aspx

⁴ iSCSI and FCoE offload is available on select QLogic Ethernet adapters. Please check specifications for details.



QLogic FastLinQ 3400 10GbE Intelligent Ethernet Adapters

ESG Lab began by testing the QLogic **QLE3442-SR Dual-port 10GbE Adapter**. Network traffic was simulated using IxChariot, an application performance and network assessment tool from Ixia. IxChariot simulates real-world applications, enabling users to measure device, system, and network performance under realistic conditions.

Each VMware server was configured with four VMs. Each of the VMs ran two threads of IxChariot. Thus, for unidirectional tests, there were a total of eight threads moving traffic. For bidirectional simultaneous transmit and receive tests, there were a total of sixteen threads moving traffic—eight threads for each direction.

We first ran unidirectional tests, measuring the maximum amount of data transferred in one direction. Each test used a different data I/O size, starting at 1 kilobyte (kB) and progressing through 2kB, 4kB, 8kB, 16kB, 32kB, and finally ending at 64kB. Either the hypervisor or the offload engine was responsible for segmenting the large I/O blocks into network packets, encapsulating the packets, and transmitting the packets to their destination. We ran each set of tests twice, first without stateless offload, and then a second time, with stateless offloaded from the host CPU to the QLE3442.

As shown in Figure 4, at every data block size, the QLE3442 delivered more throughput when tunneling was offloaded to the Ethernet adapter. The most dramatic results were delivered with 32kB and 64kB I/O blocks, demonstrating the value of transmit large segment offload (LSO) and TCP segment offload (TSO).

Figure 4. QLogic QLE3442 10Gbps Simultaneous Transmit and Receive Throughput



QLogic QLE3442 10Gbps Simultaneous Transmit & Receive Throughput

Source: Enterprise Strategy Group, 2016

The relative performance gains for each I/O size are shown in Figure 5. The smallest gains were realized with 1kB data blocks (7%), while the largest gains were realized with 64kB data blocks (32%). This is to be expected, as more work is required to segment a large block into packets. On average, offloading tunneling resulted in a 19% improvement in the amount of data transferred.

Figure 5. QLogic QLE3442 10Gbps Simultaneous Transmit and Receive Offload Throughput Improvement



QLogic QLE3442 10Gbps Simultaneous Transmit & Receive Offload Throughput Improvement

Source: Enterprise Strategy Group, 2016

In addition to increasing the amount of data that can be transferred, efficiencies are increased with offload enabled by reducing the overall load on the host processor. We computed the amount of data that can be transferred for each unit of CPU work, measured as Mbps/CPU %. As shown in Figure 6, the QLE3442 was able to transfer more data per unit of CPU work for almost every I/O block size.

QLogic QLE3442 10Gbps

Figure 6. QLogic QLE3442 10Gbps Simultaneous Transmit and Receive Offload Efficiency



Source: Enterprise Strategy Group, 2016

Next, we calculated the percentage improvement in efficiency achieved when tunneling with offload is enabled. The largest gains were obtained at 8kB I/O block size. On average, the QLE3442 delivered a 6% improvement in efficiency, as shown in Figure 7.

Figure 7. QLogic QLE3442 10Gbps Simultaneous Transmit and Receive Offload Efficiency Improvement



Source: Enterprise Strategy Group, 2016

QLogic FastLinQ 45000 Series 25GbE Intelligent Ethernet Adapters

Next, ESG Lab repeated the entire test sequence using **QLogic QL45212 25GbE Dual-port Adapters**. The amount of data transferred is shown in Figure 8. As with the 10GbE adapters, the QL45212 delivered improved throughput with tunneling offload. When the host CPU encapsulated the data, the host was able to sustain about 38Gbps. When offloading encapsulation to the NIC, the system was able to sustain between 43 and 46Gbps.

Figure 8. QLogic FastLinQ QL45212 25Gbps Simultaneous Transmit and Receive Throughput



QLogic FastLinQ QL45212 25Gbps Simultaneous Transmit & Receive Throughput

Source: Enterprise Strategy Group, 2016

We computed the percentage increase in throughput, as shown in Figure 9. On average, the QL45212 tunneling offload delivered an 18% increase in throughput.

