Data Center Networking – Designing Today's Data Center

There is nothing more important than our customers.

Data Center Networking – Designing Today's Data Center

Executive Summary

Demand for application availability has changed how applications are hosted in today's data center. Evolutionary changes have occurred throughout the various components of the data center, including server and storage virtualization and also network virtualization. The focus of today's enterprise is to increase business agility and the data center is a key asset that is gaining a lot of attention.

One of the most evolutionary changes to present itself is server virtualization. Motivation for server virtualization began with the opportunity to realize cost reductions in infrastructure and utilities, followed by an increase in redundancy and recovery. Virtualization benefits have evolved to include scalability and elasticity.

To realize the maximum benefits of virtualization, the rest of the data center has to evolve too. Data center LAN technologies have taken a similar path; one of providing redundancy and then creating a more scalable fabric within and between data centers.

There are three drivers to consider for next generation data center networks:

- Support for virtualization initiatives
- Reducing the number of tiers to improve performance
- Support for SAN convergence

In today's highly virtualized and dynamic IT infrastructure, organizations are continuously challenged to provide maximum scale and performance along with cost effective and resilient infrastructures in the data center. Virtualization has dramatically changed the requirements. Yesterday's highly segmented data center networks do not support key virtualization benefits such as dynamic virtual machine provisioning (VMotion/XenMotion). Flattening the network solves the initial problem, but introduces other design challenges. Virtualization's key benefits to an organization, namely reducing capital and operational expenses and improving resiliency, are firmly dependent upon next generation data center architectures.

Virtualization's promise of reduced capital investment and lowered operational expenses is also tied to reducing the number of tiers in the data center. Reducing the number of tiers in the data center not only reduces equipment (CAPEX/OPEX), but it also increases application performance by reducing latency. As enterprises are increasing the available bandwidth, they are also leveraging the opportunity to reduce the complexity of the topology. While reducing the number of devices decreases complexity, the third driver – SAN convergence – introduces new challenges and considerations for the next generation data center network.

There are many debates on the benefits of SAN convergence. SAN convergence is a highly debated topic, and the standards are newly ratified or in the process of ratification depending on your choice of SAN. In general, the primary driver of SAN convergence is infrastructure consolidation – data and storage sharing the same infrastructure with a common interface on the server. iSCSI was the industry's first converged SAN technology, and Fibre Channel over Ethernet (FCoE) is the new technology in the hype cycle. Another paper in this series will examine the challenges associated with each and provide information on solutions offered by Enterasys today and a vision of SAN/LAN convergence.

This paper, the second of a three part series, will establish guidelines on how to design, build and operate a data center infrastructure. We will focus on a data center network that addresses the requirements highlighted above and provide an overview of the differentiated solutions provided by Enterasys. The third paper will provided information specific to solutions from Enterasys in the area of SAN convergence.

Data Center Evolution

The evolution of the data center is in full swing. Advances have been made in several technological components of the data center and more are set to be ratified in the coming years. Businesses that are re-engineering their networks have many new options to consider. Each technology area is an improvement and changes the way we think about building networks.

Businesses require next generation networks to evolve their architecture from redundancy to resiliency. Redundant architectures provide insurance for failures, but also add cost without providing daily operational benefit. Businesses want to get the best performance from all assets and expect the new data center network to automatically adapt and provision services even when there are failures – in effect, to become resilient.

Resilient network architectures also help the business meet additional requirements, including:

- Improving application performance
- Enabling regulatory compliance
- Providing business (IT) agility
- Lowering CAPEX and OPEX

Advances in network products have benefited from Moore's Law: higher port densities, faster fabrics, smaller footprints and lower power consumption. Leveraging these advances helps improve system performance and reduce the hardware footprint. Virtualization further reduces the physical hardware required to support an organization's applications. Combining the two enables a new trend in data center networks – flattening the network. Reducing the footprint and total number of systems not only makes the network perform better, it contributes to its resiliency.

The network is simpler, but delivering the resilient network requires consideration around two key areas:

1. Virtualization of the data center edge switch
2. Implementing a meshed network topology

Each area alone will improve the operation of your data center. Implementing both will yield a benefit greater than the sum of the parts in terms of operational efficiency, capital investment required and satisfied end users.

