

**DATA CENTER**

## **Data Center Evolution without Revolution**

Brocade networking solutions help organizations transition smoothly to a world where information and applications can reside anywhere.

**BROCADE**

**Preparing a data center for virtualization and cloud computing does not have to be an all-or-nothing proposition. Organizations can control risk and reduce costs by choosing a cost-effective, phased network architecture upgrade that enables new levels of application flexibility and business agility.**

**Today, Brocade® provides a complete family of Ethernet, storage, and converged networking solutions that help forward-thinking organizations build strategic data centers to meet their business requirements today and well into the future.**

The proven cost savings and high return on investment from virtualization and cloud computing in the data center have captured the imagination of many CFOs and CIOs. This is especially the case in a challenging economic period, when few organizations are adding headcount and most IT budgets remain flat. The pressure remains to reduce complexity and optimize existing resources for maximum value.

In most cases, organizations can reduce data center costs by moving more applications onto fewer servers and can reduce software licensing fees and administrative resources by migrating to a Software-as-a-Service (SaaS) model. When they add in faster deployment times for new applications and services while ensuring 24x7 reliability, organizations have a compelling business case for virtualization and cloud computing in the data center. At some point in the future, this type of highly virtualized, services-on-demand delivery model enabled through cloud computing will be the IT gold standard.

Before this transition occurs, however, the underlying network architecture must provide greater availability, performance, and security while simplifying data center operations. Data centers will need network architectures that support more virtualized environments and cloud operations models that reduce the time to deploy new applications.

Preparing the virtualized data center for cloud services doesn't require a complete replacement of the data center network from access layer to core. Rather than rip and replace their network infrastructures, organizations can follow a more attractive—and cost-effective—strategy of evolving their data centers.

The economic leverage from virtualization and cloud computing does not require network convergence of storage and IP traffic on the same physical network. In fact, critical applications typically should not be deployed on a converged physical network. However, other applications with higher outage tolerances can be deployed on a physically converged network. Therefore, the network architecture must be able to accommodate applications where storage and IP traffic are on physically separate networks and where they are converged.

Slowly evolving the network enables organizations to consolidate their resources as well as support server virtualization and cloud computing wherever and whenever it makes sense. A complete revolution in the data center, with its associated costs, disruptions, and time requirements, is unnecessary in almost every case. Instead, the transition to a highly virtualized data center can be smooth, non-disruptive, and full of business advantages.

## **VIRTUALIZATION AND CLOUD COMPUTING IN THE DATA CENTER**

As organizations add Virtual Machines (VMs) and shift to a cloud-based operations model, the underlying network architecture becomes more critical because of high-availability requirements, potential security threats, an increase in data traffic, and network management challenges created by VM mobility. The network must be hardened against failures that disrupt traffic yet adaptable and flexible enough to support new business requirements.

In terms of the physical network infrastructure, virtualization and cloud computing complement each other in the following ways:

- Cloud computing services can reside on top of virtual data centers. Virtualization can support cloud architectures.
- Cloud computing software can orchestrate virtual server deployments. Cloud management software can control virtualization services.
- Cloud computing adds another virtualization layer between end users and the entire IT environment, which enables a pay-per-use model.
- Both virtualization and cloud computing require robust physical network infrastructures. They rely heavily on the network and require new ways of thinking about network architecture and design.

As more server virtualization and cloud computing projects are deployed in the data center, the existing network design must adapt gracefully, one rack at a time, while continuing to operate with non-virtualized applications. Not all applications require virtualization, so the network infrastructure must be flexible enough to support applications on both dedicated and virtualized hardware.

Many organizations are considering the option of network convergence, which combines SAN traffic and IP network traffic on the same physical network as a way to reduce the amount of cabling in server racks and the number of top-of-rack switches. The underlying driver for network convergence is cost reduction. If half as many network interfaces and switch ports are required on a converged network, the costs are much lower. However, there are important tradeoffs that will limit adoption of this option, because the network architecture has to support both network convergence and non-converged traffic.

Even though network convergence is expected to become more prevalent, it is not inevitable that all organizations and all types of applications will be good candidates for convergence. For instance, high-availability and high-performance applications that are not well suited for convergence will remain on separate networks. Network latency, change control processes, regulatory compliance, and security are other factors that might require IP and storage traffic to remain separated. The good news is that each of these architectures can be enhanced to support the virtualized data center.


