A Novel Network Approach for Cloud Computing Models

The advent of cloud computing changes the approach to data center networks in terms of throughput and resilience. The ability to scale, control, visualize and customize the cloud network is an important evolution to a "data center-in-the-box" approach. 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 models. New computing models for virtualization and cloud environments require a scalable network infrastructure with a very large number of 1/10 Gigabit Ethernet (GbE) connections and future higher-speed 40/100 GbE connections. Clouds also require a high level of network availability since a network failure can affect thousands of servers. Lastly, new levels of open integration are required to achieve orchestrated management.

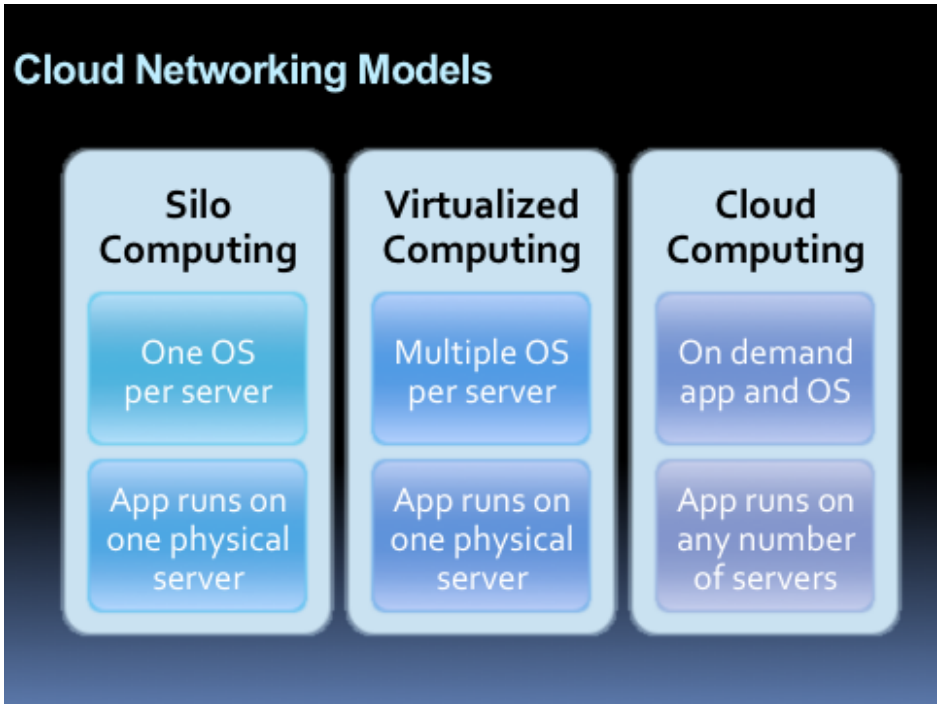

The Shift to Green Clouds and 10GbE:

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 thousands of servers still co-existing with classical enterprise designs. In many ways, the difference between building a Cloud versus traditional silos 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, easier to manage and can be done without sacrificing any reliability. Custom "hooks" can also be added with the right planning.

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, acquisition costs as well as power consumption are significant. In today's market, green clouds and power efficiencies are becoming a more significant part of the equation. Depending on the location, power costs can be as low as $0.03/KWh (US Pacific Northwest) to $0.30/KWh in Europe. Power efficient 10 GbE such as Arista's 7100 Series of switches can reduce total power consumption. This can result in significant cost savings of 10%-20% 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 though 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, which separate persistent state from the server resource pool, require a non-blocking network fabric, that is robust and able to isolate application domains.

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	Enables dynamic application provisioning and deployment

Cloud Networking™ Attributes

Arista 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 by separating process and state. This, in turn, enables administrators to provision clouds, add new servers or services, and update software images with ease.

Two –Tiered Cloud Topologies

In terms of cloud topology, what matters to customers are 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 leafs to spines, with active multi-pathing across links and multiple 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 switching 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 achieving 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 a 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 600 nanosecond to 2 microseconds are paramount unlike most industry switches of today that are on an average ten times higher. A cloud network must also pay attention to the interconnections of 1/10 GbE aggregation and design for terabit class non-blocking fabrics with uncompromised throughput and capacity. Link aggregation of multiple 10GbE and future options for 40GbE and 100GbE is essential for distribution as it avoids over-subscription in the data center.

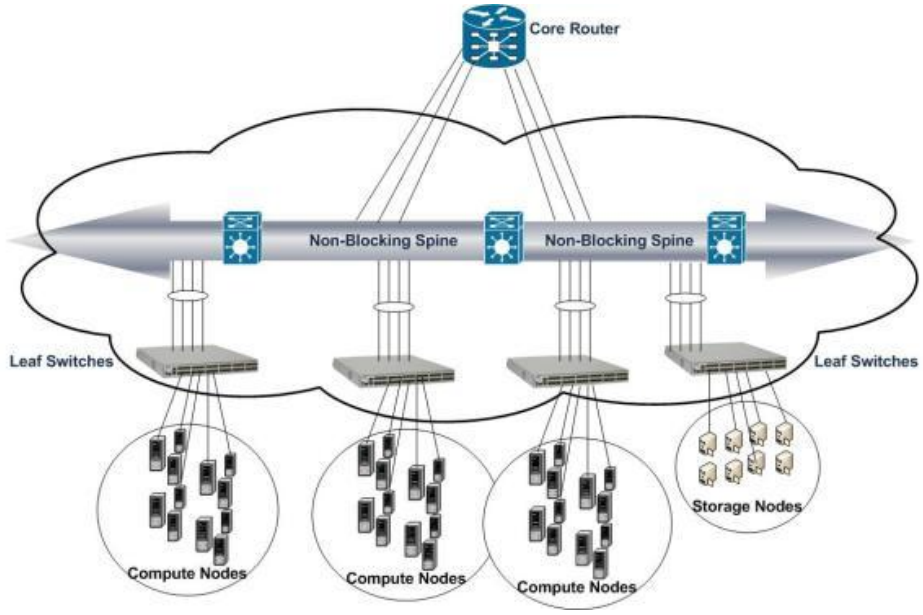

At the leaf-level, most clouds use one GbE but migration to 10GbE is imminent. Next generation servers and storage systems will enable the network to run at higher and higher rates, therefore, it is expected that one GbE ports will rapidly transition to 10 GbE in order to realize the full potential of the server hardware. Server vendors are adding 10G NICs to the motherboard with low-cost cables, reducing the 10G connection cost. A price point of twice the typical GbE connection ($200-$250/port) is possible for 10GbE. Furthermore, the performance of servers is expected to more than double. For 5-10K node topologies, higher-speed Ethernet (10G now, 40G and 100G future) and larger-density spine switches will be essential to provide scalable non-blocking bandwidth. Cost-effective, scalable solutions with resilience and open interface access are key differentiators.

Summary

Andy Kessler, in his editorial for WSJ, demystified the 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. 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, with Arista customers looking for an integrated cloud container approach towards data centers. This is a dramatic departure from closed networks of the past.