



ETHERNET FABRICS 101

HANDBOOK

ABSTRACT

This unique handbook will help you better understand Ethernet fabrics and best practices for deploying them in your environment. It also sheds light on the business and technology trends behind Ethernet fabrics, and includes an assessment to gauge your network's readiness, multiple business use cases, as well as a glossary of key terms, technologies, and standards. Ethernet Fabrics 101 is a great resource for data center architects and managers ready to learn about the next stage of virtualization.

CONTENTS

Get to Know Ethernet Fabrics

Start preparing your networks for Ethernet fabrics.4

What's Happening in Your Data Center?

Four significant trends are the driving force behind Ethernet fabrics.

Learn what they are and their impact on your data center.5

A Short History of Data Center LANs and Predictions for the Future

What does your data center look like today, and how is it likely to look in

five years? Think flat networks with at least some services from the cloud.7

Ethernet Fabrics 101

Understand what Ethernet fabrics are and how they can impact your

business and your IT department.11

Ethernet Fabric Assessment

Complete this short exercise to help gauge your network's current

virtualization maturity stage and uncover potential gaps as you

move toward an Ethernet fabric.15

Transitioning to Ethernet Fabrics

When and how do you transition to an Ethernet fabric? Four use cases show how you can evolve

portions of your current network architecture to an Ethernet fabric, and which applications and

services are good early candidates.18

A Dictionary for Ethernet Fabrics

View brief definitions of terms, technologies, and standards for the

Ethernet fabric-based data center.23

GET TO KNOW ETHERNET FABRICS

Start preparing your networks for Ethernet fabrics.

We all know the wise maxim, “Less is more,” and it is finally true of networks.

Ethernet fabrics decrease the number of hops to create a flatter, faster architecture that is more efficient while remaining highly scalable and resilient. This innovative network advancement takes the most prized qualities of Ethernet, adds much-needed intelligence, and delivers the services needed for today’s virtualized data center and changing business requirements. At last, organizations can realize the full potential of server virtualization, which has been a boon to optimizing server resources but a configuration and monitoring nightmare.

Ethernet fabric-based networks address your data center pain points by providing:

- Intelligence
- Reduced network complexity
- Simplified management
- Elasticity
- Improved performance and scalability

Networking vendors across the industry are talking about Ethernet fabrics, and some even have product offerings. The market is still young, though, and the new term has made many promises and sparked dozens of questions about the technology and architectural requirements.

This handbook cuts through the hype and addresses those questions while helping you understand what Ethernet fabrics are, how they work, and what standards enable them. The sections “Ethernet Fabrics 101” and “A Dictionary for Ethernet Fabrics,” for example, cover convergence, Transparent Interconnect of Lots of Links (TRILL), Data Center Bridging (DCB), Priority-based Flow Control (PFC), and other IEEE standards, as well as aging standards, such as Spanning Tree Protocol (STP).

In “A Short History of Data Center LANs and Predictions for the Future,” you can discover how Ethernet fabrics relate to traditional data center LAN architectures, and “Transitioning to Ethernet Fabrics” describes how to enable solutions for large-scale virtualization, cloud infrastructure, and large, flat data center networks. These sections show how Ethernet fabrics and enhanced IP storage networks allow for lossless Ethernet transport that improves performance and reliability for iSCSI, NAS, Fibre Channel over Ethernet (FCoE), and more.

Brocade recognizes that Ethernet fabrics can solve the increasing challenges of virtualization, while helping to unlock its benefits. For that reason, Brocade is working to simplify the transition to this

new network architecture. In “Ethernet Fabric Assessment,” you can take a short quiz developed in conjunction with Forrester Consulting to determine your data center’s virtualization maturity stage and readiness for Ethernet fabrics.

If you are curious about Ethernet fabrics and want to advance your data center’s transformation, this handbook can help you understand the technology behind Ethernet fabrics and the best practices for deploying them within your environment.

To access additional information, join the conversation about Ethernet fabrics, visit www.brocade.com/ethernetfabrics.

WHAT’S HAPPENING IN YOUR DATA CENTER

Four significant trends are the driving force behind Ethernet fabrics. Learn what they are and their impact on your data center.

At one time all data operations were handled from a server room and a switch closet. Now, even small and mid-sized organizations—not just large corporations—have data centers. And every data center, regardless the size, endures upgrades and overhauls because this area is easily the most dynamic within the entire organization.

Four significant trends are making the data center a hot zone. **More devices, a deluge of data**, the **decreasing costs to transfer data**, and **server virtualization** are triggering a transformation in the data center that will lead to multiple data center advancements during the next five years.

A sampling of analyst numbers reveals the true significance of these four trends. According to IMS Research, 22 billion devices will be connected to the Internet by 2020. These devices have an unprecedented ability to create and consume data. Alongside these devices creating data, organizations are digitizing and storing an incredible amount of raw data. IDC and EMC forecast that by 2020 the world will have 35 zettabytes of data on hand. Such astronomical growth is possible because the cost of transferring data has decreased while the speed to transfer it has continued to increase.

