BROCADE

WHITE PAPER

Brocade Dedicated IP Storage Networks for VM Farms—Architecture and Deployment Guide

Brocade recommends deploying a dedicated IP storage network based on Brocade® VCS® Fabric technology, to meet the business demands placed on the growing Virtual Machine (VM) deployments in business-critical functions. Segmenting priority application storage traffic in this way addresses inherent performance limitations of general-purpose Local Area Networks (LANs) while providing resilient data access and automating the ability to scale with application Service Level Agreements (SLAs) and the number of VMs. The following white paper is for executive decision makers, architects, network administrators, and storage administrators who wish to improve their SLAs and application performance. It describes the required architecture, deployment considerations, and best practices for moving to a dedicated IP storage network that is optimized for virtualized applications.

Dedicated IP Storage Networks for Business-Critical Applications

The traditional approach to IP storage networking can create a bottleneck to both the full utilization of high-performance storage and to the application SLAs that storage administrators must maintain. To match the performance needs of VM farms, Brocade recommends moving storage traffic to a dedicated IP storage network based on the more robust and manageable Brocade VCS Fabric technology, in order to take performance and ease of use to the next level.

In our rush to maximize capital investments by scaling application environments with VM farms we have learned that business-critical applications have SLAs that legacy architectures were not designed to support. We have created a new storage problem where accessing data is the new bottleneck and it must be solved with a storage solution.

Evolving Storage and Networking Environments

The modern data center model for supporting business-critical virtualized applications in a large scale VM farm has pushed the approach of directing important application traffic through the shared company IP network to its limits. As with anything that is pushed to its limit, we are beginning to see an increase in problems including:

- **Missed SLAs:** Latency issues are on the increase, due to congestion caused by small traffic spikes, leading to application issues that include:
  - **Downtime:** Congestion-caused latency itself can disrupt application performance, but it can also lead to zombie VM conditions, which are hard to diagnose and often require a manual restart.
  - **Cascading Application Failure:** Application communication delays can become failures that cause buffer overloads and hurt the rest of the traffic that is dependent on the same ports. The higher, or more efficient, the infrastructure utilization, the greater the risk that this will occur.
  - **Productivity loss:** Widespread interruptions might affect email, Unified Communications (UC), or an inventory application, but in all of these cases, employee productivity suffers.
- **Difficult Management:** The existing tools do not meet the needs of the administrators:
  - **Complex scaling:** The speed of growth and rate of change in today's application environments is exceeding the Spanning Tree Protocol, such that it is necessary to plan a single active path and backup paths in case of link failure. (Backup path links are not active and not available for transport, therefore they are underutilized.)
  - **Complex troubleshooting:** Troubleshooting a mixed traffic environment is difficult at best, leading to long ticket times and fixes that cause other problems.
- **Business losses:** As business-critical applications sit on a general-purpose LAN and formerly "common" tools such as file sharing and collaboration become more important to productivity, data access performance issues potentially can cause business losses with greater and greater impact.

To make things worse, storage has done its part to remove bottlenecks with a new generation of highly scalable, extremely high performance storage arrays that have caught and passed the rest of the infrastructure in ability to handle the volume and speed of data access traffic. This significantly increases requirements on the network infrastructure as they become more common, and they're exposing the unintended degradation in overall performance being caused by the sheer volume of other application traffic which is not storage bound talking over the same network.

The New Paradigm for IP Storage Networks

If these problems sound familiar, that is because they have been known and solved before for mission-critical enterprise applications through the advent of the dedicated fibre Channel SAN. The new network bottleneck for important applications isn't about bandwidth—it's about how to ensure the data your users and applications need is being served with the right level of performance. These problems will be solved again with a similar storage network solution, the dedicated IP storage network optimized for storage traffic. The dedicated IP storage network should be based on Brocade VCS Fabric technology to eliminate any single point of failure and to automate with zero touch scale-out and self-healing. As a dedicated storage network, it can be optimized by storage administrators for storage best practices to ensure packet delivery, minimize latency, and be optimized for specific storage best practices for software, applications, arrays, and environments. Another benefit is that the storage administrator—who will now

be responsible for all of the dedicated storage traffic—can manage it with the same familiar Brocade Network Advisor application that they use to manage their mission-critical SANs.

