

**Cloud Networking:
A Network Approach that Meets the Requirements of Cloud Computing**

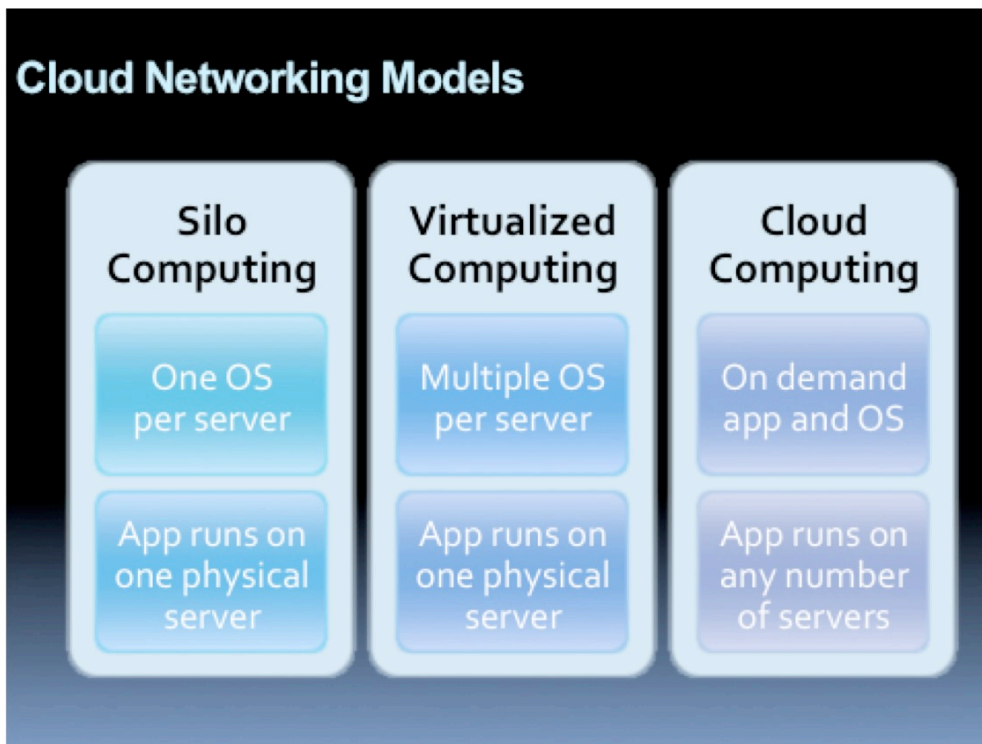
CQ2 2011

ARISTA

Arista's Cloud Networking

The advent of Cloud Computing changes the approach to datacenters networks in terms of throughput, resilience, and management. Cloud computing is a compelling way for many businesses, small (private) and large (public) to take advantage of web based applications. One can deploy applications more rapidly across shared server and storage resource pools than is possible with conventional enterprise solutions. Deploying modern web applications across a cloud infrastructure enables a new level of agility that is very difficult to accomplish with traditional silo computing model. New Computing models for virtualization and cloud require an ultra-scalable network infrastructure with a very large number of 1/10 Gigabit and future 40/100G higher-speed Ethernet connections. Clouds necessitate a high network availability since a network failure could affect thousands of servers and new levels of open integration. Finally, clouds require automated ways to rapidly bring-up, provision, virtualize, and administer the network.

In short, the ability to scale, control, automate, customize, and virtualize the cloud network is an important evolution to "data center in the box" approach.



The Shift to Green Clouds and 10Gigabit:

Cloud networking is different from traditional enterprise network designs. Large content providers are building public clouds of 100,000+ servers in physical containers while smaller private clouds are being constructed for thousand servers still co-existing with classical enterprise designs. In many ways, the difference between building a Cloud verses traditional racks is the difference between pre-fabricated houses and full-custom house construction. While full-custom house construction is more flexible and allows every decision to be made custom (proprietary), it is complex and expensive. There is no question that prefabricated houses are more cost-effective, quicker to build and easier to manage and can be done without sacrificing any reliability. With the right planning you can even add some custom "hooks".