The desire to control costs and increase utilization is driving the fourth trend—server virtualization. According to a 2011 survey of Brocade customers, more than 90 percent will have implemented server virtualization in their data centers by 2013.

Signs of a Coming Network Redesign

The increasing use of server virtualization, which removes the hardware dependency that existed between applications and the underlying hardware, is causing data center architects to rethink the current, traditional three-tier network design and consider a migration to a flatter network design.

The cost and time savings of server virtualization are tremendous, and virtualizing applications unleashes great opportunities. The difficulty, though, is that server virtualization changes the dynamics of network traffic from a north-south pattern to a multi-directional pattern. Moreover, Gartner predicts that by 2014 more than 80 percent of network traffic will be server to server¹. This will force next-generation data centers to change in an unprecedented fashion.




Figure 1. Example of a disaggregated application.

A network redesign is required for the following reasons:

- As more Virtual Machines (VMs) are added to each physical server, the traffic load per server increases accordingly. Application service levels become more difficult to maintain due to increased server-to-server traffic latency, as data is forced to run north-south before reaching its destination.
- Web 2.0 and application disaggregation (Figure 1) that depend on server-to-server communication can become poor performers in a three-tier environment.
- VM mobility is supported only at Layer 2 and is limited to the size of the Layer 2 network. Network traffic and Spanning Tree Protocol (STP), which allows only one active path between switches, are causing the build out of large, unwieldy, hard-to-manage Layer 2 networks.
- The increasing complexity of a three-tier architecture, and its inability to fully utilize the entire network due to STP limitations, creates cost inefficiencies and prevents IT organizations from quickly responding to changing business demands.

Next, in “A Short History of Data Center LANS and Predictions for the Future,” we will look at what is happening within data centers today and what is forecast to occur during the next couple of years.

¹Munch, Bjarne (Gartner). Your Data Center Network Is Heading for Traffic Chaos (ID G00210674). April 27, 2011.

A SHORT HISTORY OF DATA CENTER LANS AND PREDICTIONS FOR THE FUTURE

What does your data center look like today, and how is it likely to look in five years? Think flat networks with at least some services from the cloud.

In enterprise data centers around the world, today's network is transitioning from the classic hierarchical, three-tier network design to a flatter design that can better address the challenges created by virtualization, Web 2.0, and cloud computing. Evolving business demands are also causing changes to ripple through data centers as CEOs and CIOs look for ways to better leverage their existing IT resources to gain a competitive edge.

Looking ahead three to five years, organizations will be able to provision their own applications and services on demand from their own data center, as well as take advantage of dynamic, on-demand applications and services available from the cloud. But first, the network must pass through several stages along the path to the cloud.

Classic Three-Tier Architecture

The dominant architecture today—and one that has been around more than a decade—is the conventional three-tier, or hierarchical, data center network architecture (see Figure 2). This architecture includes the familiar LAN access, aggregation, and core tiers. It dates back to the era when clients consumed applications running on dedicated physical servers, and network traffic typically flowed from the client, through the data center network tiers, to the application, and back out. This traffic pattern is typically referred to as north-south. This environment tolerates oversubscription in the switching components because, on average, each server connection utilizes a relatively small portion of network bandwidth. To help ensure application availability, network resiliency is delivered through redundant switching components and network connections.

Virtualization and Web 2.0 are causing breaks in the classic three-tier architecture, which will lead to the next data center stage—a flat network design.




Figure 2. Classic hierarchical, three-tier architecture with VMs.

Flat Network Design

The more powerful **flat network design** supports higher traffic loads and increasing east-west traffic in virtualized environments, while avoiding network congestion. Collapsing network layers also reduces complexity, which lowers overhead costs and reduces risk. But challenges still remain.

This flatter design requires high-density, high-bandwidth, and low-latency network components that deliver full wire-speed connectivity. At the same time, Spanning Tree Protocol (STP) brings traffic to a halt during tree convergence, allows only one active path between switches, and requires a manual switch reconfiguration when changing inter-switch connections or attempting to move a Virtual Machine (VM)—limiting network scalability as well as IT agility and efficiencies.

To overcome scalability, management, and productivity challenges, the next step in the data center transformation is Ethernet fabrics.

Ethernet Fabrics

In Ethernet fabrics (see Figure 3), this large, flat Layer 2 network delivers high wire-speed performance and high network resiliency. In addition, Ethernet fabrics are efficient. All paths between switches are fully active, and traffic is continuously routed to use the most efficient path. The network topology is scalable, flexible, and dynamic—changing quickly as the needs of the business change. And, if appropriate to the application, IP and storage traffic can be converged over a common network infrastructure, further reducing cost.