An Important Perspective

A dedicated IP storage network will use approximately the same number of ports to support the storage workload as would have been used on the shared network so the cost is equivalent. However, the different architectures are optimized for different traffic and moving business-critical traffic to the dedicated network allows better performance and easier management of both. Companies make a substantial investment in compute and storage to run their business, and the cost of network equipment in a data center is just a component of that end-to-end infrastructure. However, the wrong design approach could limit the performance of all that attached investment, not to mention the business-critical workloads that it supports.

In most environments, a dedicated IP storage network based on Brocade VCS Fabric technology in a four-switch base solution can be racked and operational—with storage attached and handling VM traffic—within two hours of delivery, with no downtime or disruption to the environment. The entire setup may take less time than solving the issues caused by not having a dedicated IP storage network. The base solution is extremely scalable, providing zero touch scale-out simplicity, yet it starts out with significant capacity. Check the prices on four Brocade VDX® switches to see how inexpensive it is to implement this solution right away.

Getting Started

While deploying Brocade VCS Fabric technology into a dedicated IP storage network is a simple process, good planning is advisable to ensure that you get the most out of your investment.

This section walks you through some basic questions to use as you develop a deployment plan.

Dedicated IP Storage Network Ownership

It may seem obvious who is responsible for network equipment, but you might revisit this question as you begin the planning process. The objective in ensuring the performance of business-critical applications is not to have a better network. The objective is to ensure that data is being served to those applications with the performance required by business needs and committed SLAs. This solution is a data-driven solution, and storage environments have a long history with mature architectures that deliver on SLAs. The new dedicated IP storage network closely resembles the existing Brocade FC SAN in topology and management. This means that the person who is responsible for the SAN may be the optimal fit for ownership.

Furthermore, storage administrators will be responsible for all the resulting traffic—another reason it may be optimal for them to own the end-to-end solution. Finally, as a dedicated IP storage network, the network needs to be optimized for the kinds of storage best practices that the storage administrator knows and manages already. Regardless of who is assigned ultimate ownership, the storage administrator should be involved throughout the process and should assist with storage network design best practices and optimization.

Logistics

From planning, to sizing, to racking and scaling, the decision-making process involves accounting for certain logistics up front. Though this seems obvious at first glance, it is easy to miss important facts that can affect outcomes. Be aware that these logistical items can cross disciplines as well, requiring a team effort to implement. It is outside the scope of this paper to cover all of these topics, but here are some basic questions to ask:

- **Storage:** Which storage vendors are being used in the mixed IP storage network environment?
- **High-value storage:** What is the high-value, high-performance storage already in place in the environment, and which business-critical applications or SLA policies does it support?
- **Storage performance limitations:** Is any of the high-performance storage limited by the current network performance?
- **Available rack space:** What is the optimal placement for this scale-out dedicated IP storage network? Is appropriate rack space available?
- **Management tools:** Which management tools are in use today? Which tools will have, or require, visibility into the new dedicated IP storage network?
- **SLA policy:** What policy changes are needed for the new dedicated IP storage network to start receiving new VMs? Which business functions or applications are important and need to be ensured?
- **Current deployment topology:** Do you already use consolidated server, VM, and storage pools, or do you manage separate pods for line-of-business needs or physical location reasons?

Best Practices

While the storage administrators should be familiar with general storage best practices, it is also time to start collecting storage best practice recommendations from equipment, software, and service vendors. The following is a partial list of companies that recommend that IP storage traffic use a dedicated physical network, along with their recommendations:

- EMC, Clariion Best Practices for Performance and Availability; Release 30.0 Firmware Update, Applied Best Practices—Chapter 3: Network Best Practices: http://goo.gl/VY0jbr

- Apache CloudStack Recommendation: http://goo.gl/K0nKVy
- VMware, Best Practices for Running VMware vSphere on iSCSI—Security Considerations, Private Network: http://goo.gl/PKD7cT
- Citrix, Citrix Zen Server Design: Designing Zen Server Network Configurations—Chapter 6: Designing Your Storage Network Configuration: https://goo.gl/E2SOGU
- IBM, Redbook, b-type Data Center Networking: Design and Best Practices Introduction—Chapter 9: IP Storage Area Networks, 9.4.4 Challenges in iSCSI SANs: https://goo.gl/ykGE1D

Architecture

As illustrated in Figure 1, the concept of the dedicated IP storage network is simple enough. A basic switch architecture is set up in parallel with the general-purpose IP LAN, with the storage attached to the IP storage network and the servers attached to both. (See Figure 1.)

