The Impact of Virtualization on Cloud Networking

The adoption of virtualization in data centers creates the need for a new class of networking designed to support elastic resource allocation, increasingly mobile workloads, and maximum availability under production loads. Building a network that spans both physical servers and virtual machines with consistent capabilities requires a new architecture for designing and building the IT infrastructure. Performance, elasticity, and logical addressing structures must be considered, as well as the management of the physical and virtual networking infrastructure. Once deployed, a network that is virtualization-ready can offer many performance and productivity advantages over a common shared infrastructure.

Virtualization technologies from VMware, Citrix and Microsoft encapsulate existing applications, and abstract them from the physical hardware. Unlike physical machines, virtual machines are represented by a portable software image, which can be instantiated on physical hardware in seconds.

With virtualization comes elasticity where compute capacity can be scaled up or down, on demand, by adjusting the number of virtual machines running on a given physical server. Additionally, virtual machines can be migrated while in service from one physical server to another thereby simplifying maintenance and increasing up-time.

Extending this further, virtualization creates 'location freedom' enabling virtual machines to become portable across larger geographies. As cloud architectures and multi-tenancy capabilities continue to develop and mature, economies of scale can be achieved by aggregating resources across applications, business units, and even separate corporations to a common, yet segmented, infrastructure.

Elasticity, mobility, automation, and the density of virtual machines all demand new network architectures which deliver high performance, address portability, and incorporate support of the virtual machine as the new building block of the data center. Consistent network- supported and virtualization-driven policy and controls are necessary for visibility to virtual machines' state and location as they are created and then moved across a virtualized infrastructure.

The Virtual Machine Sprawl Challenge

A direct consequence of virtual server deployments is the explosion of virtual machines. New multi-core processing architectures running 10-20 or more virtual machines per server results in a significant increase in the number of elements to be managed. Simultaneously, there is a proportional sprawl of virtual switches that VMs use within each physical server. Every physical server that hosts VMs has a virtual switch, thus creating a 20-40x increase in the number of managed network elements. Essentially, the network access layer formed by virtual switches is now inside the server.

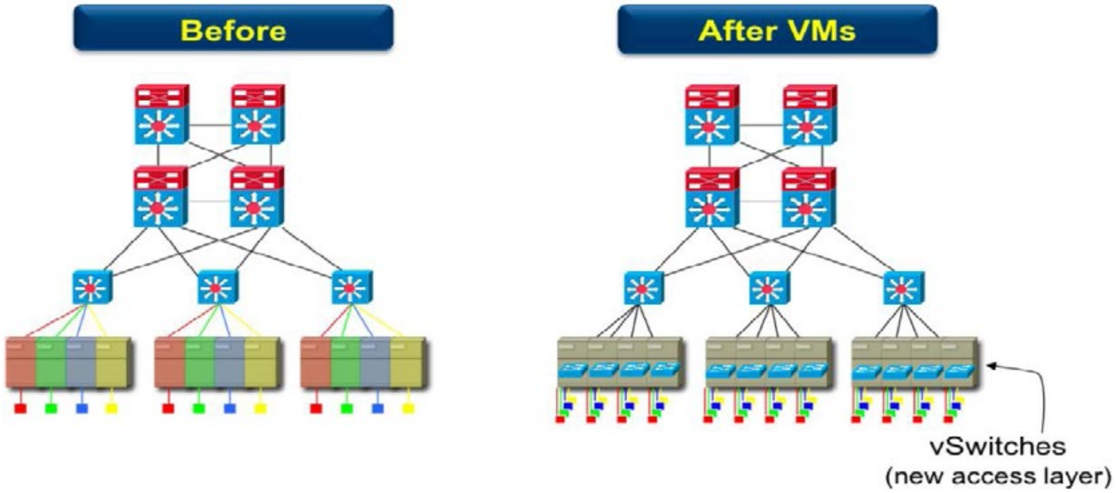

Virtualization Demands Consistent Cloud Networking

This massive virtual machine and switch infrastructure places new demands on the underlying network fabric for seamless transactions including: user-to-VM, VM-to-VM, VM-to- data, VM-to-Fault Tolerant Peer, and for VM Mobility. These new types of transactions demand a network architecture purposefully built to support these demands: a cloud network architecture. Specific challenges include:

1) **VM Explosion:** Since many virtual machines can be instantiated on a physical server, the physical NIC bandwidth utilization increases proportionally. This implies that traditional oversubscribed network topologies must be re-architected for virtualization and private clouds.