Ethernet fabrics provide unprecedented levels of network intelligence that enable seamless VM mobility and simplified management across the entire data center environment. With Ethernet fabrics, all devices in the fabric are aware of each other. As a VM moves, manual reconfiguration is no longer needed because the VM's profile information already exists in the fabric and is known by all network devices. Network administration is simplified since all switches in the fabric can be managed as a single entity or individually as needed. Last, Ethernet fabrics are self-forming, self-aggregating, and VM-aware. By simply having an administrator add a switch to the fabric, Inter-Switch Links (ISLs) are automatically configured and aggregated, and VM profile settings are automatically extracted from the hypervisor and applied.




Figure 3. Ethernet fabric design.

Private Clouds

Once an Ethernet fabric-based infrastructure is established, the next step is to begin moving toward a private cloud model (see Figure 4). In this stage, data centers have a large, flat, fully utilized Layer 2 network that provides high bandwidth and delivers a high level of network automation. In some cases—and for some applications—IT can consider connecting storage resources to the Ethernet fabric. With this infrastructure in place, IT organizations are able to scale virtualization, improve and refine automation, and rewrite IT policies and processes for a more services-oriented approach.

A private cloud model combines a services-oriented approach with policies, processes, and infrastructure automation in which infrastructure—compute, storage, and networking—is procured by the project for business units and implemented rapidly. Infrastructure resources can be delivered as services to any application utilizing the common resource pool, rapidly allocated (hours or days), and charged back to business units based on usage.




Figure 4. Private cloud model.

Extended Private Clouds

To create an extended private cloud (see Figure 5) and seamlessly extend VM mobility to other data centers, organizations will need to extend the Layer 2 network over distance and be able to accelerate and access applications over WAN connections. Multiprotocol Label Switching (MPLS) technology, Ethernet fabric innovation, 100 Gigabit Ethernet (GbE), and high-speed data transfer technology are key enabling technologies for extended private clouds.

Using Ethernet fabrics as the foundation, the extended private cloud allows IT organizations to expand their pools of compute, storage, and network resources by leveraging infrastructure investments across multiple data center locations. IT gains more flexibility to scale applications and meet rapid increases in demand. IT can also quickly relocate applications and data to capitalize on lower energy costs at different times, or better manage data center maintenance projects that would normally take an application offline temporarily.




Figure 5. Extended private cloud model.

Hybrid Clouds

The next model, a hybrid cloud (see Figure 6), builds on the concept of private and extended private clouds. A hybrid cloud model, with its Ethernet fabrics, requires a network infrastructure built with standards-based technology; open support for hypervisors, servers, and storage; and tight integration with systems management applications.

In a hybrid cloud, resources from the private and public clouds are combined so that businesses can be more agile and responsive. The hybrid cloud allows IT organizations to leverage other data center locations and resources from service providers in the public cloud. Applications, such as e-mail, data storage, and Customer Relationship Management (CRM), are served through the cloud at lower costs. Financial data or Enterprise Resource Planning (ERP) remain on the private cloud. New services and applications can be provisioned quickly when needed—and de-provisioned when no longer necessary. This model gives IT organizations more options to cost-effectively meet spikes or seasonal demand for applications and storage, supports speedy application deployments, and ensures high service levels in the event of a local outage.




Figure 6. Hybrid cloud model.

Expect the Unexpected

Of course, change is the only constant in life. Even though the vision of these future data center models is consistent with what analysts and other industry leaders are forecasting, new technologies are sure to appear and alter the landscape again.

Architecting a data center that meets the challenges of today and transitions smoothly to technologies that meet the challenges of tomorrow is no easy task. Architects need to invest in technologies and platforms that are efficient, intelligent, resilient, and can adapt or be easily upgraded to provide long-term value. At the same time, the executive management and technology leaders of every organization must form a strong partnership to align business goals with technology capabilities. Only by working together will the organization reap the clear benefits of this transformation: greater business agility and increased cost efficiency. With technology changing so fast, an organization's competitiveness rests squarely on its ability to take advantage of the business benefits provided by technology improvements.

ETHERNET FABRICS 101

Understand what Ethernet fabrics are and how they impact your business and your IT department.

In "A Short History of Data Center LANs and Predictions for the Future," you read about the upcoming data center transformation and how networks will ultimately move to a cloud-based architecture. To make your move to the cloud, you will first need to transition your infrastructure to an Ethernet fabric. But what is an Ethernet fabric?

IDC and Brocade define Ethernet fabrics as switch networks that are much more flexible and simpler than today's common Layer 3 routed networks. Ethernet fabrics also provide greater scalability than classic hierarchical switched networks that use Spanning Tree Protocol (STP). These networks can be loosely described as "Layer 2 routing," and allow organizations to build large, flat, and fast data center networks that can be routed based on location-independent Layer 2 MAC addresses.

Two technologies are behind the significant performance, management, and efficiency improvements in Ethernet fabrics—Data Center Bridging (DCB) and Transparent Interconnect of Lots of Links (TRILL).

