


# Accelerating Network Virtualization Overlays

Assists and Offloads for VMware® VXLAN



- VMware® Virtual Extensible LAN (VXLAN) overlay network technology addresses key network scalability challenges associated with a hyperscale cloud infrastructure.
- QLogic® FastLinQ® 3400/8400 Series 10GbE Adapters offer the most complete portfolio of overlay network offloads.
- QLogic VXLAN offloads enhance VMware ESXi 5.5 to scale network performance by more than 32%.

## KEY REQUIREMENTS FOR VIRTUALIZED CLOUD-SCALE NETWORKS

A virtualized, multi-tenant environment must allow the unlimited transparent migration of workloads across physical servers—while controlling the cost and maintaining the Quality of Service (QoS) the customer requires. Most importantly, virtualized data centers need the flexibility of provisioning resources that span multiple geographic locations. At the same time, virtualized data centers must maintain isolation between tenants and still allow seamless management of the multi-tenant environment.

Virtualized cloud networks must also accomplish the following:

- Handle MAC address growth in conjunction with the explosive growth of VMs in the cloud data center.
- Accommodate increasingly larger number of VLANs to handle VM traffic segregation in a VLAN “sprawl” situation.
- Provide isolation of the physical L2 network. Each virtual tenant network has the illusion of being on its own physical network without impacting network performance and scalability.

## CHALLENGES IN VIRTUALIZED CLOUD-SCALE NETWORKS

In today’s cloud-scale networks, multiple organizations share the same physical infrastructure. Utilizing common processing and networking resources on an as-needed basis has become a standard business practice. Some cloud networks support implementations with dedicated physical servers for each customer, while other cloud network implementations support virtual servers per customer (on a common physical server).

A single network environment that hosts multiple customers (tenants) allows the customers to reduce up-front costs for processing or networking resources, yet provides them with the flexibility to increase or reduce the resources as needed. Such multi-tenant environments are increasingly using these new architectures due to the advantages of server virtualization. However there are some challenges to overcome in order to deliver a truly scalable, elastic, and secure virtualized cloud-scale infrastructure.

Figure 1 highlights some of the benefits of a virtualized cloud-scale network as well as the shortcomings this environment commonly faces. The benefits are outlined in green (+), and the shortcomings are outlined in red (-).

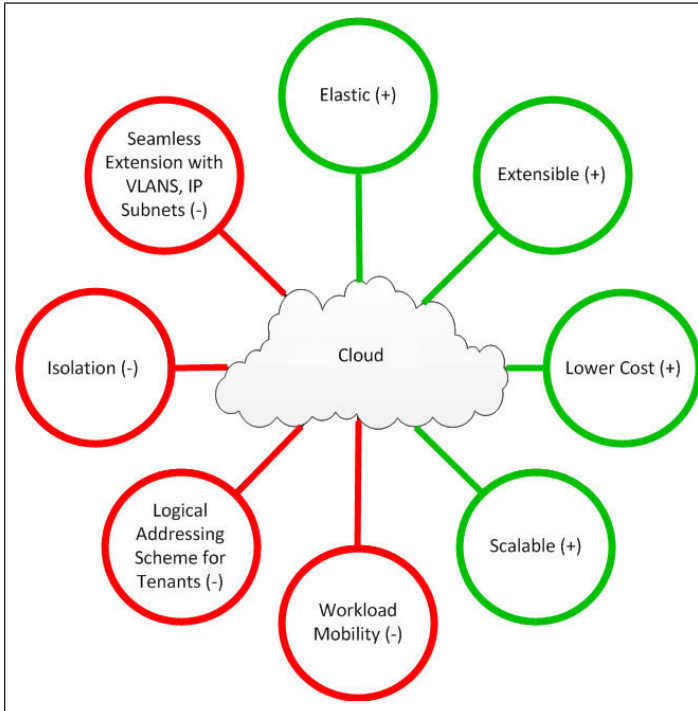


Figure 1. Benefits and Shortcomings of Virtualization in a Multi-Tenant Cloud Environment

## SOLUTIONS FOR VIRTUALIZED CLOUD-SCALE NETWORKS

To provide workload mobility and migrate across geographic locations, one cloud network solution is to decouple the physical and logical addressing schemes. The tenant uses the logical address while the network infrastructure sees the physical address. This decoupling enables the flexibility required by the virtualized cloud data center for creating a faster, fatter, and flatter network.