In a cloud, the networking layer is an integral part of the computing cloud and is part of the overall

solution rather than a separate piece. The operational and acquisitions costs as well as power consumption are significant. In today's market, green clouds and power efficiencies are a bigger and growing part of the equation. Depending on the location, power costs can be as low as \$0.03/KWh (Pacific northwest) to \$0.30/KWh in Europe. Power efficient 10 Gigabit Ethernet such as Arista's 7050, 7100, and 7500 family can reduce total power consumption and therefore cost by 10-20%. This can result in significant cost savings in a typical cloud where costs can mount into thousands of dollars.

A New Cloud Stack:

Unlike the traditional OSI stack model of Layers 1 through 7 with distinct separation between network layers (Layer 2/3/4) and application layers (Layer 7), Cloud Networking transcends the layers and blurs these boundaries by coupling the network infrastructure with machines and modern web applications. Stateless Servers that separate persistent state from the server resource pool require a non-blocking network fabric that is robust, resilient and able to isolate application domains.

The New Cloud Stack

Stack Layer	Examples	Benefits
Application	SAAS, PAAS, Web apps, Internal apps	On-Demand Scheduler maximizes application access
OS	Any version of Linux, Windows, Solaris	Any version of Linux, Windows, Solaris
Hypervisor	ESX, Hyper-V, KVM, XVM	Decouples App + OS from Hardware
Server	Bare-Metal Stateless Server	Minimizes Server Administration Cost
Storage	Network Attached File Storage	No separate SAN needed
Network	Cloud Networking, CloudVision®	Enables dynamic network bring-up, provisioning, configuration, and deployment

Cloud Networking Attributes

Arista's Cloud Networking™ goes beyond classical networks to redefine scalability, administration and management processes. Solving these problems requires a new design approach for the cloud network fabric, starting with the software architecture.

Some of the key considerations are:

Scalability. The cloud network must scale to the overall level of throughput required to ensure that it does not become the bottleneck. This means the cloud networking fabric must handle throughputs that will reach trillions of packets per second in the near future.

Low Latency. The cloud network must deliver microsecond latency across the entire network fabric since low latency improves application performance and server utilization.

Guaranteed Performance. The cloud network must provide predictable performance to service a large number of simultaneous applications in the network, including video, voice and web traffic.

Self-Healing Resilience. Cloud networks operate 24x7, so downtime is not an option. This requires a network architecture that offers self-healing and the ability for transparent in-service software updates.

Extensible Management. Real-time upgrades and image/patch management in a large cloud-network is a daunting challenge to network administrators. A vastly simpler approach is required to handle networks of this size, which automates provisioning, monitoring,

maintenance, upgrading, and troubleshooting.

Cloud networks simply cannot be achieved with today's monolithic networking software stacks. Arista's Extensible Operating System (EOS™) can distribute software and configuration information across the entire fabric, enabling seamless consistency across the infrastructure. Because of its finely granular modular architecture and by separating process and state, Arista EOS is extremely resilient and provides in-service software upgrades for specific pieces of the functionality without requiring any downtime. Finally, with its extensibility, EOS offers customers the ability to tailor the behavior of the network infrastructure to the requirements of their own management workflows.

Scaled management for Cloud infrastructures

Arista's Extensible Operating system (EOS) provides unique management capabilities that are optimized for cloud environments: VMTracer, ZTP, CloudVision, and LANZ. Together, they provide the tools necessary for IT administrators to effectively bring up, manage, provision, and monitor their network in a dynamic and automated manner that meets the requirements of today's cloud infrastructures.

Dynamically adapting the network to changes in the virtualized environment. Arista VM Tracer interfaces with VMWare vSphere to provide complete visibility into the virtual network, and respond automatically and dynamically to changes in the virtual infrastructure. The ensuing virtualization of the network is an indispensable complement to server virtualization in virtualized environments.