TRILL, a new standard protocol recently ratified by the Internet Engineering Task Force (IETF), provides a multi-path Layer 2 alternative to the single path and network bandwidth-limiting STP currently deployed in data center networks. TRILL offers the following advantages:

- Eliminates the need for STP and utilizes formerly unused links
- Creates a loop-free infrastructure that continuously analyzes and routes traffic to the most efficient path

DCB, another enhancement to Ethernet, provides additional benefits:

- Provides granular control for reliability and improved Quality of Service (QoS)
- Lays the foundation for reliably transmitting storage traffic (lossless Ethernet)

DCB and TRILL are two technology advancements behind Ethernet fabrics. To get a broader perspective on Ethernet fabrics, let's compare five characteristics of today's typical Ethernet networks to Ethernet fabric networks.

Flatter Architecture

Classic Ethernet networks are hierarchical with three or more tiers. Traffic has to move up and down a logical tree to flow between server racks, adding latency and creating congestion on Inter-Switch Links (ISLs).

STP prevents loops by allowing only one active path, or ISL, between any two switches. This means that ISL bandwidth is limited to a single connection, since multiple paths between switches are prohibited. Enhancements to Ethernet tried to overcome this limitation. Link Aggregation Groups (LAGs) were defined so that multiple links between switches were treated as a single connection without forming loops. But a LAG must be manually configured on each port in the LAG and is not very flexible.

Ethernet fabrics prevent loops without using STP. Flatter networks include self-aggregating ISL connections between switches, which eliminate manual configuration of LAG ports while providing non-disruptive, scalable bandwidth within the fabric. Ethernet fabrics support any network topology (tree, ring, mesh, or core/edge) and avoid bottlenecks on ISLs as traffic volume grows—since all ISLs are active.

Intelligence

Classic Ethernet switches require configuration of each switch port. Configuration requirements include setting network policies such as QoS, security, Virtual LAN (VLAN) traffic, and others. When only physical servers were connected to the network, this model was sufficient. But today, server virtualization requires multiple Virtual Machines (VMs) to be configured on each switch port. When a VM migrates, for either load balancing or routine maintenance, the port configuration has to move to a new network port or the migration fails. This requires manual configuration.

Ethernet fabrics have built-in intelligence, which allows common configuration parameters to be shared by all switch ports in the fabric. In the case of VM migration, the network policies for that VM are known at every switch port, so migration does not require any changes to network configuration. VMs can move seamlessly and transparently to any compute node in the primary data center or another data center. Network policies such as QoS and security also automatically follow VMs when they move.

In an Ethernet fabric, switches share configuration information and sophisticated automation. When a device connects to an edge port of the fabric, all switches know about that device. As the device sends traffic to other devices, the fabric can identify the shortest loop-free path through the fabric and forward frames with the lowest-possible latency. New traffic types such as VM migration and storage traffic are latency-sensitive. The Ethernet fabric ensures that this traffic gets to its destination with minimal latency.

Scalable

Classic Ethernet allows only one path between switches. LAGs allow several physical links to act as a single link, but these links must be manually configured on every port in the LAG—an inefficient and time-consuming process. If a new switch is added for more connectivity, it becomes increasingly more complex to manually configure multiple LAG connections.

Ethernet fabrics overcome this challenge. When a new switch connects to the fabric, no manual configuration is required for the ISLs. The switch joins the fabric and learns about all the other switches in the fabric and the devices connected to the fabric. This eliminates the need for manual configuration of policies and special LAG configurations on specific ports.

If multiple ISLs are connected between two switches, a logical trunk automatically forms. Traffic is load-balanced in hardware so that utilization is near line rate on every link for high efficiency and performance. Should a link in a trunk go offline, traffic on the remaining links is not affected and incoming frames are automatically distributed on the remaining links without disruption to the devices sending them.

Efficient

Classic Ethernet uses STP to create a loop-free path, forming a logical hierarchical switch tree. Even when multiple links are connected for scalability and availability, only one link or LAG can be active. This lowers utilization. When a new link is added or removed, the entire network halts all traffic for tens of seconds to minutes while it configures a new loop-free tree. This is highly disruptive for storage traffic, VM migration, and so on. In the case of storage traffic, traffic disruption could cause a server crash.

The Ethernet fabric does not use STP to prevent loops. It uses link-state routing with Equal-Cost Multi-Path (ECMP) routes, which always take the most-efficient path through the network. When a link is added or removed in the fabric, traffic on other links continues to flow non-disruptively. Link resiliency is assured and full utilization of all links between switches is automatic when the topology is changed.

Simplified Management

Classic Ethernet switches require management. Each switch has to be configured, and each port has to be configured for protocols (STP, RSTP, MSTP, LAG, and so on), VLANs, network policies, QoS, and security. As more server racks are added, more switches are added at the top of the rack, middle of the row, or end of the row. Each requires configuration, and none can share configuration parameters.

