



TECHNOLOGY STRATEGY BRIEF

Enterasys Data Center Fabric

Enterasys Data Center Fabric

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 drawing a lot of attention.

One of the most revolutionary changes 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 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, data center organizations are continuously challenged to provide maximum scale and performance along with cost effective and resilient infrastructures. 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, 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 expense 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 their bandwidth use, 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.

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; Fibre Channel over Ethernet (FCoE) is the new technology in the hype cycle.

This paper is focused on the Enterasys Data Center Fabric Architecture that addresses the needs of next generation data center networks.

Benefits

Optimized for seamless application availability

- Guaranteed bandwidth for applications
- Application awareness for granular visibility and control in the fabric
- Resilient fabric core and server access technology
- Active fabric that maximizes performance and minimizes latency

Automation & visibility

- Automatic provisioning of physical and virtual switch fabric
- Virtual machine VM tracking throughout the fabric
- Automatic restoration of fabric links

Better service level

- Guaranteed bandwidth and minimum delay for sensitive applications

Reduced TCO

- Lower CAPEX & OPEX with a converged fabric supporting virtualization
- High degree of automation minimizes OPEX
- Standards-based: choice of server, storage, virtualization technologies
- Longer lifecycles and new standard support with flexible CoreFlow2 technology

The Enterasys Data Center Fabric Architecture

Using the Enterasys Data Center Fabric Architecture customers can smoothly migrate today's data center network to a fully converged fabric that addresses their needs around virtualization, improved performance and SAN convergence.

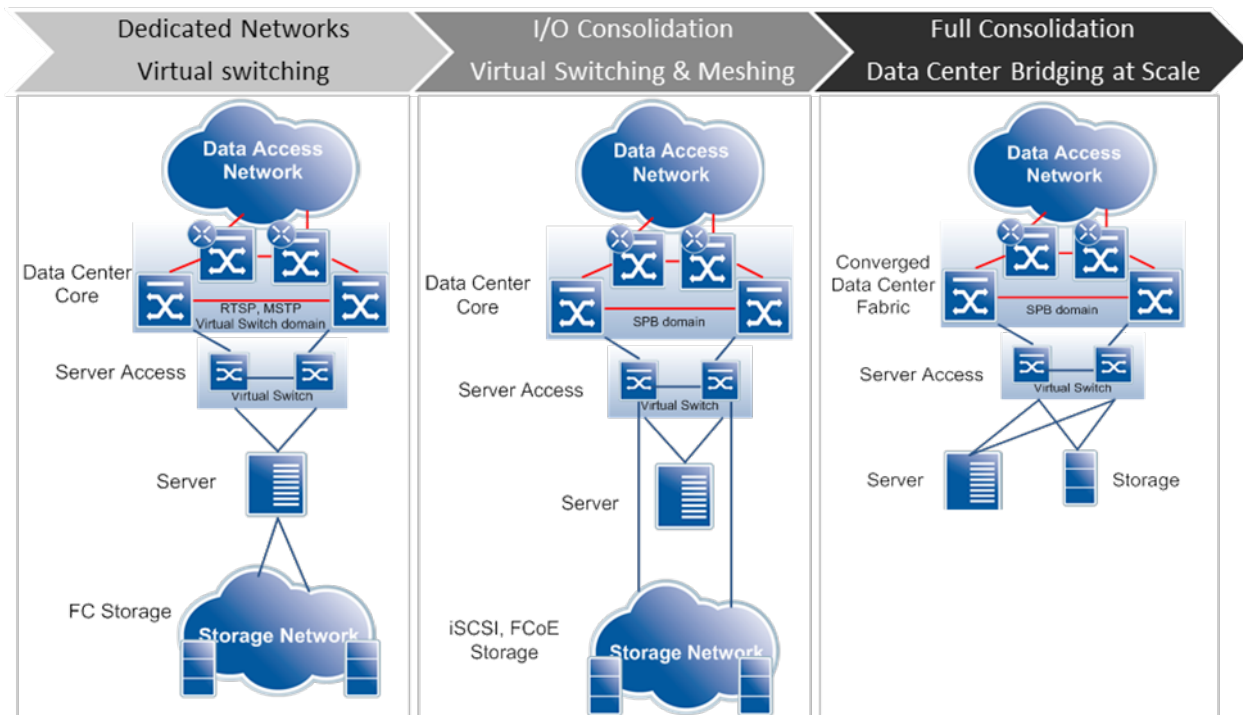


Figure 1 - Data Center Fabric over time

The main technology components of the Enterasys architecture include:

