

Sponsored by:

NETWORKWORLD

This story appeared on Network World at <http://www.networkworld.com/news/tech/2011/110103techupdate-network-trends.html>

Trends reshaping networks

By Ken Won, Director of Product Marketing, Force10 Networks , Network World
January 03, 2011 02:45 PM ET

This vendor-written tech primer has been edited by Network World to eliminate product promotion, but readers should note it will likely favor the submitter's approach.

Server and storage environments have seen a lot of changes in the past ten years, while developments in networking have remained fairly static. Now the demands of virtualization and network convergence are driving the emergence of a host of new network developments. Here's what you need to know and how to plan accordingly.

* **Virtualization.** Virtualization has allowed us to consolidate servers and drive up utilization rates, but virtualization is not without its challenges. It increases complexity, causing new challenges in network management, and has a significant impact on network traffic.

Prior to virtualization, a top of rack (ToR) switch would support network traffic from 20-35 servers, each running a single application. With virtualization, each server typically hosts 4-10 VMs, resulting in 80 to 350 applications being supported by a single ToR rather than multiples switches, consolidating the network traffic. As a result, the ToR is much more susceptible to peaks and valleys in traffic, and given the consolidation of traffic on one switch, the peaks and valleys will be larger. Network architectures need to be designed to support these very large peaks of network traffic.

* **Flattening the network.** It is not possible to move VMs across a Layer 3 network, so the increased reliance on VMs is driving the need to move toward flattening the network – substituting Layer 2 architecture for older Layer 3 designs. In addition, a flatter network reduces latency and complexity, often relying only on ToR switches or end of row (EoR) switches connected to core switches. The result is lower capital expenses as fewer switches need to be purchased, the ability to migrate VMs across a larger network, and a reduction in network latency.

* **TRILL.** To facilitate implementation of Layer 2 networks, several protocols have emerged. One major change



is the replacement of Spanning Tree Protocol (STP). Since there are usually multiple paths from a switch to a server, STP handled potential multipath confusion by setting up just one path to each device. However, STP limits the network bandwidth and as the need developed for larger Layer 2 networks, STP has become too inefficient to do the job.

Enter the Transparent Interconnection of Lots of Links (TRILL). TRILL is a new way to provide multipath load balancing within a Layer 2 fabric and is a replacement for STP. TRILL has been defined by IETF and maps to 802.1q capabilities within the IEEE. TRILL eliminates the need to reserve protected connections for future use, and thereby stranding bandwidth.

* **Virtual physical switch management.** In a virtualized environment, virtual switches typically are run on servers to provide network connectivity for the VMs in the server. The challenge is each virtual switch is another network device that must be managed. Additionally, the virtual switch is often managed through the virtualization management software. This means the virtualization administrator is defining the network policies for the virtual switch while the network administrator is defining the network policies for the physical switch. This creates potential problems as two people are defining network policies. Since it is critical to have consistent security and flow control policies across all switches, this conflict must be resolved.

EVB (Edge Virtual Bridging) is an IEEE standard that seeks to address this management issue. EVB has two parts, VEPA and VN-Tag. VEPA (Virtual Ethernet Port Aggregator) offloads all switching from the virtual switch to the physical switch. All network traffic from VMs goes directly to the physical switch. Network policies defined in the switch, including connectivity, security and flow control, are applied to all traffic.

If the data is to be sent to another VM in the same server, the data is sent back to the server via a mechanism called a hairpin turn. As the number of virtual switches increases, the need for VEPA increases because the virtual switches require more and more processing power from the server. By offloading the virtual switch function onto physical switches, therefore, VEPA removes the virtualization manager from switch management functions, returns processing power to the server, and makes it easier for the network administrator to achieve consistency for QoS, security, and other settings across the entire network architecture.

In addition to VEPA, the IEEE standard also defined multi-channel VEPA, which defines multiple virtual channels, allowing a single physical Ethernet connection to be managed as multiple virtual channels.

The second part of EVB is VN-Tag. VN-Tag was originally proposed by Cisco as an alternative solution to VEPA. VN-Tag defines an additional header field in the Ethernet frame that allows individual identification for virtual interfaces. Cisco has already implemented VN-Tag in some products.

* **VM migration.** In a virtualized data center, VMs are migrated from one server to another to support hardware maintenance, disaster recovery or changes in application demand.

When VMs are migrated, VLANs and port profiles need to be migrated as well to maintain network connectivity, security and QoS (Quality of Service). Today, virtualization administrators must contact network administrators to manually provision VLANs and port profiles when VM's are migrated. This manual process can greatly impact the data center flexibility as this manual process could take minutes, hours or days, depending on the network administrator's workload.

Automated VM/network migration addresses this problem. With automated VM/network migration, the VLAN and port profiles are automatically migrated when a VM is migrated. This eliminates the need for network administrators to do this manually, ensuring that VM/network migration is completed immediately.

* **Convergence.** The other major trend underway in data center networking is fabric convergence. IT managers want to eliminate separate networks for storage and servers. With fabric convergence they can reduce

management overhead and save on equipment, cabling, space and power. Three interrelated protocols that enable convergence are Fibre Channel over Ethernet (FCoE), Ethernet itself (which is being enhanced with Data Center Bridging (DCB), and 40/100GB Ethernet.

Storage administrators initially gravitated to Fibre Channel as a storage networking protocol because it is inherently lossless, as storage traffic can't tolerate any loss in transmission. FCoE encapsulates Fibre Channel traffic onto Ethernet and allows administrators to run storage and server traffic on the same converged Ethernet fabric. FCoE allows network planners to retain their existing FCoE controllers and storage devices while migrating to a converged Ethernet network for transport. This eliminates the need to maintain two entirely separate networks.

DCB comes into the picture because it enhances Ethernet to make it a lossless protocol, which is required for it to carry FCoE. A combination of FCoE and DCB standards will have to be implemented both in converged NICs and in data center switch ASICs before FCoE is ready to serve as a fully functional standards-based extension and migration path for Fibre Channel SANs in high performance data centers.

Another advancement being driven by the rise in server and storage traffic on the converged network is the move to 40GB and 100GB Ethernet. With on-board 10GbE ports expected to be available on servers in the near future, ToR switches need 40GB Ethernet uplinks or they may become network bottlenecks.

Preparing for the Future

Some of the protocols discussed are still in development, but that doesn't mean you shouldn't begin planning now to leverage them. Here are some ways to prepare for the changes:

- * Understand these technologies and determine if they are important to implement.

- * Evaluate when to incorporate these new technologies. While you may not have the need today, you should architect the data center network to allow for future upgrades.

Plan for open standards. Data center network architectures should be based on open standards. This gives customers the most flexibility. Proprietary products lock customers into a vendor's specific products. Products based on open standards give customers architectural freedom to incorporate future technologies when they become available and when it makes sense for the business.

- * Plan on upgrading to 10GbE sometime in the next year or two, then have a way to upgrade to 40/100GbE when prices fall. The 10GbE price per port is expected to fall significantly next year, driving a significant increase in 10GB port installations.

It is clear the network is becoming a dynamic part of a virtual computing stack. Plan now to stay ahead of the dynamic changes reshaping the data center.

[Read more about lans & wans](#) in Network World's LANs & WANs section.

All contents copyright 1995-2011 Network World, Inc. <http://www.networkworld.com>