Enterasys Design Center Networking – Connectivity and Topology Design Guide

Demand for application availability has changed how applications are hosted in today's datacenter. Evolutionary changes have occurred throughout the various elements of the data center, starting with server and storage virtualization and also network virtualization.

Motivations for server virtualization were initially associated with massive cost reduction and redundancy but have now evolved to focus on greater scalability and agility within the data center. Data center focused LAN technologies have taken a similar path; with a goal of redundancy and then to create a more scalable fabric within and between data centers.

Business requires next generation networks to change focus from redundancy to resiliency.

While it may seem that redundancy and resiliency are one and the same, they are not. Redundancy simply requires duplication of systems. Resiliency is the ability of the solution to "adapt" to the consequences of failure. Today's data center must meet a number of business requirements and overcome several design obstacles in order to truly achieve resiliency.

Business requirements

- Improve application performance
- Regulatory compliance
- Business (IT) agility

Design obstacles

- Density increases with a rapid pace
    - On an ongoing basis new applications are deployed on new server systems
    - Increases in server performance results in a large number of virtual machines per server
    - Increases in the number of virtual machines per server increases the traffic per server
- Dynamic application provisioning and resource allocation

Resiliency is not achieved by simply implementing new technologies. It also requires investment in architectures and tools along with a ready workforce that can operate these networks without requiring extensive vendor-specific training.

This paper will provide the reader with key concepts for designing a standards-based data center fabric to meet the requirements of today and tomorrow.

Table of Contents

Data Center Network Design Goals

Derived from the business objectives and the requirements of the applications hosted today in the data center the common design goals include:

- Performance
- Scalability and agility
- Flexibility to support various services
- Security
- Redundancy/High availability
- Manageability
- Lower OPEX and CAPEX
- Long term viability

There is no single solution that can be applied to all. What we will propose is a set of guidelines from which a solution can be designed which will meet the unique needs and goals of the organization. Additionally, the design architecture will emphasize criteria which are standard-based without compromising critical functionality.

Data center LANs are constantly evolving. Business pressures are forcing IT organizations 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 from centralized servers to a private cloud is well underway and will be augmented by hybrid and public cloud computing services.

With data center traffic becoming less client-server and more server-server centric, new data center topologies are emerging.

With data center traffic becoming less client-server and more server-server centric, new data center topologies are emerging. Yesterday's heavily segmented data center is becoming less physically segmented and more virtually segmented. Virtual segmentation allows for the reduction of physical equipment, leading to both capital and operational expense (CAPEX/OPEX) savings.

New Enterasys connectivity solutions provide the ability to compress the traditional 3-tier network into a physical 2-tier network by virtualizing the routing and switching functions into 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 as it reduces latency throughout the fabric.