The change seems simple, but the network performance enhancements, management simplification, automation, and new increased resiliency are life-changing in the data center. The introduction of Brocade VCS Fabric technology provides a flatter, multipath, and deterministic mesh network architecture that is ideal for ensuring data service to a scaled-up virtualized application environment.

The Brocade Dedicated IP Storage Network Best Practices

For the best performance, predictability, and reliability, the dedicated IP storage network based on Brocade VCS Fabric technology meets the following criteria:

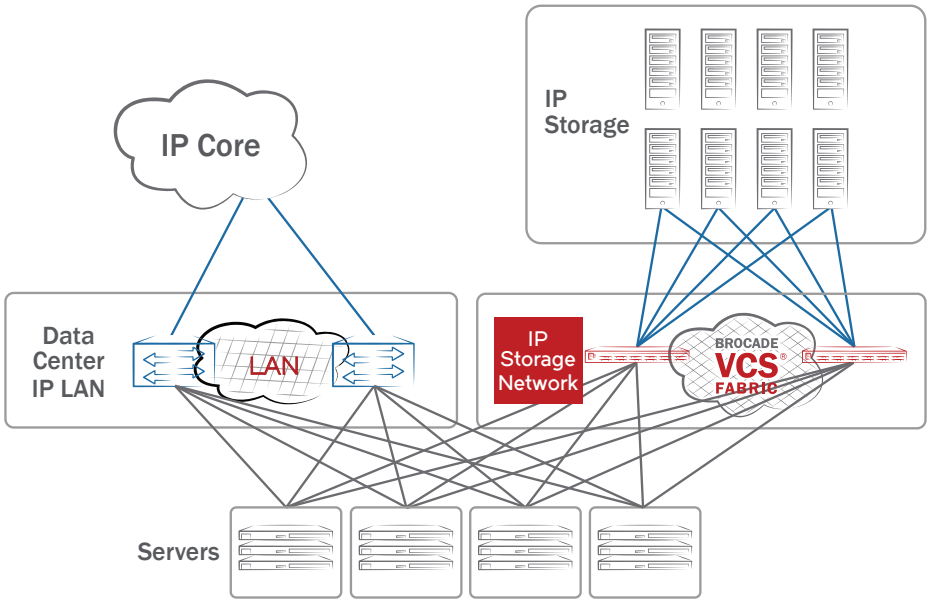

Figure 1. The dedicated Brocade IP storage network.

- **Predictable performance:** To achieve the most predictable performance and optimal data transport, the network based on Brocade VCS Fabric technology is redundant and nonblocking at physical, data link, and network layers (L1 to L3). Switches have full line rate capability and deep enough buffers for spikes in traffic, while delivering packets at minimum and consistent latency.
- **Highly automated and simple-to-deploy:** Brocade VCS Fabrics are self-provisioning and self-healing, to drive down the cost of services and ongoing maintenance. All links are active, which drives up utilization, and traffic is rerouted automatically so that the network is efficient and resilient.
- **Purpose-built for next-generation data centers:** Chassis-based High Availability (HA), In-Service Software Updates (ISSU), and available fixed-configuration redundant power supplies and fans provide the HA required for business continuity. Important features such as ISSU allow maintenance to be performed without taking switches offline, which keeps services highly available at all times.

Brocade VCS Fabric technology, included with the Brocade VDX Family of switches, provides the performance, predictability, and reliability needed to meet today's stringent IP storage network demands.

Key Technologies

This section highlights key concepts and technologies that underpin the solution architecture and provides links to more information.

Brocade Network Advisor

Brocade Network Advisor helps proactively manage end-to-end network health and performance and aids troubleshooting. Administrators can quickly identify network issues with

customizable dashboards and can drill down to isolate and fix problems. Brocade Network Advisor supports the entire Brocade IP and SAN portfolio, for unified network visibility and control.

For more information, go to: http://www.brocade.com/products/all/management-software/product-details/network-advisor/index.page

vRealize Integration

See and manage the virtual and physical networks together with vRealize integration. VMware vRealize Suite is a cloud management platform that is purpose-built for the hybrid cloud. It provides a comprehensive management stack for IT services on vSphere, other hypervisors, the physical infrastructure, and external clouds, all with a unified management experience.

