WHITE PAPER

Fibre Channel Storage Area Network Design Considerations for Deploying Virtualized SAP HANA

Business Case

Enterprises are striving to virtualize their mission-critical applications such as SAP HANA, Oracle Applications, and others to meet their goals in terms of performance, availability, and cost. The benefits of virtualization help organizations optimize and utilize resources, but they also include new challenges. Often, organizations that move to virtualized environments are still relying on storage systems that were not designed with the flexibility and ease of deployment and management that a virtualized data center requires. As a result, their application performance is impacted and they are forced to sacrifice efficiency and deployment time for the agility and cost advantages that virtualization offers. To compensate, organizations commonly over-provision storage to address latency and application performance, however this introduces increased costs and does not solve the need for the scalability, speed, and resilience required by virtualized environments. Enterprises need a storage and networking solution that is easy to deploy and user-friendly while also enabling the ability to support application performance and evolve as the environment scales.

Solution Overview

This solution provides a reference architecture to deploy virtualized SAP HANA in a VMware environment based on VMware vSphere and VMware NSX, offering customers agility, resource optimization, and ease of provisioning. This solution enables SAP customers to provision SAP HANA instances faster and more effectively by using vSphere Virtual Machines (VMs). VMware selected Brocade, a leading technology partner, to provide the products and architecture for a Fibre Channel-based Storage Area Network (SAN) together with Pure Storage, another leading technology partner, to provide an all-flash high-performance storage array. This work involved defining and testing a VMware Software-Defined Data Center-based reference architecture to deploy virtualized SAP HANA on SAP HANA Tailored Data Center Integration (TDI)-Certified hardware. Definition and testing of this solution followed the process and best practices as per Application Workload Guidance (formerly VMware Validated Design). This white paper provides the best practices and architecture to deploy the Fibre Channel SAN used in this solution.

The design details of the SAN solution is based on Brocade® Gen 6 Fibre Channel products to satisfy the storage requirements as specified by this solution. Please refer to reference section at the end of this paper for more detailed information about this solution.

Scope

This solution focuses on virtualized SAP HANA with a TDI environment. Best practices for virtualizing SAP HANA for scale-up and scale-out configurations are utilized in this design. The solution is sized, tested, and validated on SAP-certified partner hardware for different use cases. This reference architecture provides information specific to SAP HANA and leverages existing VMware Validated Designs for general-purpose guidelines.

The physical architecture includes Dell R630 and R730 PowerEdge Servers with up to 1 TB memory, a Pure Storage TDI array, and Brocade SAN and network components.

Use Cases

This design is targeted for the following use cases:

- Rapid provisioning to increase timelines for testing, training, and SAP upgrades
  - Clones SAP systems quickly and easily
- Enhanced application availability
  - High out-of-the-box availability without complex configuration and setup to protect against VMware ESXi server failure; no downtime due to ESXi server maintenance
  - Continuous utilization monitoring across resource pools and intelligent allocation of available resources among VMs, based on predefined rules that reflect customer business needs and changing priorities
  - VMware vSphere Fault Tolerance for SAP Central Services
- Server consolidation and data center energy cost reduction; runs SAP in VMs consolidated onto fewer physical servers that use less energy overall
- Unified performance monitoring of the virtualized SAP environment; coordinates performance reporting and analytics across the solution infrastructure stack, including the SAP and database tier, the guest operating system, the hypervisor layer, and storage
- Enhanced security that leverages VMware NSX; enables restricted communication between application components

Audience

This white paper is intended for administrators, architects, and System Engineers who are responsible for designing and deploying infrastructure for enterprise mission-critical applications.

Design Overview

A prototypical enterprise SAP customer with example requirements is used as a basis for all design activities in this solution. Requirements, assumptions, and constraints are carefully logged so all logical and physical design elements can be easily traced back to their source and justification.

Storage Requirements for SAP S4/HANA

All SAP HANA worker VMs have a database log, data, root, local SAP, and shared SAP volume. The storage capacity sizing calculation of these volumes is based on the overall amount of memory needed by SAP HANA's in-memory database. SAP has defined very strict performance KPIs that should be met when configuring a storage subsystem. With conventional arrays of spinning disks, this often resulted in low storage utilization in order to achieve required IO performance and latency. Incorporating all-flash arrays dramatically changes both IO performance and latency resulting in increased demand for a high-performing SAN.