Automating network bring-up and provisioning. Arista ZTP (Zero Touch Provisioning) fully automates the bring-up of network switches and other components of the cloud infrastructure. By using standard based protocols and offering advanced configuration and scripting capabilities, ZTP empowers the IT administrator to tailor the configuration of cloud racks or clusters in the most complex cloud environments without requiring any human intervention.

Controller-based administration and management of the network elements. Arista CloudVision provides a framework for standard-based single-point of management and administration of the network elements. In addition, Arista EOS is compatible with the OpenFlow protocol, which allows Arista network elements to interface with a host of OpenFlow controllers that are available today.


Advanced monitoring and troubleshooting of the cloud infrastructure. Arista LANZ (Latency Analyzer) measures latency on all ports and all network elements in a cloud network with microsecond granularity, providing unprecedented visibility into impending congestion events and empowering the IT administrator with a powerful troubleshooting and diagnostics tool.

Two-Tiered Cloud Topologies

In terms of cloud topology, what matters to customers is economics, performance and reliability. Typically servers are single port connected to a leaf switch for access, which is then connected to multiple load-sharing spine switches. Redundancy and Improved cross-sectional cloud bandwidth can be achieved via dual-homed connections from leaves to spines, with active multi-pathing across links and chassis. In a leaf-switch topology the connections between the server and the switch can be low cost copper cables. The connections from the leaf-switches to the core spine-switch are typically fiber, although in a container environment all connections may be low-cost copper cable.

In cloud networking designs, many customers desire to build fairly large Layer 2 clouds since they are easier to manage than Layer 3 clouds. Other customers organize Layer 2 subnets in more

conventional sizes of hundreds to thousands of nodes. Arista's products have sufficient Layer 2 and Layer 3 forwarding entries to support either one of these cloud architectures. A big issue in many customers' minds is actually a *balanced* cloud network of latency, reliable packet buffering, non-blocking throughput and total scalability. Solving for one metric without the other is not good enough. For clouds demanding large data analytics queries with wide range of networking protocols, including UDP and Multicast, larger packet buffers may provide performance benefits as they avoid packet loss. For specific market data and technical computing applications, consistent switch latencies of 500 nanosecond to 2 microseconds are paramount unlike most industry switches of today that are on an average ten times worse. A cloud network must also pay attention to the interconnections of 1/10Gigabit aggregation and design for terabit class non-blocking fabric with uncompromised throughput and capacity. Link Aggregation of multiple 10 Gigabit and future options for 40G and 100G is essential for distribution as it avoids over-subscription in the data center.



At the leaf-level, most clouds use one Gigabit Ethernet but migration to 10 Gigabit Ethernet is imminent. With next generation servers and storage systems ability to drive the network at higher and higher rates it is expected that one Gigabit Ethernet ports will increasingly transition to 10 Gigabit Ethernet in order to realized the full potential of the server hardware. Server vendors are adding 10G NICs to the motherboard with low-cost cables, reducing the 10G connection. A price point of twice the typical gigabit Ethernet connection of \$200-\$250/port is possible for 10GE. Furthermore the performance of the servers is expected to more than double. For 5-10K node topologies higher-speed Ethernet (10G now, 40 and 100 Gigabit future) and larger-density spine switches will be essential to provide scalable non-blocking bandwidth.

Summary:

Andy Kessler in his editorial for WSJ demystified Cloud well in layman's terms in The War for the Web:

- The Cloud. The desktop computer isn't going away. But as bandwidth speeds increase, more and more computing can be done in the network of computers sitting in data centers - aka the "cloud."...
- The Edge. The cloud is nothing without devices, browsers and users to feed it....
- Speed. - Speed. Once you build the cloud, it's all about network operations....
- Platform. ...Having a fast cloud is nothing if you keep it closed. The trick is to open it up as a platform for every new business idea to run on, charging appropriate fees as necessary.....

This resonates with Arista Networks and more importantly Arista customers looking for that integrated cloud container approach towards Data Centers and a dramatic departure from closed networks of the past.