An Ethernet fabric shares configuration information among all switches in the fabric. The fabric appears to network administrators as one large switch and allows increased traffic visibility and control of the network while reducing network management and overhead. Each physical switch in the fabric is managed as if it were a port module in a chassis. When a new switch joins the fabric, it automatically receives common information about devices, network policies, security, and QoS. This simplifies network configuration, reduces mistakes, and reduces operating cost. No manual intervention is necessary.

A Business Perspective

So far, this discussion of Ethernet fabrics has focused on the technical aspects of this technology. Now let's look at how Ethernet fabrics will impact your business. From a business point of view, Ethernet fabrics will eliminate the typical infrastructure barriers that prevent organizations from realizing the full potential of virtualization and, most important, reacting quickly to business opportunities. Consider these scenarios:

- The sales and business development teams have a new service offering idea. You need to put it into production.
- A new sales promotion is launching or the buying season is approaching, and Web traffic is expected to increase dramatically.
- Your company just announced a merger and you are responsible for integrating all systems, platforms, and networks.
- Executive management and sales are scheduled to receive new mobile devices with CRM and analytic applications in two months.
- The compliance deadline for encryption is four weeks away.
- Poor application performance in branch offices must be improved through load balancing.

With Ethernet fabrics, IT can respond to these requests with a “Yes, we can deliver that in a few hours,” rather than outlining why these changes will take significant amounts of time, resources, and planning to complete.

Upgrading to Ethernet fabrics also leads to cost savings. Businesses can expect to lower capital expenses, reduce operational expenses, and eliminate lost revenue due to unplanned outages.

For example:

- **Decreasing capital expenses.** Consider for a minute the built-in redundancies and idle links in your current network. Ethernet fabrics can reduce overall purchases for new equipment by as much as 30 percent. Organizations can purchase switches with fewer ports because they will be able to use all the links instead of having half of them sit idle. Ethernet fabrics double the ISL resiliency and support twice the bandwidth between servers.
- **Reducing operational expenses.** Ethernet fabrics’ sophisticated automation decreases the time and resources needed to configure and administer the network and manage VMs. Ethernet fabrics also decrease overall network management. Organizations can expect up to a 30 percent reduction in operational costs due to time savings.
- **Eliminating lost revenue.** When a network outage occurs, businesses lose potential revenue, and they tarnish their good reputation. While you cannot quantify the loss of potential revenue or a poor reputation, in this era of social media, one bad day for the network can open a business up to considerable public criticism.

Ethernet fabrics provide organizations with the competitive advantages they need to support a more agile, streamlined business that can easily support its users and be cost-effective. The key takeaways you need to know about Ethernet fabrics are:

- The network is resilient without being redundant.
- Latency is extremely low (as much as 10 times lower than hierarchical networks) while scalability is quite high.
- Multi-path technology delivers increased bandwidth between servers.
- The overall architecture is flat, fast, and efficient.

ETHERNET FABRIC ASSESSMENT

Complete this short exercise to help gauge your network’s current virtualization maturity stage and uncover potential gaps as you move toward an Ethernet fabric.

Has virtualization caused problems on your network or challenges for your staff? Do you want to understand what stage your data center network is at? Brocade and Forrester Consulting have defined four stages of infrastructure maturity, which track where you are on the virtualization path. Once you understand where you are in terms of virtualization maturity, you can better assess your need and readiness for Ethernet fabrics.

The four stages of virtualization maturity are:

Stage 1: Acclimation

- Comfortable with virtualization as a concept and tool
- Deployed for testing and development
- Some production deployments, but tactical
- No change to operations processes
- Limited virtualization tool deployments

Stage 2: Strategic Consolidation

- Comfortable with concept, use, maturity, stability
- Mindset has shifted from server to virtual server
- Spread production deployments widely
- Beginning deployment for some business-critical disaster recovery
- Experimenting with live migrations of Virtual Machines (VMs) and balancing resource pools with utilization

Stage 3: Process Improvement

- Using live VM migration and gaining confidence with utilization and resource availability feedback
- Deployed for business-critical disaster recovery
- Beginning to bifurcate applications between priority and nonpriority
- Developing new operational efficiencies
- Process improvement spreading against network, storage, and security development

Stage 4: Pooling and Automation

- Full confidence in automation and resource optimization
- Automation policies in production
- Some mission-critical disaster recovery deployments
- Pooling and internal cloud development
- Charge back/utility tracking
- Service Level Agreements (SLAs) and Quality of Service (QoS) focus

Your readiness for Ethernet fabrics depends on the progress you have made in virtualizing your data center. Organizations that are further along in their virtualization deployments will be more ready—and have a greater need—for Ethernet fabrics than those in the early stages of virtualization. If your data center has not reached Stage 2—Strategic Consolidation—you might not be ready nor have a need for Ethernet fabrics. Organizations that are at Stage 2 or beyond will gain significant advantages from migrating to Ethernet fabrics. To find out which stage your organization is at, take this assessment and add your implementation and process scores, then multiply the total by 2.

