Accelerating Network Virtualization Overlays

Assists and Offloads for Windows Server 2012 R2 Hyper-V NVGRE

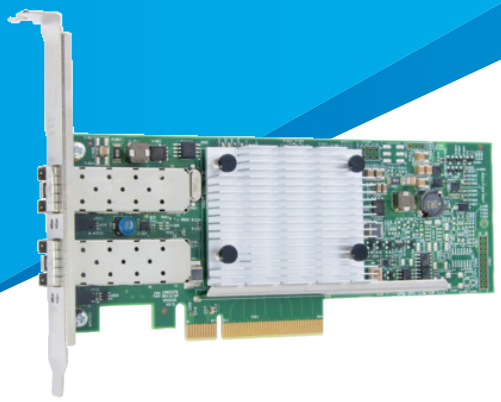

- Microsoft® Network Virtualization using Generic Routing Encapsulation (NVGRE) Overlay Network technology addresses key network scalability challenges associated with a hyperscale cloud infrastructure.
- QLogic® FastLinQ® 3400/8400 Series 10GbE Adapters offer the most complete portfolio of Overlay Network offloads.
- QLogic NVGRE offloads enhance Microsoft Windows Server 2012 R2 Hyper-V to scale network performance by more than 145%.

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 increasingly use these new architectures due to the advantages of server virtualization. However, this requires overcoming some challenges 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 (-).

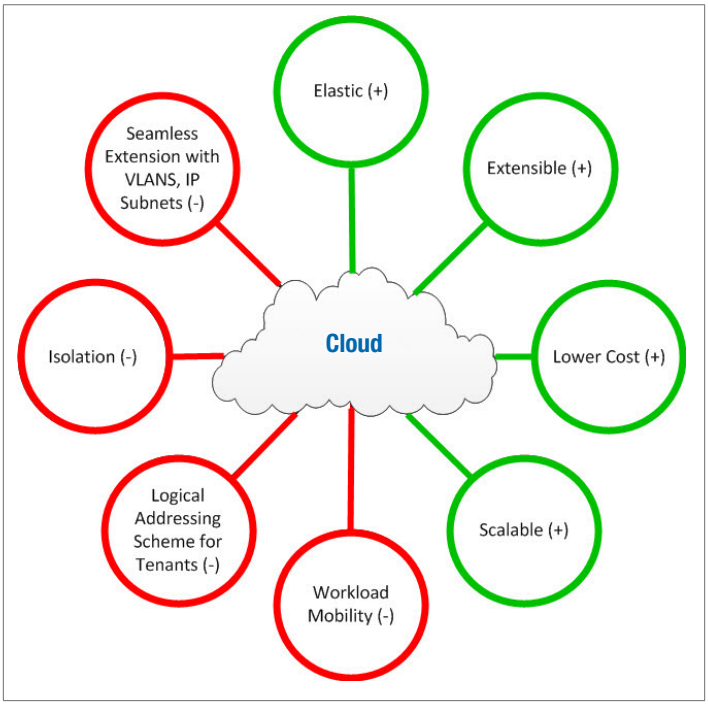

Figure 1. Benefits and Shortcomings of Virtualization in a Multi-Tenant Cloud Environment

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 NVGRE

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

Scaling the cloud network with NVGRE is the first step toward enabling logical, software-based networks created on demand, allowing enterprises to leverage the capacity wherever it is available. In other words, NVGRE 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. NVGRE accomplishes these goals by using an overlay of tunnels, where layer 2 Ethernet frames are encapsulated within layer 4 UDP packets.

Technically, this is achieved by deploying NVGRE technology in the hypervisor vSwitch (Windows Server 2012 R2 Hyper-V using NVGRE 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 NVGRE TECHNOLOGY AND QLOGIC SOLUTIONS

The main drawback of NVGRE is that encapsulated traffic by-passes the normal stateless offload features of an NVGRE-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 NVGRE tunneled traffic of all ports.

QLogic adapters, with advanced assists and offloads for NVGRE technology, enable the 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 Microsoft® Windows Server 2012 R2, Hyper-V acceleration is enabled for various packet task offloading on Virtual Machines running Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.