SAP has published guidelines for calculating and sizing the correct storage requirements for SAP HANA TDI configurations. These guidelines change from time to time; therefore, it is always recommended to download the latest SAP HANA storage requirements white paper. It is also recommended to involve the storage vendor offering an SAP HANA TDI storage solution for SAP HANA Scale-Out configurations to ensure they have followed the correct space and IP sizing guidelines.

The SAP HANA nodes connect to the storage arrays through a Fibre Channel SAN. SAN components require a minimum 8 Gbps link speed, and the SAN topology must follow best practices for all redundant components and links. The connectivity, which includes host HBAs, SAN ports, switches, and array front-end ports, requires careful planning. 16 Gbps Fibre Channel (or faster) is recommended for the best performance and all Fibre Channel ports can negotiate to lower speeds.

Table 1: Shared Storage Physical Design Specifications.

Attribute	Specification
Vendor and model	Pure M50 Storage Array
Type	Active/Active
ESXi host multi-pathing policy	Round Robin
Minimum/maximum speed rating of switch ports	16 Gbps

Building the Storage Area Network

The SAN used for this SAP S4/HANA Solution is based on the Brocade G620 switch, though it is fully interchangeable with other switches in the Brocade Gen 6 Fibre Channel portfolio depending on the required scale. The entire Brocade Gen 6 portfolio also includes Brocade Fabric Vision technology, which provides unprecedented insight and visibility across the storage network. With enterprises migrating to high-performance flash storage for their mission-critical applications, Brocade Gen 6 switches are best suited to build the Fibre Channel SAN to support applications such as SAP S4/HANA.

Brocade Gen 6 is backward-compatible with previous generations providing 16 Gbps and 8 Gbps connectivity. While this solution is tested with a Pure Storage M50 array connected with 16 Gbps speed, there are features listed below unique to Gen 6 FC that benefit this deployment.

- Provides high scalability in an ultra-dense, 1U, 64-port switch to support high-density server virtualization, cloud architectures, and flash-based storage environments
- Increases performance for demanding workloads across 32 Gbps links and shatters application performance barriers with up to 100 million IOPS
- Provides proactive, non-intrusive, real-time monitoring and alerting of storage IO health and performance with IO Insight, a Gen 6 platform capability that provides built-in instrumentation to directly measure device IO performance
- Enables VM visibility in a storage fabric to monitor VM performance, identify VM anomalies, and optimize VM performance
- Increases resiliency by automatically discovering and recovering from device or network errors
- Leverages Brocade Fabric Vision™ technology to simplify administration, quickly resolve problems, increase uptime, and reduce costs

Fibre Channel SAN Design Considerations

When designing the SAN for an all-flash array, understanding the application workloads and the intended scalability, redundancy, and resiliency requirements are the main factors to consider.

When deploying a SAN, the main design consideration is the adequate sizing of the ISLs between the edge switches where the servers are connected and the core switches where the storage arrays are connected. In implementations where the storage arrays are deployed to serve specific latency- and IO-intensive applications such as SAP Hana, Oracle and other business-critical systems connecting both servers and Fibre Channel storage array to the core backbone switch(es) can be advantageous.

When deploying dedicated SANs for Fibre Channel storage array services, the SAN design can be tailored directly based on the application workloads and scalability requirements.

Storage Performance and Scaling

A best practice for SAN infrastructure design is to plan for a three- to five-year lifespan for the solution to be deployed. Considerations for the length of the life of a SAN infrastructure include a combination of:

- Equipment depreciation considerations
- The limited predictability of business transformation/development
- Technology refresh and improvements

Having clear indications or understanding of application needs over the same period determines the range of scalability and flexibility necessary in the design.

SAN Topology Redundancy and Resiliency

An important aspect of SAN topology is the resiliency and redundancy of the fabric. The main objective is to remove any single point of failure. Resiliency is the ability of the network to continue to function and/or recover from a failure, while redundancy describes duplication of components, even an entire fabric, to eliminate a single point of failure in the network. Brocade fabrics have resiliency built into Brocade FOS, the software that runs on all Brocade B-Series switches, which can quickly "repair" the network to overcome most failures. For example, when a link between switches fails, FSPF quickly recalculates all traffic flows. Of course, this assumes that there is a second route, which is when redundancy in the fabric becomes important.