## **SEPARATE OR CONVERGED NETWORKS FOR VIRTUALIZED DATA CENTERS**

When and how far to converge the IP and Fibre Channel traffic—or not to converge—is a decision that should be made in the context of the unique requirements of each organization. In the short term, organizations can phase-in network convergence and reduce complexity while still supporting virtualization and cloud computing services. Potential convergence options include:

- No convergence
- Management convergence
- Layer 2 technology convergence and physical network separation
- Access layer convergence
- Entire network convergence

## No Convergence

Data centers can retain the classic architecture of two separate networks. To support virtualization, some management features of the LAN and SAN must be added to manage server and storage nodes and to support the automated changes when VMs move between physical servers (see Figure 1).



**Figure 1.**

Next-generation data center, non-converged.

## Management Convergence

Converging management means adding orchestration tools that coordinate LAN and SAN behavior. Managing the SAN and LAN from a single management platform requires that tools simultaneously communicate with LAN and SAN switches. It does not require convergence of IP and Fibre Channel traffic on the same physical network.

## Layer 2 Technology Convergence and Physical Network Separation

In this option, organizations can retain physical network separation for storage and IP traffic but use the same Layer 2 (data link) protocol. Although the SAN and IP networks would use Ethernet switches, the storage and IP networks are deployed, maintained, and managed as physically separate networks. Physical separation implies separate switches for the IP and SAN traffic, and no network path exists between them. Technology convergence through the use of the same Layer 2 protocol can reduce cost and simplify training. It also allows the use of other storage protocols such as iSCSI, Network File System (NFS), and Common Internet File System (CIFS), while ensuring high availability via physically separated storage and IP networks.

Layer 2 technology convergence for SAN access relies on emerging technologies such as Fibre Channel over Ethernet (FCoE) via a physically separate lossless Ethernet network through Data Center Bridging (DCB). Servers will converge IP and storage traffic at Layer 2 using Ethernet and Fibre Channel technology. The DCB switches integrate with existing Fibre Channel storage while also supporting backend NAS protocols such as NFS or CIFS, iSCSI, and FCoE storage traffic.

FCoE can be terminated natively by attaching FCoE storage to the switches, or it can be bridged into the existing Fibre Channel storage fabric. The extent of the converged network will be limited to the network segments that need it, and this convergence will not impact the aggregation or core of the IP network—or disrupt the core of the SAN.


**Access Layer Convergence**

Physically converging IP and Fibre Channel traffic inside a server, external network adapter, and top-of-rack switch can reduce the number of cables and switches required. Converged switch products are available to separate IP and Fibre Channel traffic at the top of the rack, forwarding each traffic type to the appropriate layer (the aggregation switch for IP traffic and the core switch for Fibre Channel traffic).

**Entire Network Convergence**

Full physically converged IP and storage networks require a network that can handle IP and Fibre Channel traffic on any port. The converged network model merges what was once two physically independent networks into a single network. This changes the fundamental architecture of the SAN and LAN and reduces availability for the storage traffic. It also has implications for security as the threat surface of the SAN does not include the threat surface of the IP network. This tradeoff for reducing initial equipment cost is not appropriate for all applications.

Physically converging IP and storage traffic does not merge their network services. Having server and storage virtualization layers at the same tier of the data center does not merge either their software or hardware components. Organizations still need separate interfaces into the orchestration software to manage, monitor, and configure the SAN and LAN services in the converged network. Similarly, the interfaces for storage and server virtualization will remain independent (see Figure 2).



**Figure 2.** Next-generation data center, fully converged.

## **DESIGNING THE VIRTUALIZED DATA CENTER ONE STEP AT A TIME**

Organizations can implement a virtualization and private cloud computing architecture (converged or not converged) that makes sense for their particular business requirements. The initial step is to decide how the data center should look in the future by determining the extent of LAN/SAN convergence, the number of layers within each network, the number of switching tiers in each layer, and the management model for virtualization and cloud computing services.

Organizations should be able to answer the following questions to make sure that all of their future investments bring them closer to that goal:

- Will the final network have physical access layer switches in each rack, or will that function reside inside virtual servers?
- What level of oversubscription is acceptable at each stage of the network, given the target server consolidation ratio?
- How much bandwidth would each VM receive if they all tried to communicate simultaneously?
- Will there be an aggregation layer on the LAN, or will large virtual servers connect directly to a high-port-count collapsed access/aggregation layer?
- Will there be a single orchestration tool and hypervisor vendor, or will there be multiple solutions for specific applications and departments?
- Will there be continued use of any existing equipment in that design? And if not, why not?