Forrester's Infrastructure Virtualization Maturity Assessment

$$\left(\begin{array}{|c|} \hline \text{Implementation} \\ \hline \text{score} \\ \hline \end{array} + \begin{array}{|c|} \hline \text{Process} \\ \hline \text{score} \\ \hline \end{array} \right) \times 2 = \begin{array}{|c|} \hline \text{Total maturity} \\ \hline \text{score} \\ \hline \end{array}$$

Stage 1 = 1-25 points

Stage 2 = 26-50 points

Stage 3 = 51-75 points

Stage 4 = 76-100 points

Criteria	Score explanations	Scores
Implementation		
What percentage of your test or development environment is virtual?	1 = 0-25% 2 = 26-50% 3 = 51-75% 4 = 76-100%	
What percentage of your production environment is virtual?	1 = 0-25% 2 = 26-50% 3 = 51-75% 4 = 76-100%	
What percentage of your mission-critical servers is virtual?	1 = 0-25% 2 = 26-50% 3 = 51-75% 4 = 76-100%	
Do you have an executive sponsor for your virtualization implementation?	0 = No 3 = Yes	
Do you boot all VMs from networked storage?	0 = No 2 = Yes	
What is your virtual server host utilization target?	1 = < 10% 2 = 10-30% 3 = 31-60% 4 = > 60%	
How many virtual machines do you deploy on one physical host?	1 = < 10 2 = 10-20 3 = 21-30 4 = 31+	
	Total implementation score:	
Processes		
Do you utilize live migration?	0 = No 1 = Yes	
Do you utilize automated resource scheduling?	0 = No 2 = Yes	
Do you utilize VM templates to propagate changes into production?	0 = No 1 = Yes	
Have virtual servers reduced the number of people or tools required to deploy new systems?	0 = No 2 = Yes	
Have you used virtual servers to simplify day-to-day tasks like patching or system changes?	0 = No 2 = Yes	

(Processes continued)

Criteria	Score explanations	Scores
Are you virtualizing applications even if they require a dedicated VM host?	0 = No 1 = Yes	
Have you financially accounted for the benefits of virtualization to your organization?	0 = No 1 = Yes	
Have you set up improved SLAs for your virtual environment (e.g., better availability)?	0 = No 2 = Yes	
Do you charge back or allocate costs based on virtual resource consumption?	0 = No 2 = Yes	
Are you using virtualization-optimized management tools for VM backups?	0 = No 1 = Yes	
Are you using virtualization-optimized management tools for VM monitoring?	0 = No 1 = Yes	
Are you using virtualization-optimized management tools for VM migrations?	0 = No 1 = Yes	
Are you using virtualization-optimized management tools for capacity planning?	0 = No 1 = Yes	
Are you using virtualization-optimized management tools for high availability?	0 = No 1 = Yes	
Does every VM that you deploy start with an approved template from a formal library that is maintained and updated centrally?	0 = No 1 = Yes	
Have you implemented a self-service portal for provisioning VMs?	0 = No 1 = Yes	
Do your testing and development VMs all have expiration dates?	0 = No 1 = Yes	
Have you implemented a “virtual first” policy?	0 = No 2 = Yes	
Do you have a virtual infrastructure architect on staff?	0 = No 1 = Yes	
	Total process score:	

TRANSITIONING TO ETHERNET FABRICS

When and how do you transition to an Ethernet fabric? Four use cases show how you can upgrade your current network architecture to an Ethernet fabric, and which applications and services are good early candidates.

If you take only one idea away from this handbook, remember this: Ethernet fabrics do not require you to rip up and replace your network.

You can leverage your existing infrastructure—including routers, firewalls, 1 and 10 Gigabit Ethernet (GbE) platforms, and other Layer 3 devices—and migrate to Ethernet fabrics when you are ready, and target

where you need the increased flexibility and ease of management. You set the pace and specific areas that you want to pilot and deploy to production. With that clear, let's look at a viable roadmap that will help transition your network from a rigid hierarchical, multitier design to a more flexible architecture that supports a highly mobile Virtual Machine (VM) environment, on-demand services, and increased business agility.

As with any major infrastructure transition, you will want to migrate your infrastructure in stages, identifying applications and projects that will most benefit from Ethernet fabrics. The use cases outlined below are the most common approaches to transitioning to Ethernet fabrics. Each use case is driven by business factors and priorities, and the common desire to preserve existing investments as much as possible.