The key to high availability and enterprise-class installation is redundancy. By eliminating a single point of failure, business continuity can be provided through most foreseeable and even unforeseeable events. At the highest level of fabric design, the complete network should be redundant, with two completely separate fabrics that do not share any network equipment (routers or switches).

Servers and storage devices should be connected to both networks utilizing some form of Multi-Path IO (MPIO) solution that enables data to flow across both networks seamlessly in either an active/active or active/passive mode. MPIO ensures that if one path fails, an alternative is readily available. Ideally, the networks would be identical, but at a minimum they should be based on same-switch architecture. In some cases, these networks are in the same location. However, to provide for disaster recovery, two separate locations are often used, either for each complete network or for sections of each network. Regardless of the physical geography, there are two separate networks for complete redundancy.

In summary, recommendations for the SAN design are to ensure application availability and resiliency via the following:

- Redundancy built into fabrics to avoid a single point of failure
- Servers connected to storage via redundant fabrics
- MPIO-based failover from server to storage
- Redundant fabrics based on similar architectures
- Separate storage and server tiers for independent expansion
- Core switches of equal or higher performance compared to the edges
- The principal switch defining the highest performance in the fabric

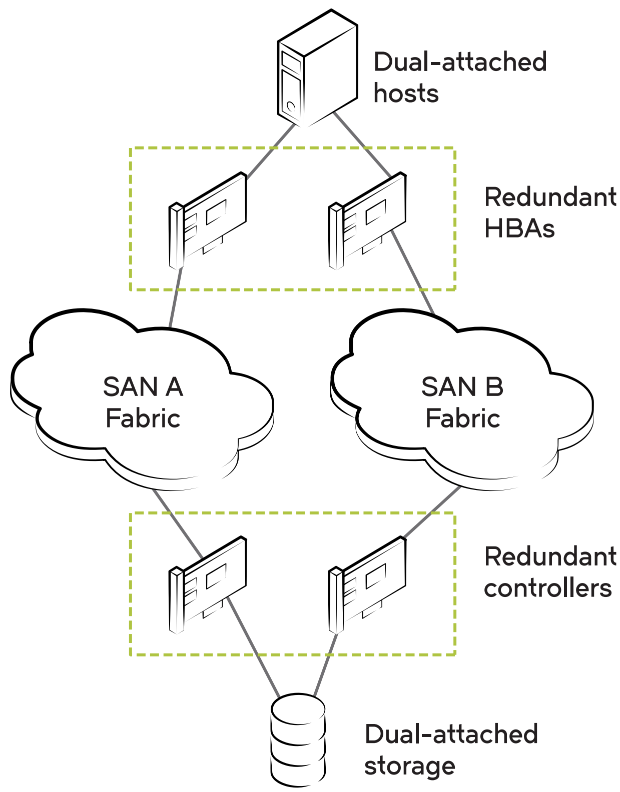

Figure 1: Connecting devices through redundant fabrics.

SAN Design Guidelines

Designing SANs fully dedicated with Fibre Channel storage array using the three SAN designs reviewed in the previous section is often a straightforward sizing exercise. Calculate the ISL bandwidth necessary for the server density and number of arrays in each Fibre Channel storage array cluster combined with the anticipated server and storage growth during the lifecycle of the computing infrastructure.

- For core-edge designs, a good rule of thumb is to design with ISL bandwidth equal to total capacity of all array ports which are accessible (zoned) to hosts attached to each ToR switch.
- For collapsed SAN designs, calculate the number of host and storage ports for the lifetime of the infrastructure.

Brocade Best Practice for Zoning

Zoning is a fabric-based service in SANs that groups host and storage nodes need to communicate. It only allows nodes to communicate with each other if they are members of the same zone. Nodes can be members of multiple zones, allowing for a great deal of flexibility when implementing a SAN using zoning.

Zoning not only prevents a host from unauthorized access of storage assets, but it also stops undesired host-to-host communication and fabric-wide Registered State Change Notification (RSCN) disruptions.

Brocade recommends that users always implement zoning, even if they are using LUN masking. Also, PWWN identification for zoning is recommended for both security and operational consistency. For details, please refer to the Brocade SAN Fabric Administration Best Practices Guide.