Virtual Switching Systems

Virtual switching represents an evolution in data center switching that provides data center architects with a new set of tools to improve application availability and response time and simplify the edge network topology. Virtual switching is gaining acceptance in data centers as it offers resiliency for server interconnects that previously required manual configuration in the servers. Today's implementation of virtual switching allows servers to view two physical switches as a single unit, solving the following challenges:

- Automating link aggregation across physical switches and servers
- Meshing L2 network uplinks to data center aggregation/core switches
- Enabling non-stop forwarding of application traffic in event of a single device failure

Enterasys virtual switching merges physical switches into a single logical switch to increase bandwidth and creates an active mesh between servers and switches in the data center. This enables the real-time delivery of applications and services and simplifies the management of the network infrastructure. In a Top-of-Rack scenario (see Figure 1: Enterasys Virtual Switching) this is represented by two physical switches. In an End-of-Row scenario the Enterasys S-Series chassis system implements a virtualized chassis system first realized in the Enterasys N-Series. Enterasys is an innovator in network fabrics with many industry patents in the field, having delivered the industry's first meshed network in 1996.

Enterasys virtual switching solves the aforementioned challenges and provides:

- Automated "host-specific" network/security profiles per virtual host, per port
- Maximum availability and failure tolerance with seamless failover capabilities
- Established technology with more than 3 million switch and router ports deployed
- Proven Enterasys OS code base, now and for the future

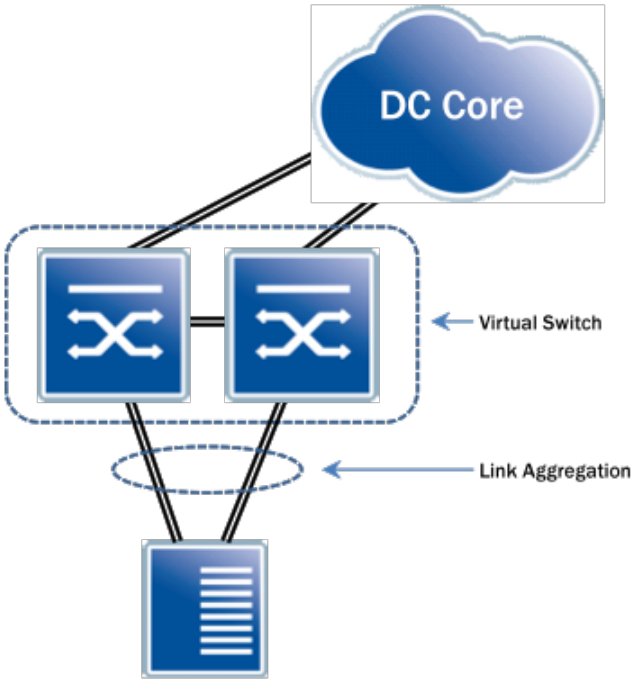

Figure 1: Enterasys Virtual Switching

Meshed Networking

Resiliency is required across the entire network, not just a particular layer if we are to meet the needs of the applications driving the business. In the past, networks were designed with active and passive links. While this provided redundancy, changes in the network topology would create service outages while the network settled on a new logical configuration. Technologies have evolved and many networks today segment their logical topologies using standards such as IEEE 802.1Q-2005 Multiple Spanning Tree Protocol (MSTP) to enable multiple topologies to best use all possible links.

While MSTP allowed for all links to be leveraged, not all links are leveraged equally. This is because the segmentation still has active/redundant links within each VLAN grouping. Next generation networks need to support an active/active configuration which can:

- Contain failures so only directly affected traffic is impacted during restoration
- Enable rapid restoration of broadcast and multicast connectivity
- Leverage all of the available physical connectivity, with no lost bandwidth
- Enable fast restoration of connectivity after failure

There are two competing standards in ratification that will increase the resiliency of tomorrow's data center LAN:

- Shortest Path Bridging (SPB) – IEEE 802.1aq work group
- Transparent Interconnect of Lots of Links (TRILL) – IETF TRILL work group

Enterasys is committed to open standards, and these proposed protocols show promise for delivering a more reliable and interoperable data center. SPB, due to its interoperability with RSTP/MSTP and standardization by the IEEE, has inherent advantages and is expected to be ratified in 2011, with fully interoperable implementations available in early 2012.