Virtualization Awareness

- Addressing the need for visibility and automation when virtual servers get (re-)deployed

Data Center Bridging

- Effectively supporting I/O and SAN convergence in the data center fabric

Virtual Switching

- Increasing the available bandwidth and enabling a resilient link to servers and blade center switches

Fabric Core Mesh

- Scaling the data center fabric core to a larger aggregated capacity with lower latency

Application Awareness

- Enabling application visibility and control in the data center fabric

Virtualization Awareness

Virtualization is one of the most revolutionary changes to the data center in the past decade. Server and storage virtualization enable rapid changes on the services layer, but the dynamic nature of virtualization places requirements on the data center network. “Motion” technologies create rapid configuration changes on the network layer as servers/VMs are added or moved amongst physical machines.

To deliver network services in real-time within a virtualized environment, [Enterasys Data Center Manager](#) integrates with the [Enterasys Network Management Suite](#) (NMS) to bridge the divide between virtual machine and network provisioning applications. Enterasys DCM is a powerful unified management solution that delivers visibility, control and automation over the whole data center fabric, including network infrastructure, servers, storage systems and applications, across both physical and virtual environments.

Enterasys DCM requires no special software or applications loaded onto hypervisors or virtual machines. The solution interfaces directly with the native hypervisor and hypervisor management systems. Server and VM visibility and control are provided with no bias to the server or operating system vendor. Enterprises have the freedom to choose the server and hypervisor vendor that best fits their requirements, not the vendor that will lock them into a one shop solution. DCM is unique in the industry in supporting all major virtualization platforms, including Citrix XENServer and XENDesktop, Microsoft Hyper-V and VMware vSphere, ESX, vCenter and VMware View.

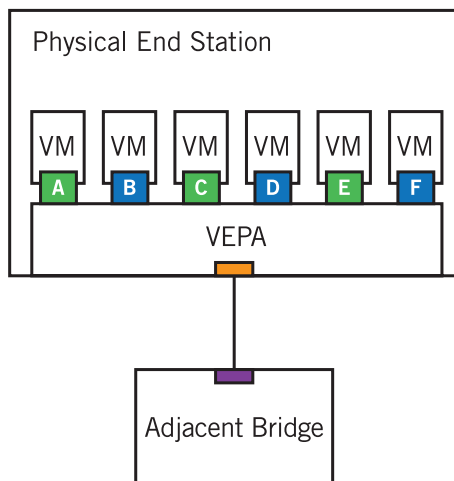


Figure 2 – VEPA packet flow

Enterasys DCM integrates with existing workflow and lifecycle tools to provide cradle-to-grave visibility into virtualized and physical assets and to automate the physical and virtual network configurations for virtual machines. Instead of requiring new software installed on the hypervisor, Enterasys DCM leverages each vendor's APIs and Enterasys published APIs to provide automated inventory discovery and control over the hypervisor switch configuration, as well as management of the physical network configuration.

Enterasys expects that the hypervisor switch (vSwitch) as a dedicated software component will disappear over time and be replaced by standard based mechanism like IEEE 802.1Qbg – Virtual Ethernet Port Aggregator (VEPA).

Implemented on the NIC card of a server and accelerated by mechanisms like SR-IOV Single Root I/O Virtualization to offload the server CPU, VEPA enables the physical switch to make the forwarding decision. DCM will then enable the management of these solutions in a holistic manner.