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

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

QLogic FastLinQ 3400/8400 Series Adapter with assists and offloads for NVGRE technology enables efficient distribution of network processing for NVGRE 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 Windows Server 2012 R2 Hyper-V NVGRE

Stateless Offloads	Standard Adapter without Acceleration	QLogic FastLinQ 3400/8400 Series Adapter with Acceleration	Benefits of Acceleration
Large Send Offload (LSO)	✘	✓	Offloading reduces interrupts, which saves CPU cycles
TX Checksum Offload (CSO)	✘	✓	Offloading calculations saves CPU cycles
RSS/TSS Queue Acceleration	✘	✓	Spreads workload over multiple CPU cores, avoiding single-core bottlenecks
Virtual Machine Queue (VMQ) Acceleration	✘	✓	Spreads workload over multiple CPU cores, avoiding single core bottlenecks

PERFORMANCE METRICS OF NVGRE WINDOWS SERVER 2012 R2 WITH ALL PACKET TASK OFFLOADS ENABLED

QLogic performed benchmarks between two identical x86 servers running Windows Server 2012 R2 Hyper-V setup for NVGRE encapsulation. The test drove network I/O using the industry standard IxChariot application from multiple Windows Virtual Machines. Various scenarios and transfer sizes were tested. Table 2 summarizes the test results.

Table 2. Test Results

Traffic	10GbE without NVGRE Acceleration	10GbE with NVGRE Acceleration	Performance Advantage
Receive (RX) Only	4.2 Gbps	8.8 Gbps	108% (Figure 2)
Bi-Direction (BiDir)	7.2 Gbps	17.6 Gbps	145% (Figure 3)

The QLogic FastLinQ 10Gb 3400/8400 Series Adapters showed a performance increase of 145% with NVGRE bidirectional traffic over a configuration without any acceleration (see Figure 3).

TEST CONFIGURATION

- **Operating System:** Windows Server 2012 R2 Hyper-V
- **Server:** Standard x86 Server – Dual Socket 2.9 GHz CPUs (six cores each)
- **Virtual Machines:** Four Windows Server 2012 R2 – 2v CPUs, 4GB RAM each
- **10GbE Adapter:** QLogic FastLinQ 3400 Series

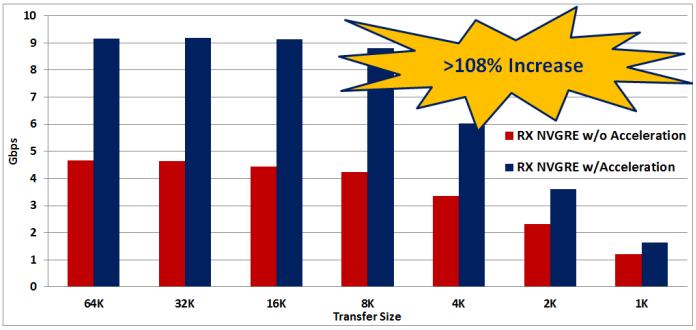

Figure 2. NVGRE RX Traffic Performance

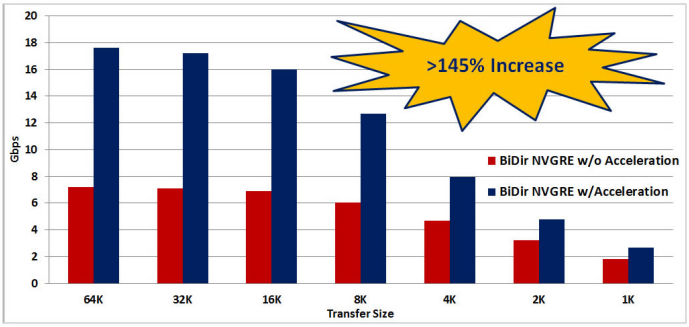

Figure 3. NVGRE BiDirectional Traffic Performance

BENEFITS OF ACCELERATING NETWORK VIRTUALIZATION OVERLAYS

QLogic FastLinQ 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 NVGRE overlays with stateless offload acceleration:

- QLogic FastLinQ 3400/8400 Series Adapters with NVGRE acceleration can be deployed on an existing NVGRE 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 145% using QLogic FastLinQ 3400/8400 Series Adapter NVGRE acceleration, which increases the number of tenants with the same physical infrastructure and boosts their Return On Investment (ROI) on NVGRE deployments.
- Cloud network administrators using QLogic FastLinQ 3400/8400 Series Adapter NVGRE acceleration can provision additional bandwidth for resource-intensive applications or over-provision VMs with bandwidth for high-peak scenarios.
- NVGRE and QLogic FastLinQ 3400/8400 Series Adapters support a larger number of VLANs (16 million) over a wider non-geographic network, which allows tenant VMs to be deployed anywhere.
- NVGRE and QLogic FastLinQ 3400/8400 Series Adapters enable network administrators to create or migrate VMs dynamically over geographically separate locations, which increases 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