Each of these standards is aimed at simplifying the network topology and providing an active mesh between the edge and core of data center networks. The two standards are not far apart, but there are sufficient differences between them that make them non-interoperable. Pre-standard implementations have not been able to demonstrate interoperability. At this time, Enterasys recommends that customers wait for the IEEE and IETF to work through differences and the standards bodies to consolidate and focus on a single standard. While that occurs, there are many benefits to leveraging existing technology while planning for future standards to ratify.

Today Enterasys recommends a Layer 2 meshed design that derives its resiliency by leveraging industry standard protocols including IEEE 802.1Q-2005 (MSTP) and IEEE 802.1D-2004 (RSTP). These protocols not only offer interoperability with existing data network equipment, but their maturity also provides for a rich set of tools and a trained workforce that can implement and maintain the network topology. Interoperability and maturity provide for greater uptime and lower OPEX when compared to proprietary solutions.

It is recommended for customers building new data center networks that they build the layer 2 core with standard RSTP/MSTP protocols today that will enable an easy, non-disruptive migration towards new standards. Customers should also verify that the vendor hardware they are purchasing today will support emerging topology standards.

Deployment Scenarios

Enterasys is a leader in data network consolidation and has enabled many customers to consolidate their data centers to a two-tier network design. The Enterasys S-Series provides the ability to compress the traditional three-tier network into a physical two-tier network by virtualizing the routing and switching functions within a single tier. Virtualized routing provides for greater resiliency and fewer switches dedicated to just connecting switches. Reducing the number of uplinks (switch hops) in the data center improves application performance, reduces CAPEX and reduces mean-time-to-repair (MTTR).

We realize that no two data centers are the same, and that data center managers have preferences in how they physically deploy their switches. Enterasys provides flexible solutions to meet customer needs no matter their preference.

Access Layer (Top of Rack Option)

Enterasys recommends a virtual switching option for the Top of Rack network design option. Virtual switching unifies two physical switches as a single logical switch to provide resilient network connectivity to the servers. Resilient connectivity from server to virtual switch is attained by:

- Automating link aggregation from each server across physical switches (IEEE 802.1ad)
- Automating host-specific network and security profiles per host (including virtual hosts) per port

Access Layer (End of Row Option)

The best practice option that Enterasys recommends is an End of Row chassis-based switch for server connectivity. The Enterasys S-Series has all the advantages of Top of Rack virtual switching without requiring an independent chassis. The S-Series chassis implements a distributed switching architecture without a supervisor engine. In essence, the S-Series chassis is a virtual switch cluster with fully redundant switching and power systems. For example, an S8 positioned in this solution includes the same protocols as the Top of Rack solution above, as well as providing the same benefits, but in a flexible platform that allows networks to scale for future connectivity requirements.

Core

Enterasys recommends a Layer 2 meshed network today, leveraging multiple connections from the access layer and the core devices. IEEE 802.1Q-2005 and IEEE 802.1D-2004 provide the ability to distribute traffic across the entire mesh. This results in maximum interoperability with other devices while building the infrastructure towards a future platform supporting a full meshed design when standards ratify.

Enterasys does provide the option for a Layer 3 meshed network design, to accommodate design principles or to support network expansion or segmentation. In this scenario, Enterasys would recommend the use of IETF OSPF-ECMP to provide a resilient Layer 3 mesh. Leveraging ECMP (Equal Cost Multi-Path) allows for distribution of traffic across the available links in the mesh and, as with the Layer 2 protocols, provides for autonomous recovery from link failures. Enterasys would recommend use of IETF VRRP to provide gateway resiliency.

Consolidation of vSwitch with Physical Network

The virtual network represents a new challenge to organizations that have implemented virtualization solutions. The virtual network lacks visibility and control into traditional network applications, leaving network operators at a disadvantage in meeting compliance objectives or SLAs.

Delivering network services in real-time in a virtualized environment, Enterasys Data Center Manager (DCM) integrates with a traditional network management tool, Enterasys Network Management Suite (NMS). DCM bridges the divide of managing and provisioning virtual network and physical networks to support data center applications. Enterasys DCM is a powerful unified management solution that delivers visibility, control and automation across the entire data center fabric.