Architecture: 1/10 Gbps Top-of-Rack (ToR) access

Use cases: Server consolidation, data center virtualization (Stage 1, 2, or 3 in the Ethernet fabric assessment)

In this scenario, you are adding servers with 1 or 10 Gbps LAN connectivity and need the additional bandwidth. Adding a switch that has 1 or 10 GbE and, most important, Ethernet fabric-ready capabilities to the top of the server rack will provide the additional capacity you need. The 1/10 GbE Top-of-Rack (ToR) switch can serve as a classic low-latency, high-density switch until you are ready to turn the Ethernet fabric features on. With an Ethernet fabric-ready switch, you simply “turn on” Ethernet fabric capability on each switch. The fabric automatically forms, allowing you to launch your pilot or move into production easily, with little or no downtime.




Figure 7. Use case: 1/10 Gbps Top-of-Rack (ToR) access architecture.

Benefits:

- Preserves existing core and access architecture; coexists with vendor-agnostic ToR switches
- Supports 1 Gbps and 10 Gbps server connectivity
- Enables seamless transition from traditional Layer 2 at ToR to Ethernet fabric-enabled Layer 2 at ToR

Once the Ethernet fabric is enabled, the in-rack network has:

- 100 percent link utilization
- Seamless intra-rack VM mobility

Architecture: 10 Gbps aggregation, 1 Gbps Top-of-Rack (ToR) access

Use cases: Server consolidation, data center virtualization (Stage 1, 2, or 3 in Ethernet fabric assessment)

This architecture allows you to use 10 GbE and Ethernet fabrics as the aggregation technology while preserving your investment in 1 Gbps servers and ToR switches, with or without “stackable” technologies. In this scenario, you do not expect to upgrade to 10 Gbps servers in the near future, nor want to continue using your stackable ToR solution, but you are experiencing bottlenecks at the aggregation layer. You can add 10 GbE Ethernet fabric switches in the aggregation layer to eliminate those bottlenecks and the need for Spanning Tree Protocol (STP) at the same time.




Figure 8. Use case: 10 Gbps aggregation; 1 Gbps ToR access architecture.

Benefits:

- Enables building-block scalability
- Provides flexible subscription ratios
- Supports optimized multi-path network
- Eliminates the limitations of STP in the aggregation layer
- Reduces operational costs through simplified management
- Decreases capital costs through increased link utilization

Architecture: 1/10 Gbps access, convergence plus Fibre Channel Storage Area Network (SAN)

Use cases: Server consolidation, data center virtualization (Stage 1, 2, or 3 in Ethernet fabric assessment)

This architecture leverages all the benefits of Ethernet fabrics by consolidating the access and aggregation tiers into one Ethernet fabric tier while providing the most flexible storage options. For servers with existing Fibre Channel connectivity, direct server-to-SAN connectivity remains. For servers needing to connect to the SAN but lacking Fibre Channel adapters, the connection can be made through the Ethernet fabric. In all cases, servers communicate across the data center and out to the clients using the Ethernet fabric.

This architecture will support your plans to consolidate servers, deploy Ethernet-attached storage, and match the storage option (Ethernet or Fibre Channel) that best meets the workload and business requirement.




Figure 9. Use case: 1/10 Gbps access; convergence plus Fibre Channel SAN architecture.

Benefits:

- Introduces a flatter, simpler network design with a logical two-tier architecture
- Increases Layer 2 scalability and flexibility, supporting more VM mobility options, server-to-server communication, and seamless network expansion
- Optimizes multi-path network; all paths are active, no STP, and no single point of failure exists
- Supports maximum storage flexibility

Architecture: Ethernet fabric architecture for an extended private cloud

Use cases: Server consolidation, data center virtualization (Stage 1, 2, 3, or 4 in Ethernet fabric assessment), Ethernet-attached storage, combined Fibre Channel and Ethernet-attached storage, extended private cloud

The data center transition to Ethernet fabrics is complete. Now both the SAN and LAN are interconnected through the Ethernet fabric, or the Fibre Channel SAN might still be in use. All ToR switches have Ethernet fabric capabilities, and they connect to both 10 Gbps servers and, in some cases, 8 Gbps or 16 Gbps Fibre Channel switches. Multiple private clouds are connected using storage and fabric extension technologies to decrease latency and deliver seamless connectivity and VM mobility across the rack, the data center, or over distance.




Figure 10. Use case: Ethernet fabric for an extended private cloud.

Benefits:

- Provides maximum storage flexibility
- Optimizes resource capacities and power usage
- Enables greater flexibility and application availability

A DICTIONARY FOR ETHERNET FABRICS.

View brief definitions of terms, technologies, and standards for the Ethernet fabric-based data center.

Cloud computing—Logical computational resources (data, software) accessible via a network (through a WAN or the Internet), rather than from a local computer. Data is stored on server farms, usually in the same country as the service provider.

Control plane—Layer 2 link-state routing protocol ISIS.

Convergence—The ability of a single network infrastructure to support the needs of multiple technologies. In data center networking, convergence means combining LAN and storage traffic on one infrastructure or infrastructure type, such as Ethernet.