The key to achieving “revolution through evolution” in data center networking is to move toward a target design along a well-planned path and to use incremental steps to control risk. Organizations should create separate timelines for each project and consider the following actions when making decisions about the steps and the scope of change in each one:

- On the server side, the first step is to classify applications and select the correct virtualization strategy for each. Some may benefit from hypervisors, while others would use application-specific solutions to scale horizontally. Some application groups might be Microsoft-centric and fit optimally into Hyper-V servers, whereas others might work best with Novell Xen or VMware ESX Server.
- Next, classify the existing server platforms that are modern and robust enough to be used in the target design. All candidates should appear on the list of supported servers from the hypervisor vendor and provide sufficient RAM capacity, CPU power, and SAN interfaces. If a server is otherwise sufficient but lacks Fibre Channel SAN connectivity, organizations should consider purchasing a dual-port HBA since adding a card is often less expensive than replacing a server.
- If enough servers can be leveraged in the new design, organizations should perform a rolling migration, purchase a few new servers loaded with virtualization software, and migrate existing applications onto the new servers. This frees up the old hardware, which can be reloaded with virtualization software and then used as the migration target for the next set of applications. Eventually, all applications will be migrated to a virtualization platform, with most applications residing on existing hardware resources.

Organizations can apply similar strategies in the LAN. For instance, assuming that the target design supports server virtualization, an organization can converge the management of the LAN and SAN by integrating them with virtualization orchestration software, isolate SAN and IP traffic on separate physical networks, and implement a two-tier network to achieve significant cost reduction. Many aggregation switches have oversubscribed blades, and many access switches do not have sufficient 10 Gigabit Ethernet (GbE) support. Organizations can replace these switches with higher-performance chassis switches that are connected to the existing core. In a similar way, router cores can be replaced one at a time over a timeframe that enables organizations to eliminate the aggregation switch layer over time.

Organizations can also utilize MPLS/VPLS configuration where appropriate. As virtual servers are deployed, they connect directly to the router via blade server switching blades or a Layer 2 switch at the top of the rack, thereby flattening the network and extending the reach at Layer 2. Tunneling Layer 2 traffic over MPLS/VPLS also enables solutions such as long-distance VM mobility for disaster recovery. Once again, the aggregation layer is removed over time.

If organizations can eliminate an entire tier of switches, they can often save enough on operating costs to pay for a significant portion of the network upgrade to support lower oversubscription ratios in the access layer switches and to cover their investments in virtualization orchestration software.

After organizations have made all of the hardware replacements and reconfigurations, the data center servers will be fully virtualized. Both the SAN and LAN will be managed by open interfaces supporting multiple orchestration software applications. Organizations can select one orchestrator for the entire data center or different solutions for each application. Over time, they can connect this network to external cloud providers to provide capacity for peak loads or to more cost-effectively meet disaster recovery requirements.

## **REVOLUTION THROUGH EVOLUTION**

Organizations can implement a virtualization and private cloud computing architecture at a pace that makes sense for them. Many are already using virtualization projects (and their upgrade requirements) to replace aging LAN infrastructure. The cost of maintaining the older devices is generally higher than replacing them, offering an opportunity to cost-effectively collapse the number of tiers in the IP network.

As a piece of hardware is depreciated and slated for retirement, new components seamlessly move in to replace it. New platforms are often acquired for growth or to improve performance for critical applications. Those applications can migrate onto the new platforms and, unless their hardware resources were completely obsolete, those resources can be repurposed into the server pool for lower-tier applications. Eventually, by rolling in new platforms and repurposing old ones, organizations can transform their data centers. They will follow a revolutionary approach to IT management, but do so one evolutionary step at a time.

To learn more, visit [www.brocade.com](http://www.brocade.com).

**Corporate Headquarters**

San Jose, CA USA  
T: +1-408-333-8000  
info@brocade.com

**European Headquarters**

Geneva, Switzerland  
T: +41-22-799-56-40  
emea-info@brocade.com

**Asia Pacific Headquarters**

Singapore  
T: +65-6538-4700  
apac-info@brocade.com

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 06/10 GA-WP-1490-00

Brocade, the B-wing symbol, BigIron, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, and Turbolron are registered trademarks, and Brocade Assurance, DCFM, Extraordinary Networks, and Brocade NET Health 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.

**BROCADE**