Enterasys DCM is unique in that it requires no special software or applications to be loaded onto the hypervisor of the customer's choice, whether it is Citrix XenServer, Microsoft HyperV or VMware ESX. The solution interfaces directly with the hypervisor systems or their management applications via their native APIs. This provides server and VM visibility and control with no bias to the server or operating system vendor. This allows enterprises the freedom to choose the server vendor that best fits their requirements, not locking them into a single vendor solution.

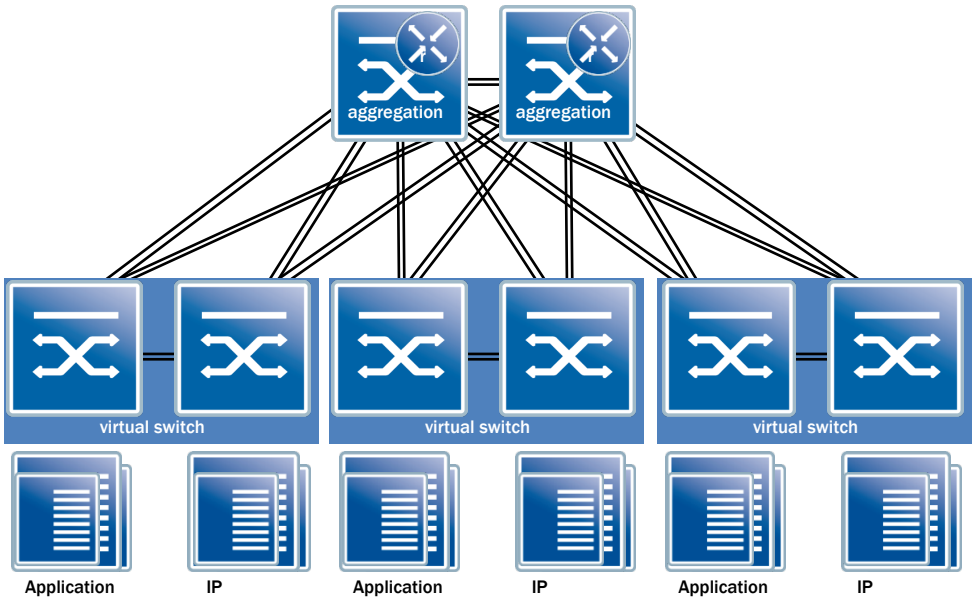

Figure 2: Enterasys Reference Design

The Enterasys reference design delivers a resilient data center LAN environment that includes:

- Automated link aggregation across physical switches (virtual switching)
- Meshed L2 network uplink to data center aggregation/core switches
- Automated "host-specific" network/security profiles per virtual host, per port
- Support for thousands of virtual hosts per edge system

The premier Enterasys product for the data center – the S-Series modular switching line – can be deployed in the data center, core and distribution layer of the network. This capability dramatically reduces the overall cost to manage and maintain a network infrastructure. One product platform, one operating system and a firm commitment to standards delivers a lower total cost of ownership.

Conclusion

Data center LANs are constantly evolving. What worked yesterday has a limited lifetime in today's environment. Business pressures are forcing IT to adopt new application delivery models. Edge computing models are transitioning from applications at the edge to virtualized desktops in the data center. The evolution of the data center to a private cloud network is well underway.

Enterasys offers a flexible product portfolio to fit how you design your data center, versus fitting your data center into someone else's design. Our product portfolio is focused on standards-based technologies, with a long life cycle that will incorporate new standards. Our data center products are focused on:

- Support for virtualization initiatives
- Support for SAN convergence
- Virtual switching
- Meshed networking

Enterasys delivers a simplified data center LAN that improves application performance and increases business agility, providing customers with a future-proofed approach to data center design best practices. To learn more, visit http://www.enterasys.com/solutions/DataCenter.aspx.

Contact Us

For more information, call Enterasys Networks toll free at 1-877-801-7082, or +1-978-684-1000 and visit us on the Web at enterasys.com

Thought Leadership
Patented Innovation

01/11

Delivering on our promises. On-time. On-budget.