Converged Enhanced Ethernet (CEE)—Considered the precursor to Data Center Bridging (DCB), CEE is a set of proposals to enhance Ethernet protocol to better support the transport of storage traffic.

Data Center Bridging (DCB)/Data Center Bridging Capabilities Exchange Protocol (DCBX)—DCB is an enhancement to Ethernet to improve traffic reliability, provide more granular control (such as improved QoS), and lay the foundation for reliably transmitting storage traffic.

Data plane—A TRILL protocol.

Distributed intelligence—Network-wide knowledge of all members, devices, VMs, port settings, and more

Ethernet—A ubiquitous networking technology.

Ethernet fabrics—A new network architecture for providing resilient, high-performance connectivity between clients, servers, and storage

Fabric-based infrastructure—A Gartner term that refers to creating a fabric for everything—including compute, memory, network, storage, and power.

Fibre Channel—A protocol typically used for storage traffic.

Fibre Channel over Ethernet (FCoE)—Encapsulation protocol that enables the transport of Fibre Channel storage traffic over a new lossless Ethernet medium.

Flat network—A network in which all hosts can communicate with each other without requiring a Layer 3 device.

Hierarchical networks—Traditional three-tier networks that include core, distribution, and access layers.

Hybrid cloud—Two or more clouds (private, community, or public) that remain unique entities but are bound together, offering the benefits of multiple deployment models.

Link-state routing—Allows routers to calculate the best path to any router on the network by discovering RBridge peers, determining RBridge Virtual LAN (VLAN) topology, establishing Layer 2 delivery using shortest path calculations, and informing routers of their closest neighbor on the network.

Multiple Spanning Tree Protocol (MSTP)—Protocol that allows separate spanning trees per VLAN.

OpenFlow—An emerging standard for software-defined networking that makes the network more automated and easier to manage.

OpenStack—A series of management tools that makes the network infrastructure effectively transparent to the world.

OpenFlow and OpenStack consortia—Working groups that are aligning to disconnect the hardware from the provisioning and managing framework, in the same way that virtualization has disconnected applications from hardware.

Private cloud—Infrastructure operated solely for a single organization, whether managed internally or by a third party and hosted internally or externally.

Public cloud—Cloud computing in the traditional mainstream sense, whereby resources are dynamically provisioned to the general public on a fine-grained, self-service basis over the Internet, via Web-based applications and services, from an off-site third-party provider who bills on a utility basis.

Rapid Spanning Tree Protocol (RSTP)—Protocol that provides significantly faster spanning tree convergence after a topology change, introducing new convergence behaviors and bridge port roles.

Routing bridges (RBridges)—RBridges are a new type of Layer 2 device that implements the TRILL protocol, performs Layer 2 forwarding, and requires little or no configuration. Using the configuration information distributed by the link-state protocol, RBridges discover each other and calculate the shortest path to all other RBridges on the VLAN.

Server virtualization—A software implementation of a machine, such as a server, that executes programs like a physical server.

Spanning Tree Protocol (STP)—Protocol that provides a loop-free topology for a LAN or bridged network.

Storage fabrics—Commonly called a Storage Area Network (SAN).

TCP/IP—Protocol of choice for peer-to-peer and client server networking.

Transparent Interconnect of Lots of Links (TRILL)—Provides a Layer 2 multi-path, multi-hop alternative to the single path and network bandwidth-limiting STP currently deployed in data center networks.

802.1aq: Shortest Path Bridging (SPB)—Enables shortest path trees, thereby providing the ability to use all available physical connectivity because of loop avoidance.

802.1Qau: Congestion Notification (QCN)—End-to-end congestion management that enables throttling of traffic at the edge nodes of the network in the event of traffic congestion.

802.1Qaz: Enhanced Transmission Selection (ETS)—Provides the capability to group each type of data flow, such as storage or networking, and assign an identification number to each group. It manages bandwidth on the Ethernet link by allocating portions (percentages) of the available bandwidth to each of the groups, known as traffic class groups.

802.1Qbb: Priority-based Flow Control (PFC)—Establishes eight priorities for flow control based on the priority code point field in the IEEE 802.1Q tags. This enables the control of individual data flows on shared lossless links. The PFC capability allows Fibre Channel storage traffic encapsulated in FCoE frames to receive lossless service from a link that is shared with traditional LAN traffic, which is loss-tolerant.

ABOUT BROCADE

Brocade® networking solutions help organizations transition smoothly to a world where applications and information reside anywhere. Innovative Ethernet and storage networking solutions for data center, campus, and service provider networks help reduce complexity and cost while enabling virtualization and cloud computing to increase business agility. Learn more at www.brocade.com.

For additional information or to join the Ethernet fabrics conversation, visit www.ethernetfabric.com.



© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 11/11 GA-BR-1627-01

Brocade, the B-wing symbol, DCX, Fabric OS, and SAN Health are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.