Fibre Channel SAN Availability and Reliability

This solution is using Fibre Channel SAN-based on Brocade Gen 6 technology with Fabric Vision technology and industry-standard 32 Gbps connectivity. The key features available in Fabric Vision technology and its integration with VMware vRealize suite help build a scalable and reliable solution architecture through automation and high availability.

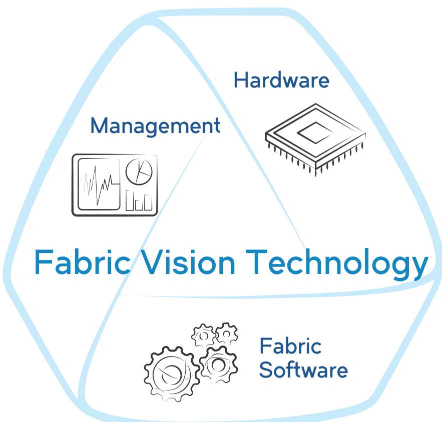

Figure 2: Brocade Fabric Vision Technology: Hardware, Software, and Management working together.

Automation

Monitoring and Alerts Policy Suite (MAPS): A policy-based monitoring and alerting tool that proactively monitors the health and performance of the SAN infrastructure based on pre-defined policies that cover more than 170 customizable rules, ensuring application uptime and availability. Administrators desiring a pristine network can set an "Aggressive" policy level that has rules and actions with strict thresholds to minimize the possibility of data errors. With the ability to tailor the MAPS policies, organizations can monitor all-flash ports closely to faster identify any performance degradation.

Availability

Configuration and Operational Monitoring Policy Automation Services Suite (COMPASS): An automated configuration and operational monitoring policy tool that enforces consistency of configuration across the fabric and monitors changes, simplifying SAN configuration and alerting about changes. In medium- to large-sized environments, this can prevent inadvertent changes to switch configurations that may impact the preferred parameters set across the fabric to optimize performance.

Buffer Credit Recovery: Automatically detects and recovers buffer credit loss at the virtual channel level, providing protection against performance degradation and enhancing application availability.

Forward Error Correction (FEC): Automatically detects and recovers from bit errors, enhancing transmission reliability and performance. FEC can reduce latency time significantly by preventing the need to retransmit frames with bit errors.

Simplified Monitoring of Mission-Critical Applications

Organizations face a constant struggle of both managing data growth and delivering actionable intelligence from raw data—all while meeting SLAs. As a result, even well-managed IT organizations must often make difficult choices about resource allocation, weighing the benefits of focusing more resources on monitoring, for instance, and fewer on planning or optimizing. With Fabric Vision technology, organizations can achieve unprecedented insight and visibility across the storage network through critical monitoring and diagnostic capabilities.

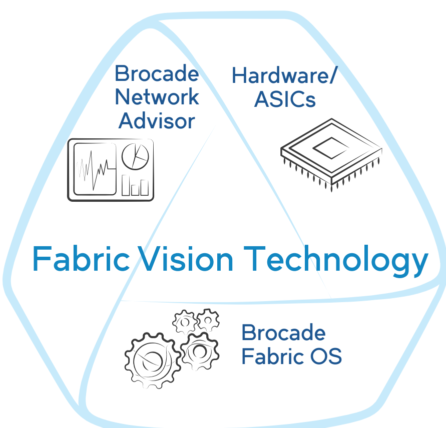

Figure 3: Brocade poducts behind Fabric Vision technology.

Dashboard: Includes customizable health and performance dashboard views, providing all critical information in one screen. Viewable dashboard "widgets" that should be monitored include errors on all-flash array facing ports, top 10 flows, memory usage, and port health.

IO Insight: Monitors individual host and storage devices to gain deeper insight into the performance of the network to maintain SLA compliance, obtain total IOs, implement first response time, IO latency (Exchange Completion Time, or ECT), and outstanding IOs performance metrics for a specific host or storage device., It also diagnoses IO operational issues and enables tuning of device configurations with integrated IO metrics to optimize storage performance (see Figure 4).