Portable VM images are several gigabytes in size and the network must now carry larger amounts of data to move and migrate VMs. Also, workload elasticity implies the allocation of virtual machines to compute resource is scaled up or down programmatically, based on policies such as load, time of day and power/cooling availability. The network has to be designed to gracefully handle the peak load within this virtualized environment.

2) **Cloud Applications:** New cloud applications integrate Web 2.0 and rich media technologies, often through network mash-ups which can be accessed by millions of users worldwide. This leads to large numbers of transactions that traverse the network with much higher downstream (VM-to-user) traffic.

Cloud application workloads are designed to distribute computing tasks across multiple layers of worker and data nodes, requiring unprecedented VM-to-VM interactions. Also, the RESTful paradigm leads to compute states being kept only on data nodes, thus demanding constant VM access to back-end databases over the network fabric.

3) **VM Migration:** The virtual machine is becoming increasingly mobile. Some use cases include: retiring/upgrading servers & operating systems, moving workloads off of low- utilization servers so they can be shut off to save power, moving VMs off of a high utilization server for application performance management, or opportunistically migrating workloads to lower-cost compute enclaves. VM mobility requires networks to have larger and flatter Layer 2 domains so that IP addresses and in-progress client transactions are not disrupted when VMs are moved.

4) Virtual Switch Management: Server administrators typically manage virtual networks, because network administrators do not have direct access to built-in virtual switches. For large-scale virtualization and private cloud environments, this creates a major challenge as consistent network-wide policies, monitoring, and diagnostics all need to be applied to large numbers of virtual switches across the infrastructure.

5) **Cloud Reach:** With emerging Infrastructure as a Service (IaaS) public cloud and virtual private cloud offerings, there is currently no way to keep network policy and accounting state intact as the virtual machine moves from one provider to the next. Network state needs to be communicated in a trusted manner to the receiving IaaS cloud provider both to facilitate provider interoperability, and to make it possible for the enterprise to retain control over its security.

Designing Virtualization-Optimized Cloud Networks

Building the combination of virtual and physical networks to support physical, virtual, and cloud deployments is non-trivial. Performance, resiliency, policy control, and management visibility must be considered in the design. The characteristics, discussed earlier, of networks that support virtualization and cloud computing require that legacy network practices be abandoned in favor of modern cloud networking architectures, such as those enabled by using Arista's 7000 Family of Data Center switches and its Extensible Operating System (EOS). Specifically, the key characteristics are:

Cost-Effective Wirespeed Performance

High performance 10GbE networking is a must for the core network (with future expansion to 40/100GbE), and in many cases the edge network, especially in the case of blade servers. As gigabytes of VM images move across the network fabric, 10Gb Ethernet is the optimal transport technology to build high performance, highly responsive networks capable of handling the peak bandwidth demands of cloud and VM workloads. Arista's non-blocking, wire-speed 10Gb Ethernet switches are ideally suited for such high-performance environments.

Symmetric cross-sectional bandwidth

The majority of network architectures today were designed to support client-server traffic and one application in particular: E-Mail. Virtualization and modern application workloads have changed the oversubscription rules that were the basis of legacy network designs. Highly utilized host links as well as the symmetry of user-to-VM and VM-to-VM traffic requires that ingress and egress switching bandwidth be highly balanced, having an ingress-to-egress bandwidth ratio of 1:1, or at most 2:1 (versus 20:1 or even 40:1 in legacy designs). For instance, a top-of-rack (leaf) switch with 40 1GbE server access ports requires 4 10GbE uplinks using a 1:1 design rule. A 48 port 10GbE switch (provisioned either as blade server access switch, or as a spine switch that aggregates multiple leaf switches), operating at a 2:1 oversubscription ratio would require 32 10GbE ports for ingress and 16 10GbE ports for egress. The wire-speed architecture of the Arista 1Gb and 10Gb switches enables the deployment of cloud networks with balanced traffic patterns. If you don't oversubscribe, you don't have to manage oversubscription.

Leaf-Spine Architecture

Constant inter-VM communication and VM mobility demand large Layer 2 domains, thus making modern two-tier leaf-spine architectures significantly more preferable to a traditional three-tier design. With this architecture, a VM communicates to any other VM in three physical hops or less.