## SCALING CLOUD NETWORKS WITH VXLAN

VXLAN is used specifically to allow non-geographically located server farms (that are not on the same IP subnet) to transparently do VM migrations and to allow tenant VM users to seamlessly interact with other pre-defined workloads, regardless of the geographic/IP subnet location of the workloads.

Scaling the cloud network with VXLAN is the first step toward enabling logical, software-based networks that can be created on demand, allowing enterprises to leverage capacity wherever it is available. In other

words, VXLAN can now help companies build true global clouds that are the sum of their parts—rather than distinct sets of parts. A true global cloud essentially decouples the physical network design from the logical network design. VXLAN accomplishes these goals by using an overlay of tunnels, where layer two Ethernet frames are encapsulated within layer 4 UDP packets.

Technically, this is achieved by deploying VXLAN technology in the hypervisor vSwitch (VMware ESXi 5.1/5.5 using VXLAN tunneling protocol) and any 10GbE adapters. However choosing a non-fully featured 10GbE adapter would compromise the performance and flexibility of the cloud infrastructure.

## CHALLENGES OF VXLAN TECHNOLOGY AND QLOGIC SOLUTIONS

The main drawback of VXLAN is that encapsulated traffic by-passes the normal stateless offload features of a VXLAN-unaware adapter due to the following reasons:

- CPUs moving the packets individually (not as a block of data) to the send queue.
- CPUs calculating each send packet's checksum value.
- A single CPU core (normally core 0) handling all VXLAN tunneled traffic of all ports.

With the VXLAN-aware QLogic FastLinQ 3400/8400 Series Adapters, both the transmit and receive direction stateless offloads become available for use, which helps resolve all the challenges outlined above and greatly increases tunneled traffic throughput.

The QLogic FastLinQ 3400/8400 Series Adapter built-in assists and offloads for VXLAN helps network architects overcome the negatives (shown in Figure 1) in a virtualized private cloud environment and improve performance via tunneled traffic offloading.

QLogic adapters, with assists and offloads for VMware VXLAN technology, enable necessary QoS for multiple VMs in a multi-core environment by removing the single-queue restrictions and enabling multiple hardware queues for multiple VMs. For VMware ESXi Hypervisor, acceleration is enabled by UDP RSS packet task offloading support in ESXi 5.1 and UDP/TCP RSS, TCP Segmentation Offload (TSO), and TX Checksum Offload packet task offloading support available in ESXi 5.5 and 6.0.

QLogic FastLinQ 3400/8400 Series Adapter's assists and offloads for VMware VXLAN technology enables efficient distribution of network transmit and receive processing for VXLAN traffic across servers with multiple CPU cores. With QLogic Server Network Virtualization Overlay acceleration, the adapter provides the ability to distribute workloads efficiently across all processor cores.

Table 1. Adapter Assists and Offloads for VMware VXLAN

Stateless Offloads	Standard Adapter without NV Acceleration	QLogic FastLinQ 3400/8400 Series Adapter with NV Acceleration	Benefits of Acceleration
TCP Segmentation Offload (TSO)	✗	✓ ESXi 5.5, 6.0	Offloading reduces interrupts, which saves CPU cycles
TX Checksum Offload (CSO)	✗	✓ ESXi 5.5, 6.0	Offloading calculations saves CPU cycles
RSS Queue Acceleration	✗	✓ ESXi 5.1 UDP ✓ ESXi 5.5, 6.0 UDP/TCP	Spreads workload over multiple CPU cores, avoiding single-core bottlenecks

**PERFORMANCE METRICS OF VXLAN ESXi 5.5 TCP/UDP RSS AND TSO/ TX CSO ACCELERATION**

QLogic performed benchmark tests between two identical x86 servers running VMware ESXi5.5 set up for VXLAN encapsulation. The test drove network I/O using industry standard IxChariot application from multiple Windows Virtual Machines. Various different scenarios and transfer sizes were tested. For detailed steps on configuration refer to the [QLogic VXLAN Deployment Guide](#).