Enterasys plans to implement VEPA on key data center platforms with software upgrades leveraging [CoreFlow2](#) technology, which enables greater visibility and control into core business application performance on the network.

Data Center Bridging

The long-term goal is to reduce TCO by establishing Ethernet as the transport for a “converged” data and storage solution. Storage connectivity in the future will be based on a single converged network – with new protocols and hardware. Enterasys offers a simple, yet highly effective approach to enable, optimize and secure iSCSI SAN or NFS NAS deployments today. Enterasys delivers today an easy and effective way to optimize communications through automatic discovery, classification, and prioritization of IP SAN traffic.

The IEEE Data Center Bridging (DCB) task group, a working group of the IEEE 802.1 work group, is focused on a set of standards that are working to make Ethernet a more viable data center transport for both server and storage traffic, especially when it comes to FCoE. DCB creates a more reliable network based on Ethernet technology that transitions from best effort to “lossless” operation and provides congestion management at Layer 2 more effectively than traditional TCP-based congestion management and flow control mechanisms. While traditional storage protocols like iSCSI and NFS will also benefit from DCB, they don't rely on it. FCoE mandates a lossless operation which in a multi-hop switch environment is only possible with the use of DCB.

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

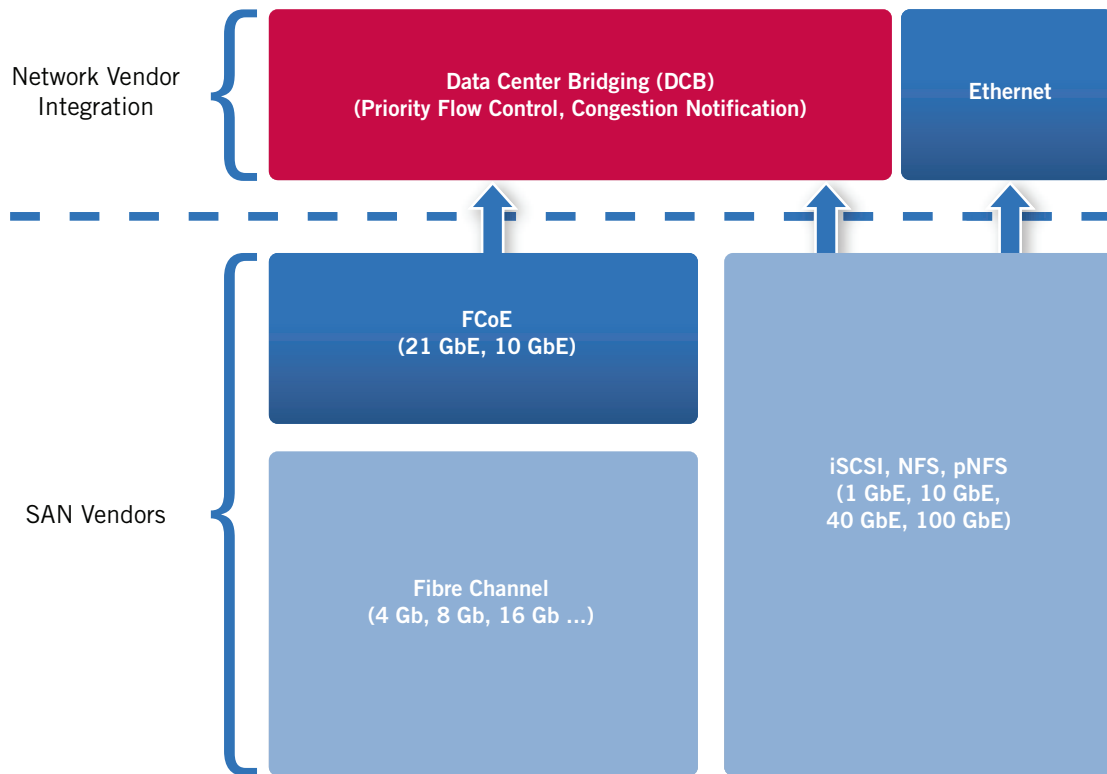


Figure 3 - DCB and various storage technologies

Enterasys plans to support this standard in two phases on key data center platforms – with both software and hardware upgrades.