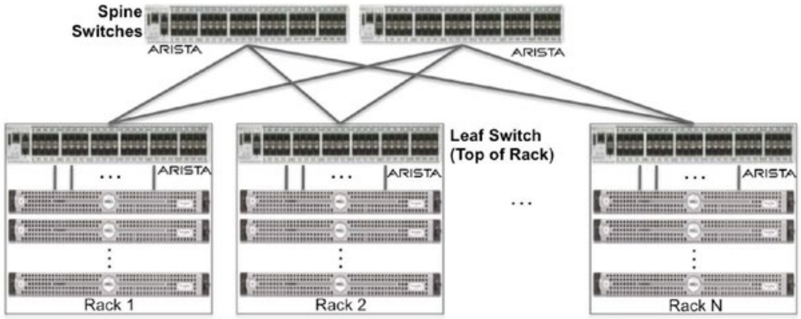

Low-latency Switching

Reducing packet forwarding latency and using proper bandwidth provisioning is critical for improving application response time. Prudent deployment of switches that leverage cut-through packet processing, in place of store-and-forward switches, provides 5-10 times the reduction in per-switch latency. Achieving system-wide latency reduction through appropriate bandwidth provisioning, while avoiding packet queuing, drops and retransmit errors, demands an architecture that utilizes a 10GbE transport substrate and a two-tier network design. Using these methods to reduce latency improves the efficiency of compute processing and generates business results faster with less power and cost. Recent developments of Virtual Machine Fault Tolerance via copying updated memory pages from the primary VM to a standby VM also place increased importance on having an ultra-low- latency infrastructure.

Resilient Networking

Principles of fault-tolerant computing ensure that workloads are not impacted when a few compute nodes (whether physical or virtual) fail. The modern cloud network needs to enable resiliency at the service level. Switches fail mostly because of outmoded operating system and software architectures. Like fault-tolerant compute principles, network operating systems also need to be engineered with fault-tolerant core operating system design principles. Additionally, a high-speed network control-plane as well as the separation of the control and data planes have become table stakes for resilient cloud networks. Arista's EOS (Extensible Operating System) is engineered from the ground-up to include these capabilities natively so that the highest levels of network resiliency can be achieved.

Virtualization

If VMs could easily move from one cloud to another there would be a larger, more strategic payoff; yet today's manually managed network is based on static port allocation and static policy definition that has no linkage between the physical and virtual networks, and no mechanisms for consistent policy enforcement in the network operating system. For seamless consistency between the virtual and the physical switches, the virtual switches provide transparent redirection using various standards-based mechanisms, including IEEE 802.1Q VLAN tagging, as well as using MAC addresses and/or tunnels that are transparent to the physical switches. Proprietary tags need to be avoided as they limit vendor-choice and interoperability.

Managing Virtualization - Optimized Cloud Networks

As VM Farms have evolved and grown in size, the networks that support them, both physical and virtual, have grown as well. Operating these networks has become increasingly challenging due to the lack of proper tools. On the server side as well as on the virtual network, useful tools have emerged to assist VMware Administrators with the day-to-day challenges of running a VM Farm. However, the equivalent tools have not emerged to help the Network Administrator resolve conflicts existing between physical and virtual networks.

Consistent network management across both physical and virtual networks demands that heterogeneous virtual switches be managed by network administrators using well-known command line interface (CLI) to simplify adoption, yet also provide more programmatic abstractions such as SNMP, XML and XMPP to enable API-based management of the network infrastructure from the Cloud OS. In order to maintain configuration and management consistency across virtual and physical networks during VM migration, it is imperative that management be consistent across physical networks, virtual machines, and cloud implementations.

Introducing Arista's VMTracer

Key requirements for enabling seamless operational management across virtual and physical networks are:

- The ability for network administrators to treat virtual switches as an extension of the physical network and manage them in a similar manner as with physical switches today
- Relieve the burden of network configuration from VM host administrators
- Provide visibility for troubleshooting and auto-discovery to the network administrator
- Provide a central way of globally discovering and enforcing consistent policy across physical and virtual switches

VM Tracer is the second tool for Network Administrators that Arista has developed to deliver operational simplicity for virtual environments. It is natively integrated into Arista 7000 switches, and along with Arista's EOS, it works with all Arista switches. VM Tracer links the physical Arista switches to VMware's vCenter and creates an adaptive infrastructure whereby the network automatically responds to sensed changes in the virtual machine network.