Table 2 shows a summary of the benchmark test results.

Table 2. VMware ESXi 5.5 TCP/UDP RSS, TSO and TX CSO Offload Test Results

Traffic	10GbE without VXLAN Acceleration	10GbE with VXLAN Acceleration	Performance Advantage
Receive (RX) Only	6.70 Gbps	8.43 Gbps	26% (Figure 2)
Bi-Direction (BiDir)	26.2 Gbps	34.6 Gbps	32% (Figure 3)

The QLogic FastLinQ 10Gb 3400/8400 Series Adapters showed a performance increase of 32% with VXLAN bidirectional traffic over a configuration without any acceleration (see Figure 3).




Figure 2. VMware ESXi 5.5 VXLAN RX Traffic Performance




Figure 3. VMware ESXi 5.5 VXLAN Bi-Directional Traffic Performance

**BENEFITS OF ACCELERATING NETWORK VIRTUALIZATION OVERLAYS**

QLogic 3400/8400 Series Adapters (available as low-profile dual 10GbE port PCIe adapters for rack and tower servers support L2 networking, iSCSI-Offload, and FCoE-Offload) combine the benefits of VXLAN overlays with stateless offload acceleration:

- QLogic 3400/8400 Series Adapters with VXLAN acceleration can be deployed on an existing VXLAN infrastructure, which reduces Capital Expenditure (CAPEX) costs and increases cloud-scale network performance.
- Cloud networks can scale up the number of VMs being deployed on their servers by up to 32% using QLogic 3400/8400 Series Adapter VXLAN acceleration, which increases the number of tenants with the same physical infrastructure and boosts their Return On Investment (ROI) on VXLAN deployments.
- Cloud network administrators using QLogic 3400/8400 Series Adapter VXLAN acceleration can provision additional bandwidth for resource-intensive applications or over-provision VMs with bandwidth for high-peak scenarios.
- VXLAN and QLogic 3400/8400 Series Adapters support a larger number of VLANs (16 million) over a wider non-geographic network, allowing tenant VMs to be deployed anywhere.
- VXLAN and QLogic 3400/8400 Series Adapters enable network administrators to create or migrate VMs dynamically over geographically separated locations, increasing flexibility.

## ABOUT CAVIUM

Cavium, Inc. (NASDAQ: CAVM), offers a broad portfolio of infrastructure solutions for compute, security, storage, switching, connectivity and baseband processing. Cavium's highly integrated multi-core SoC products deliver software compatible solutions across low to high performance points enabling secure and intelligent functionality in Enterprise, Data Center and Service Provider Equipment. Cavium processors and solutions are supported by an extensive ecosystem of operating systems, tools, application stacks, hardware reference designs and other products. Cavium is headquartered in San Jose, CA with design centers in California, Massachusetts, India, Israel, China and Taiwan.



Follow us:       

**Corporate Headquarters** Cavium, Inc. 2315 N. First Street San Jose, CA 95131 408-943-7100

**International Offices** UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan | Israel

Copyright © 2015 - 2017 Cavium, Inc. All rights reserved worldwide. QLogic Corporation is a wholly owned subsidiary of Cavium, Inc. QLogic, FastLinQ, and QConvergeConsole are registered trademarks or trademarks of Cavium, Inc. All other brand and product names are registered trademarks or trademarks of their respective owners.

This document is provided for informational purposes only and may contain errors. Cavium reserves the right, without notice, to make changes to this document or in product design or specifications. Cavium disclaims any warranty of any kind, expressed or implied, and does not guarantee that any results or performance described in the document will be achieved by you. All statements regarding Cavium's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.