For more information, go to: https://goo.gl/4y5NM

Brocade VCS Fabric Technology

Brocade VCS Fabric technology provides today's data centers with unmatched automation, efficiency, and agility. Simple, centralized management with zero-touch provisioning helps create highly scalable, flexible, and resilient environments. Brocade VCS Fabrics deliver a software-driven infrastructure to enhance IT agility and enable organizations to transition smoothly to elastic networks in highly virtualized and cloud environments.

As illustrated in Figure 2, there are a number of advantages in moving from a classic hierarchical architecture to an architecture based on Brocade VCS Fabric technology.

For more information, go to: http://www.brocade.com/solutions-technology/technology/vcs-technology/index.page

Port Profiles

In classic Ethernet, network policies and security control are assigned at the physical switch port. They are applied to all traffic entering the switch port and are tied to the physical port. Traffic is identified by the MAC address of the sending device, so a network policy includes a MAC address.

A port profile separates network policies from the physical port, acting as a portable container for a variety of policies. With Brocade VCS Fabric technology, a port profile can include one or more of the following profiles: Virtual LAN (VLAN), Quality of Service (QoS), Fibre Channel over Ethernet (FCoE), and security. In a VCS fabric, all switches and ports access a common port profile database, so they

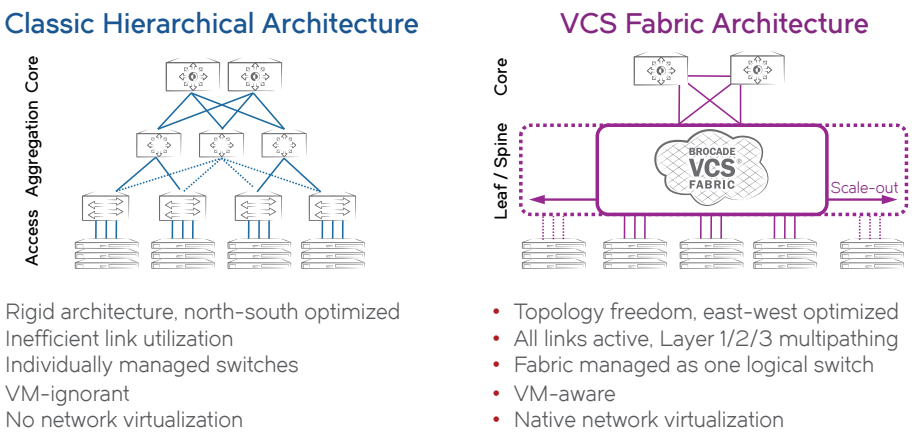

Figure 2. Traditional architecture vs. VCS Fabric architecture.

INTEGRATED OPERATIONS MANAGEMENT

Integrating with preferred tools that teams use to manage environments is key to helping the operation run smoothly. The integration of Brocade offerings with VMware offerings brings physical and virtual storage metrics to you with vRealize integration, simplifying visibility and troubleshooting:

- **VMware vRealize operations:** Provides an interface for management of a VM farm, offering visibility and actionable information to ensure SLAs, optimum resource usage, and configuration compliance.
- **Brocade content packs for vRealize operations:** Integrates with vRealize to provide storage network visibility and critical alerts that simplify management for virtual and physical environments, including Brocade VCS Fabric technology, and aid in problem resolution.
- **Brocade Network Advisor:** Complements vRealize operations with deeper storage network management and investigation for root-cause analysis and remediation.

automatically and consistently apply the correct set of port profiles to a VM MAC address, even when that VM moves to another port on a different switch in the fabric. Network management issues limiting VM mobility are eliminated.

Ports on Demand

A number of Brocade models can be purchased with a Ports on Demand (PoD) licensing feature, allowing companies to pay for the ports they use and scale as needed. PoD is ready to be unlocked in the switch firmware. Its license key may be part of the licensed paperpack supplied with switch software, or you can purchase the license key separately from your switch vendor. You may need to generate a license key from a transaction key supplied with your purchase.

Each PoD license activates the next group of ports in numerical order in either 4-port or 8- or 12-port increments, depending on the model. Before installing a license key, you must insert transceivers in the ports to be activated.

Brocade Switches

Industry-leading Brocade switches are the foundation for high-performance connectivity in storage, IP, and Ethernet fabric network environments. These highly reliable, scalable, and available switches are designed for a wide range of environments, enabling a low Total Cost of Ownership (TCO) and fast Return on Investment (ROI).