One can expect that SAN convergence will occur in at least two phases. The first phase will consolidate I/O from the server with the Ethernet data fabric. The second phase will converge the SAN across the whole or across areas of the fabric.

The first phase primarily reduces the cost for the server as a dedicated HBA is not required anymore. This saves on server energy costs as well as real estate. In this phase the reduction of cabling to the server simplifies operations and saves on costs as well with fewer switch ports required. The second phase reduces the number of network devices required for the whole fabric. The second phase requires mature standards and will take several years before it will become adopted mainstream in a data center.

While IP SAN convergence is readily available today, complete FCoE SAN convergence on networks leveraging DCB will not be realized for several years.

Virtual Switching

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, while simplifying 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 for top of rack (ToR) switch systems allows servers to view two physical switches as a single system, 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




Figure 4 - Resilient server connection

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. The Enterasys [S-Series](#) chassis system implements a virtualized chassis system that was first realized in the Enterasys N-Series.

Beyond this solution, virtual switch bonding across multiple chassis will be available on key data center platforms like the S-Series as a software option to interconnect chassis over traditional 10G and future 40G/100G Ethernet links.

Enterasys virtual switch bonding 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

Fabric Core Mesh

Resiliency and performance is required across the entire network, not just a particular layer if we are to meet the needs of the applications driving the business. New data center technologies like server virtualization and FCoE require larger scale layer 2 “flat” networks. The any to any communication pattern of presentation, application and database servers, often deployed in a scale out design, requires a non-blocking, high performance and low latency network infrastructure.

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. This is also the [best practice recommendation](#) for today's data center networks.

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:

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

There are two competing standards in the process of 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

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.

Enterasys is an innovator in network fabrics with many industry patents in the field, having delivered the industry's first layer 2 meshed Ethernet network in 1996: an active mesh on the basis of an intelligent routing protocol. Back then, it was known as SecureFast and OSPF was used as the VLSP (VLAN Link State Protocol) in order to exchange MAC address reachability.

IETF TRILL and IEEE SPB both use IS-IS as the routing protocol to achieve similar goals. The IEEE has committed itself to supporting all existing and new IEEE standards (particularly the IEEE Data Center Bridging protocols, but also existing management protocols, Ethernet IEEE 802.1ag (OAM), etc.) via IEEE SPB.

IEEE SPB uses a header from the Provider Backbone Bridging Standard called MAC in MAC encapsulation (IEEE 802.1ah) resulting in the SPB-M implementation. TRILL allows several paths (Equal Cost Multipathing) and also uses a different path for Unicast and Broad/Multicast.

Upon closer inspection, this is a problem for the same reason that some of the other IEEE protocols like Data Center Bridging have a problem with TRILL. They require the same paths between source and target. And, of course, various paths (with different latency) can also lead to out of order packet delivery, e.g. when Unknown Unicast Flooding is changed to Unicast. Using IEEE SPB a single active path is used for all traffic between a given source and target.

The Enterasys Data Center Fabric will initially use IEEE SPB, which is expected to be ratified in 2011, with interoperable implementations available in 2012. SPB will be supported by a software upgrade on key data center platforms that implement CoreFlow2 technology.