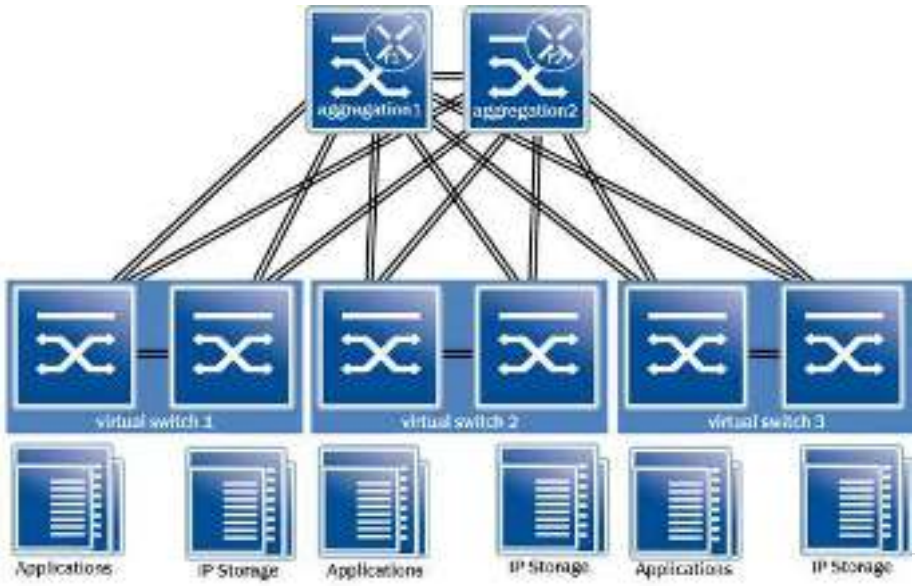

Figure 1: Two-tier Data Center Design

The design above shows virtual switches used at the data center LAN access layer providing connectivity for both applications and IP storage – iSCSI or NFS attached. The virtual switches leverage a Layer 2 meshed network for interconnectivity. The aggregation and core are merged into a single layer by virtualizing the router function in the data center LAN switch.

In addition to the transport layer, Enterasys provides an industry-leading solution for centralized command and control of the infrastructure. The Enterasys Network Management Suite (NMS) products, including Data Center Manager, simplify data center LAN management by enabling the deployment of a consistent configuration throughout the data center (both physical and virtual) and enterprise LAN. Management integration with industry-leading virtualization vendors provides multi-vendor hypervisor support that orchestrates virtual server/desktop operations with the virtual and physical networks, ultimately providing flexibility for the customer.

Data Center Connectivity Trends

10G, 40G, 100G

As the 40G/100G Ethernet standard (IEEE 802.3ba) was ratified in June 2010, the biggest market for 40G Ethernet is projected to be within the data center and for data center interconnects. Early adoption of 100G Ethernet will be used in a few bandwidth hotspots in carrier core networks and for network aggregation of 10G and 40G Ethernet links. Even with the ratification of the new Ethernet standard, there are a number of reasons 10G Ethernet is still growing and will continue to have significant growth for at least another 5 years:

- The cost of the technology is still high (as of 2013). It could take at least two more years before 100GE prices will be closer to that of 10x10GE.
- 40G/100G Ethernet is still new and it will take time until the technology is widely available. This is especially true when deploying the technology in a data center and core network, which is always planned to grow with a certain amount of multi-vendor equipment.

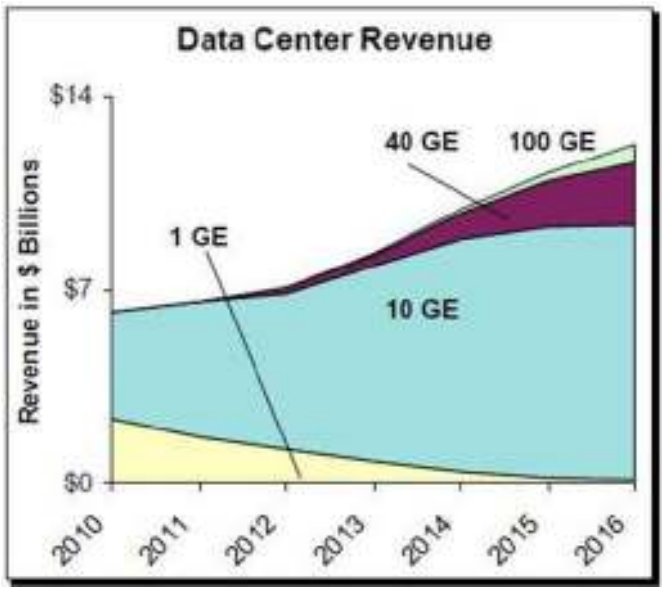

Figure 2 "source: Delloro"

The decision to implement a particular technology depends upon organizational needs, budget, and projected company business growth. However, the selected network device should at least incorporate the latest design architecture, a solid plan/road map, and enough capacity to support 40G/100GE in the future.