Switches to consider for an IP storage network Include:

- Brocade VDX 6740, 6740T, and 6740T-1G Switches
- Brocade VDX 6940 Switch
- Brocade VDX 6740, 6740T, and 6740T-1G Switches

For more information, go to: http://www.brocade.com/products/all/switches/index.page?network=ETHERNET_FABRIC

Configuration

The base recommendation is a four-switch design with a leaf-spine topology VCS fabric. This configuration is highly scalable and with PoD capability, it is possible to activate a subset of ports on each switch, rather than over-buying up front, and still scale up nondisruptively. This results in a resilient fabric of the right number of ports plus room to grow when necessary. The base of four switches also sets the ability to nondisruptively grow further, by adding pairs of switches as needed with automated configuration. (See Figure 3.)

Sizing Considerations

The number of switches to deploy and the number of ports to make active in a deployment depends on factors unique to each environment, but there are best practices and general guidelines to consider.

Determining the Maximum Number of Active Ports

In order to determine the number of storage and server ports to activate, first look at the number of storage and server connections your physical topology will require. That total is the maximum number of active ports you need right now. Your remaining choices might allow you to decrease that number.

Determining Storage Port Oversubscription

The next decision that might decrease the number of ports required is the level of oversubscription at the storage ports. This will be a function of the criticality and performance requirements of the workloads running on your VM farm. The following workload examples are general large/medium/small illustrations:

- **On-Line Transaction Processing (OLTP, or transaction-oriented processing) workloads:** These require the most bandwidth and also are generally the most business-critical workloads. The general practice for OLTP workloads is zero oversubscription, so 1:1 Top of Rack (ToR) to storage ports.
- **Virtual Desktop Infrastructure (VDI):** This is a "medium" example, where workloads require continuous connectivity and predictable performance, but port economy choices can be made. The general practice is 2:1 oversubscription, or one storage port per two ToR connections.
- **Virtual Server Infrastructure (VSI):** This is the "everything else" category, where the most tradeoffs are generally made. The level of oversubscription is a choice you can make based on the concentration of workloads, VMs, and racks vs. data paths. It could exceed 2:1.

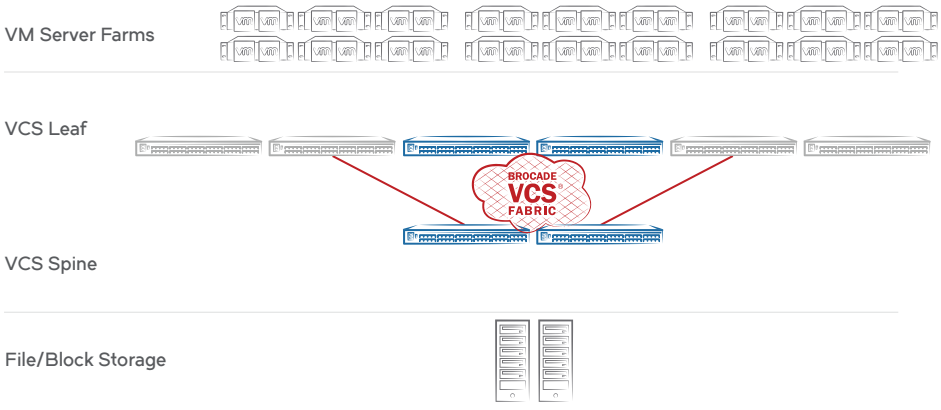

Figure 3: Scaling the Base Configuration

Racking Considerations

Once the base number of currently required active ports is determined, other general factors come into play:

- **ToR vs. Middle-of-Row (MoR):** ToR (with a switch in the server rack) is a cleaner deployment with less cabling complexity. MoR (with a separate switch that has the potential to be larger, with more available ports overall) may be appropriate if you want fewer switches or if all of your servers are not on the dedicated storage network.
- **Server rack density:** If a rack is carrying all critical workloads or is heavily thin-provisioned and requires high bandwidth, then extra Inter-Switch Links (ISL) ports should be allocated to increase bandwidth to the spine.

Configuration Best Practices

While the previous section contains guidance for determining the base topology and number of ports required to connect current known workloads, total port count and growth planning require you to consider a few more best practices.