Figure 9. QLogic FastLinQ QL45212 25Gbps Simultaneous Transmit and Receive Offload Throughput Improvement



Next, we calculated the reduction in host CPU utilization with tunneling offload, as shown in Figure 10. On average, the QL45212 delivered a 24% reduction in host CPU utilization.

Figure 10. QLogic FastLinQ QL45212 25Gbps Simultaneous Transmit and Receive Offload CPU Utilization Improvement



Source: Enterprise Strategy Group, 2016

Finally, we reviewed the efficiency gains obtained through tunneling offload. As expected, the maximum gains were obtained at large block sizes, and on average, the QL45212 delivered a 55% improvement in efficiency, as measured by Gbps/CPU utilization.

Figure 11. QLogic FastLinQ QL45212 25Gbps Simultaneous Transmit and Receive Efficiency



Why This Matters

Operating more efficiently is an ongoing mission for IT in any organization. In a perfect world, IT would spend significantly less, achieve higher levels of performance, and use fewer resources, while continuing to gain a competitive advantage. 10Gbps and 25Gbps Ethernet enables organizations to operate at maximum throughput, but using this resource has led to host processor bottlenecks. The easy response is to just buy more, but that is costly and not a sustainable business model.

ESG Lab validated that the QLogic FastLinQ 3400 Series 10GbE and FastLinQ 45000 Series 25GbE Intelligent Ethernet Adapters accelerate throughput and reduce host CPU utilization by offloading tunneling from the host CPU to the network adapter. With tunneling and stateless offload, the FastLinQ 3400 Series 10GbE Adapter delivered an average 19% improvement in throughput. The offloading of work enabled the host to operate 6% more efficiently. Likewise, the FastLinQ 45000 Series 25GbE Adapter delivered an average 18% increase in throughput. More impressively, by offloading tunneling from the host, applications could transfer 55% more data per unit of host CPU work. These gains in efficiency enable QLogic to deliver on the IT mission of doing more with less.

The Bigger Truth

IT organizations are taking to heart the lessons of the cloud- and hyper-scale data centers in driving efficiencies. As a result, surveyed organizations ranked data center modernization third in the list of CIO whiteboard initiatives, right behind cybersecurity and data analytics.⁵ However, the modern data center with a fully virtualized infrastructure presents multiple networking challenges to IT organizations, including VM isolation, performance, and scale. Legacy Ethernet adapters, which cannot offload modern tunneling protocols, are limited in their ability to deliver maximum throughput or improve the efficiency of the data center.

QLogic offers 10GbE and 25GbE adapters with tunneling offload. Performance and efficiency gains of the QLogic FastLinQ 3400 Series 10GbE and FastLinQ 45000 Series 25GbE Intelligent Ethernet Adapters come from transmit large segment offload (LSO) and TCP segment offload (TSO), along with receive side scaling (RSS), transmit side scaling (TSS), and large receive offload (LRO—25GbE adapters only). Using LSO, TSO, RSS, TSS, and LRO, the adapters increase throughput and reduce host CPU load, improving the efficiency of the virtualized server while maximizing the scalability of the data center.

ESG Lab validated the performance and efficiency gains of the QLogic Ethernet adapters by comparing tunneling offload performance with tunneling encapsulation performed by the host CPU. The FastLinQ 3400 Series, at 10GbE, delivered an average 19% increase in throughput and a 6% increase in efficiency, as measured by the amount of data transferred per unit of work by the host CPU. The FastLinQ 45000 Series 25GbE Adapters also delivered an 18% increase in throughput. More impressively, the 25GbE adapter enabled the host to operate 55% more efficiently on average.

The business world runs at the fastest speed possible and therefore, your organization should, too. Deploying QLogic 10Gbps and 25Gbps Ethernet with tunneling offload, organizations can now operate at the speed of the business without having to worry about impacting the performance of the underlying virtualization infrastructure. If you are looking to modernize your infrastructure with forward-thinking technology that enables higher levels of efficiency with more throughput, ESG Lab suggests checking out QLogic 10GbE and 25GbE networking solutions with tunneling offload.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

⁵ Source: ESG Research Report, <u>2016 IT Spending Intentions Survey</u>, February 2016.