SPB will build upon and is fully interoperable with the existing data center layer 2 LANs running MSTP and it will improve the resiliency of tomorrow's networks. Due to its nature as an IEEE standard it allows for the migration from existing infrastructures with little to no disruptions. The following benefits can be achieved with SPB:

- **Plug and Play**
 - Minimal to no configuration necessary to create an active mesh
- **Reduced number of hops**
 - With all links active in the fabric the traffic always takes the shortest path
 - Reduced latency between applications
- **Higher aggregated capacity**
 - No unused (blocked) links results in a higher aggregated fabric capacity
- **Scalability**
 - 100's to 1000's of switches within a single domain are possible
- **Resiliency**
 - fast restoration of connectivity after failure
 - under failure, only the directly affected traffic is impacted during restoration; all unaffected traffic continues along uninterrupted
 - rapid restoration for real-time traffic such as IP Voice and Video

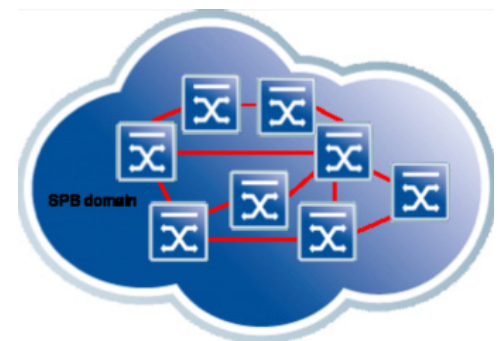


Figure 5 - Shortest Path Bridging Domain

For larger data center fabrics with complex topologies spanning multiple sites SPB will be a key element to fully leverage the benefits of virtualization and convergence. Smaller networks might not need this function. Due to increased server and switch performance, the number of nodes in the data center will decrease dramatically, so a lot of mid-size enterprises will no longer require complex topologies.

Application Awareness

Most network infrastructure components that are deployed today do not provide data to control or monitor applications. Networks are typically implemented in such a way that all services and all applications are given equal priority – or with only very basic prioritization schemes. With the increasing use of virtualization, SOA architectures, cloud computing and continued network convergence, this typical scenario cannot meet today's requirements. This affects both access networks and data center fabrics. In each area of the network it becomes more critical to properly identify applications to ensure availability by monitoring and enforcing controls. This goal cannot be achieved just by examining traffic at the transport layer.

CoreFlow2 enables IT administrators with greater visibility into critical business applications and, with this instrumentation, the ability to enable better controls to meet the SLAs the business demands. Some use cases enabled by CoreFlow2 include*:

- **SAN**
 - Enable access control for iSCSI targets with granularity to the initiator
 - Bandwidth usage monitoring per iSCSI target
- **IP Voice & Video**
 - Enable QoS and access control for RTP media streams and control data
- **Cloud**
 - Enable role-based access controls for cloud services such as www.salesforce.com
 - Bandwidth monitoring for specific sites such as www.youtube.com

** Implementation details per product category are subject to the development roadmap. Please refer to the product datasheet and release notes for details*

Expanded visibility will soon be realized in a subsequent release of Enterasys NMS. NMS will aggregate the native unsampled NetFlow records generated by CoreFlow2-powered devices to provide application-level visibility across the network. Select Enterasys products will also support application response time measurement probes distributed throughout the network infrastructure. This will further enable the IT administrator to monitor application response times in their network to meet SLAs, deliver higher application availability and enable more efficient troubleshooting.

Summary

Data center LANs are constantly evolving. What worked yesterday isn't working today and will be antiquated tomorrow. Business pressures are forcing IT to adopt new application delivery models. Edge computing models are transitioning from applications at the edge to now moving to virtualized desktops in the data center. The evolution of the data center to private, hybrid cloud services as well public cloud integration is well underway.

IP SAN convergence is here today, but FCoE convergence will lag as key technologies such as DCB are still not ready. Virtual switching will offer incremental resiliency in data centers, and meshed technologies will extend the resilience throughout the entire fabric. IEEE SPB is the near term viable candidate for data center mesh and will be a key aspect of the Enterasys data center portfolio.

Supporting an open standards approach, Enterasys delivers a simplified data center fabric solution today that improves application performance and increases business agility, providing customers with a future-proofed approach to data center architectures.

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



© 2011 Enterasys Networks, Inc. All rights reserved. Enterasys Networks reserves the right to change specifications without notice. Please contact your representative to confirm current specifications. Please visit <http://www.enterasys.com/company/trademarks.aspx> for trademark information.