Storage I/O Consolidation

Within the industry, there is a stated long-term goal to establish Ethernet as the transport for a "converged" data and storage solution, thereby reducing TCO within a converged Ethernet data center. Storage connectivity today is a mix of Fibre Channel (FC), iSCSI and NFS (both iSCSI and NFS are already Ethernet and IP based). If FC is deployed, it requires two different sets of hardware, cables, tools and skill sets. Storage connectivity in the future will be based on a single converged network with an intermediate step of a single converged interface from the server to the access switch, all with new protocols and hardware. This will results in fewer adapters, cables, and nodes, resulting in more efficient network operations.

"Gartner research shows that a converged Data Center network requires more switches and ports, is more complex to manage and consumes more power and cooling than two well-designed separate networks."

The Enterasys solution is able to co-exist with FC environments, enabling the organization to continue to leverage existing investments. Enterasys provides support for Data Center Bridging (DCB) in multiple phases with different hardware and software requirements as the underlying technology to transport Fibre Channel over Ethernet (FCoE). Industry analyst firm Gartner has published a report regarding some of the myths of FCoE technology. In announcing the report ("Myth: A Single FCoE Data Center Network = Fewer Ports, Less Complexity and Lower Costs" ID Number: G00174456), Gartner notes the traditional architecture of separate storage and network systems still has merit:

As an alternative, Enterasys offers a simple, yet highly effective approach to enable, optimize and secure iSCSI SAN or NFS NAS deployments. The Enterasys S-Series modular switch is a key component of our overall solution, delivering an easy and effective way to optimize communications through automatic discovery, classification, and prioritization of SANs. In addition, the Enterasys solution will identify and automatically respond to security threats against virtual storage nodes, enforce role-based network access control policies, and comply with regulations for monitoring and auditing.

The IEEE Data Center Bridging task group, a working group of IEEE 802.1 working group, is focused on defining a new set of standards which will enable Ethernet to effectively deliver data center transport for both server and storage traffic. Terms commonly associated with DCB are "Data Center Ethernet", also known as DCE, and Convergence Enhanced Ethernet (CEE). It should be understood that DCB is the task group and term commonly being used to describe tomorrow's Data Center LANs.

Data Center Bridging is focused primarily on three (3) IEEE specifications:

- IEEE 802.1Qaz – ETS & DCBX – bandwidth allocation to major traffic classes (Priority Groups); plus DCB management protocol
- IEEE 802.1Qbb – Priority PAUSE. Selectively PAUSE traffic on link by Priority Group
- IEEE 802.1Qau – Dynamic Congestion Notification

In addition to these protocols people often include layer 2 meshing technologies when they refer to DCE or CEE.

Right now, FCoE only addresses the first five feet of connectivity in the data center, the five feet from the server to the network access switch. The transformation to a converged data and storage environment is no small challenge and will continue well into 2015 and beyond.

Main Components of the Data Center

This main focus of this paper is on data center network infrastructure design; however, we will briefly cover some of the other data center components. A data center is a facility used to house computer systems and associated components, such as telecommunications and storage systems. It generally includes redundant power supplies, data communications connections, environmental controls (e.g., air conditioning, fire suppression, etc.) and security devices. For our purposes we will focus on the servers, storage and connectivity elements of the data center.

Servers

Servers deployed in the data center today are either full featured and equipped rack-mount servers or blade servers. A blade server is a stripped down server with a modular design optimized to minimize the use of physical space and energy. Whereas a standard server can function with (at least) a power cord and network cable, blade servers have many components removed to save space, minimize power consumption and other considerations, while still having all the functional components to be considered a computer. A blade enclosure, which can hold multiple blade servers, provides services such as power, cooling, networking, various interconnects and management. Together, blades and the blade enclosure form the blade system.

There are pros and cons for each server type. This discussion is not a focus of this paper.

Virtualization has introduced the ability to create dynamic data centers and with the added benefit of "green IT." Server virtualization can provide better reliability and higher availability in the event of hardware failure. Server virtualization also allows higher utilization of hardware resources while improving administration by having a single management interface for all virtual servers.

Storage

Storage requirements vary by server type. Application servers require much less storage than database servers. There are several storage options – Direct Attached Storage (DAS), Network