VM Insight: VM Insight uses standards-based, end-to-end VM tagging to gain VM visibility in a storage fabric. VM Insight enables storage administrators to monitor VM-level application performance and set baseline workload behavior. Using this information, they can quickly determine whether a storage fabric is the source of performance anomalies for applications running on VMs. VM Insight also enables fast correlation with other Fabric Vision technology metrics to identify the root cause of problems before operations are affected. VM Insight provides the visibility for administrators to provision and plan storage networks based on application requirements, and to fine-tune the infrastructure to meet service-level objectives (see Figure 5).

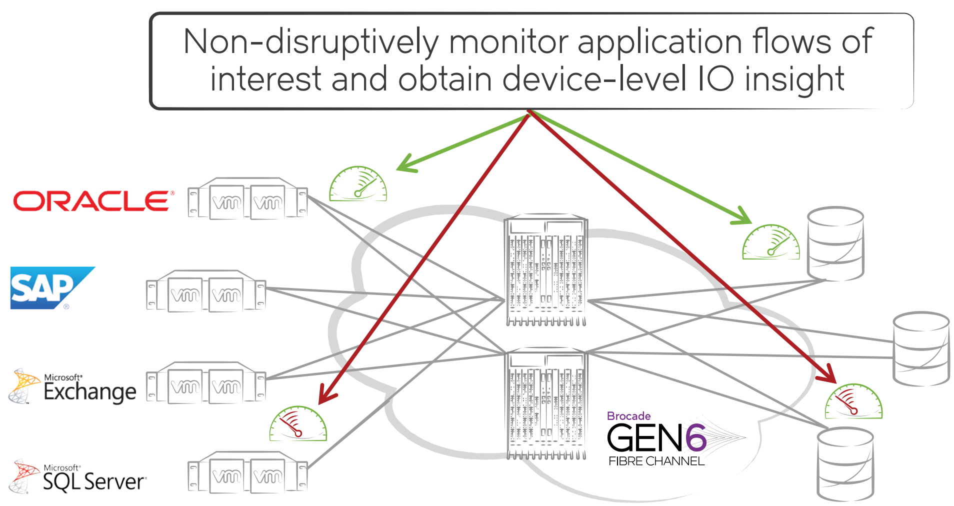

Figure 4: IO Insight: Monitor storage IO metrics to gain insights into operational issues.

Flow Vision: A comprehensive tool that allows administrators to identify, monitor, and analyze specific application and data flows to maximize performance, avoid congestion, and optimize resources. Flow Vision consists of:

- **Flow Monitoring:** Monitors specified traffic flows from source to destination through the SAN
- **Flow Generator:** Generates traffic between any two ports in a Gen 5 fabric
- **Flow Mirroring:** Captures packet data as it flows through the SAN, then displays and analyzes the captured packet's data

Flow Vision is most suited to temporary use while troubleshooting high-latency conditions, as continual use results in the collection of substantial amounts of diagnostic data. Flow Vision can be used as needed to verify optimal performance for the most demanding applications requiring optimal performance.

Fabric Performance Impact Monitoring: Identifies and alerts administrators to device or ISL congestion and high levels of latency in the fabric which can have a severe impact on all-flash array performance. FPI Monitoring provides visualization of bottlenecks and identifies slow-drain devices and impacted hosts and storage.