1. **Paired redundancy:** Deploy switches in pairs and allocate a minimum of two ports each to ISLs for datapath redundancy and bandwidth to the spine. In addition, while not strictly necessary, often it is helpful to add two ISLs between pairs of ToR, since this safeguards against datapath disruptions for VMs in each rack (in case vmkernel uplinks are misconfigured and one ToR loses connection to the spine).
2. **Operational flexibility:** Never plan to use more than 85 percent of available ports. This leaves a pool of available ports for adding more server or storage connections on the fly, or for increasing bandwidth at the spine by increasing ISLs if necessary.
3. **Planned lifecycle:** Size the total port count—the number of switches and ports enabled—based on the planned lifecycle of the environment. This most likely is driven by planned server rack replacement and related VM migrations, but it can also coincide with storage capacity additions.
4. **Switch scaling:** In order to maintain the automation and resiliency of the fabric as it grows, it is a best practice to add physical switches in pairs, connected by ISLs.

Deployment

As always, the optimal deployment model depends on the particulars of the VM farm environment. However, the most nondisruptive, "keep it simple" plan is to build the dedicated IP storage network in parallel to the existing general-purpose LAN and use it for all new VMs, leaving the existing mixed infrastructure intact. As VMs, servers, and storage reach end of life, or as conditions allow, they can be systematically moved to the new dedicated IP storage network until all IP storage is consolidated on the dedicated IP storage network. During the course of the network migration, the dedicated IP storage network can be scaled as needed to support the additional traffic.

Other Deployment Considerations

With the initial base configuration deployed in parallel to the old mixed network, there are some further considerations before finalizing this deployment.

Nondisruptive Migration

The safe and easy approach works well in most migrations. Other motivations, requiring a quicker migration, also exist for moving to a dedicated IP storage network. In these cases, migration can still be done nondisruptively but it will likely require planning. VMs must be migrated off their associated storage volumes before the arrays can be migrated to the new dedicated IP storage network. The servers hosting VMs will be connected to both the shared and dedicated networks simultaneously, so VMs can be migrated back once the storage volumes are back online in the dedicated environment. Every storage array will have best practices of its own for migration and optimization on a new network so please also refer to your storage vendor's best practices guide.

Business-Critical Prioritization

While some VM farm workloads are better suited for the wholesale migration of all applications, others may need a more selective approach. By moving only business-critical applications to the dedicated IP storage network, you can manage SLAs even more tightly. Determining what is business-critical is subjective, but you can use the "litmus test" of considering which workflows would cause a productivity loss for lines of business if they were disrupted. It is common to think of revenue systems and disaster recovery as mission-critical company functions, but what about collaboration tools, file sharing, or even email (or email for certain groups)? The business will not stop if SLAs are missed, but whole teams and projects will be slowed down by latency and performance.

Problem Applications

When latency issues are suspected, you can often quickly fix troubled applications by simply moving to the new dedicated IP storage network. Even if the slower wholesale approach is chosen to just move new VMs to the new network, you can move troubled VMs as well to fix latency issues as they arise.

High-Value, High-Performance Storage

New-generation storage arrays are pushing the limits of the general-purpose IP network by boasting huge performance and scalability leaps. For this reason, they are more likely to be constrained by network inefficiencies and should be moved immediately to the new dedicated IP storage network. This being said, such arrays often were purchased to support higher SLA standards and received their workloads through policy-based selection. For this reason, they may require more

planning than you would apply to a commoditized storage unit where the business impact might be less.

Training

It is understood that network administrators are on call and rarely feel that they have extra time for training. Yet Brocade makes use of VCS fabric automation and standard Command-Line Interface (CLI) availability, such that most network administrators can be brought up to speed within 10 minutes of hands-on experience. At most, training takes less than a day through a modular online training class. However, if ownership is turned over to these storage administrators, then they will use the same topology and Brocade management tools that they already use to manage their Brocade SAN.

Final Thoughts

To deliver on increasingly stringent application SLAs and realize the full performance of the new generation of IP storage arrays, it is increasingly apparent that mixed-network traffic on general-use corporate networks has become the new bottleneck in the data center. Moving to a modern dedicated IP storage network quickly and cost-effectively raises the bar to the next level and is necessary to scale seamlessly into the future.

For More Information

To learn more about making the move to a Brocade dedicated IP storage network, please consult your local Brocade representative.

Visit the Brocade website at:
http://www.brocade.com/solutions-technology/enterprise/ip-storage/index.page

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

BROCADE