VM Tracer works with VMware vSphere 4.0 and vSphere 4.1. It utilizes published vCenter APIs, leveraging existing standards, and thus it works across all editions of vSphere. Some of the key capabilities of VM Tracer include adaptive segmentation of VLANs, quality of service, auto discovery and VM Host View, all designed to address the challenges virtual machines bring to the physical network.

While Arista's initial vEOS tool provides network administrators a familiar command line interface and SNMP interface to the vSwitch infrastructure, Arista's VM Tracer is directly integrated into the Arista 7000 Family of switches for unprecedented visibility into the virtualized environment, seamless integration using the familiar industry-standard CLI, and automatic configuration of tasks and policy by integrating natively with VMware vCenter. This tool presents high value advantages for both the Network Administrator as well as the VMware Administrator, and makes it comfortable for the two groups to work together. Arista customers will enjoy visibility into the vSwitch, the VM farm, and policy control that is natively integrated with VMware's vCenter, as well as a dynamically responsive network within the virtualized environment.

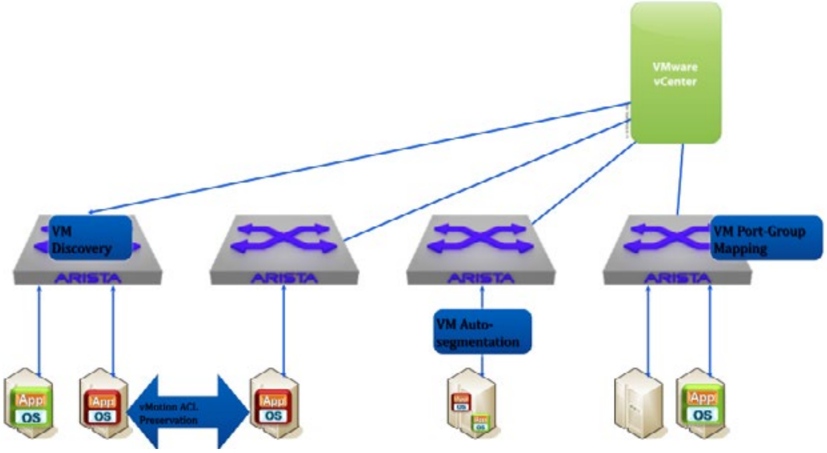

Table 1: Virtualization Support		
	Arista Virtualization	**Legacy Virtualization**
vSwitch	VMware switch you already have	Purchase new switch
Server <--> Switch	No tags	Proprietary VnTags
Visibility to virtual network	Yes (cli on switch)	Special Network Management Tools
Summary	Add value where necessary	Sell & lock-in all layers

Summary

The combination of virtualization and cloud computing creates a computing paradigm that requires careful network considerations. A cloud network must support the abstraction of virtualization services from specific physical servers. In particular, 10GbE networking and orchestration mechanisms must allow on-demand deployment of network bandwidth, virtual machine resources, and support isolation between different workloads and customers.

The migration from network silos to virtualization and clouds calls for the replacement of vastly oversubscribed high latency legacy switches with a network infrastructure designed for virtualization and the cloud, and featuring a new generation of products such as Arista's 7000 Family switches deployed in two-tiered leaf and spine designs. The Arista 7000 Family, powered by EOS and featuring the operational simplicity of administrative tools like VM Tracer and vEOS, will support the smooth migration of network infrastructures to VMware vSphere vNetwork Distributed Switches, Citrix Xen, Kernel Virtual Machine (KVM), and Microsoft Hyper-V hypervisors. Orchestration is achieved via industry standard discovery protocols, CLIs and extensible APIs. Physical, virtual, and cloud networks must not only co-exist, but also must be managed seamlessly for scalable deployments.

The combination of vEOS, VM Tracer and the Arista 7000 Family enables physical network and virtualization operators, as well as private and public cloud systems, to seamlessly migrate virtual machines across network, server, and organizational boundaries with consistent policy and accounting. It is a dramatic departure from the silos of today as workloads can now migrate from physical to virtual to cloud networking.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office
1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062