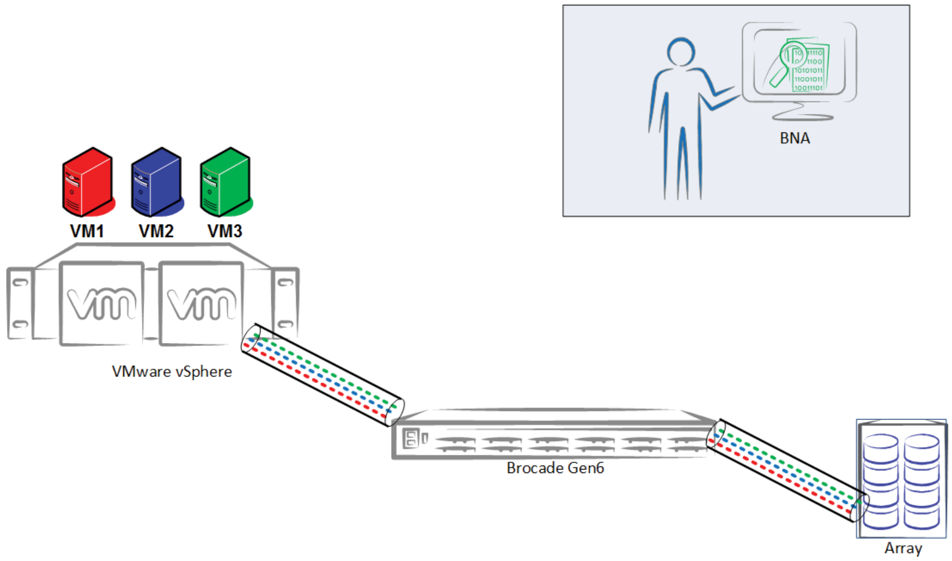

Figure 5: IO Insight delivering visibility of each VMs IO to shared LUNs.

ClearLink Diagnostics: Diagnostic Port (D_Port) provides loopback test capabilities for link latency and distance measurement at the optical and electrical level to validate the integrity and performance of optics and cabling, ensuring signal and optical quality and optimal performance across SAN and WAN connections. Pre-validating the integrity of cables and optics with ClearLink prior to deployment identifies potential support issues before they occur and enhances the resiliency of high performance fabrics. For all-flash array performance, any impurity in the physical infrastructure can impact performance.

VMware vRealize and Log Insight: Brocade has worked closely with VMware to offer deeper investigation capabilities for root-cause analysis and remediation of SANs in virtualized environments. The Brocade SAN Content Pack eliminates noise from millions of events and amplifies critical SAN alerts to accelerate troubleshooting with actionable analytics. Faster troubleshooting provides time to proactively drive the value of IT resources. Leveraging Fabric Vision technology provides thorough knowledge of the behavior and health of Brocade SAN fabrics over time to identify and remediate patterns impacting virtual machine performance and application responsiveness.

For enterprises where deeper application performance insight is necessary, the Brocade Application Monitoring Platform SAN appliance can measure application performance and traffic behavior in a Brocade Fibre Channel networks. Providing actionable intelligence to maximize performance across the data center and ensure application delivery SLAs are met.

For more information, see Brocade Applications Monitoring Platform.

Summary

Using SAP HANA in TDI deployments with SAP and VMware certified high-performance storage arrays such as all-flash storage arrays provides many benefits, including reducing hardware and operational costs, lowering risks, and increasing hardware vendor flexibility. But to benefit from these high-performance storage arrays, customers need a high-performance and highly available SAN based on newer generations of SAN technology such as Brocade Gen 6 SAN products and deployed using best practices. This will help customers to easily transition to this new application architecture using SAP HANA.

Following are some of the major components used by VMware in their lab to architect and test the solution:

- vSphere for virtualized compute and infrastructure
  - High-performance Dell R630 and R730 PowerEdge Servers
  - Pure Storage–based SAP TDI and Western Digital all-flash VMware vSAN storage
  - Brocade Generation 6 SAN fabric
- Separate vSphere clusters for management, SAP applications, and SAP HANA databases
- vSphere HA and vSphere FT for simplified high availability
- VMware NSX for networking with micro-segmentation for security on Brocade VDX switches
- vRealize Automation with Adapter for SAP Landscape Management
- vRealize Operations with Blue Medora SAP dashboards for operations and capacity planning

The solution was then successfully validated to meet and exceed the requirements. The results of the validation tests clearly show that vSphere offers a robust platform for running SAP S/4HANA production workloads.

Acknowledgements

Pure Storage
Krishna Satyavarupu

VMware
Mohan Potheri

Brocade
Marcus Thordal

Nanjunda Somayaji

Resources

For more information about the SAP HANA and VMware products and technologies discussed in this guide, review the links and references below:

The Winding Road to Virtual SAP HANA Application Workload Guidance Design for SAP S/4HANA on VMware vSphere 6.5

Application Workload Guidance and Design for Virtualized SAP S/4HANA on vSphere (Part 1, Part 2, Part 3, Part 4)

Getting Storage Right for Business-Critical Applications on VMware

Fit for the Future

The Modernization of Storage Architectures

Brocade Fabric Vision Technology

Brocade G620 Switch - Gen 6 Fibre Channel Switch

For more information about Brocade solutions, visit 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