

## **DATA CENTER**

# **Brocade SAN Interoperability with Cisco UCS**

Brocade, the leader in Fibre Channel SANs, supports the largest number of server solutions compared to all other SAN Vendors. All major server vendors support connectivity to a Brocade SAN without proprietary features and without sacrificing performance. Cisco UCS, a relatively new entrant to the server market, is also fully supported on a Brocade SAN, however there are some limitations due to the Cisco proprietary connectivity. This document outlines how Cisco UCS and Brocade work together.

# **BROCADE**

---

## CONTENTS

<b>Contents</b> .....	<b>2</b>
<b>Introduction</b> .....	<b>3</b>
<b>Cisco UCS in the Market</b> .....	<b>3</b>
<b>UCS Connectivity to SAN</b> .....	<b>3</b>
NPV/Access Gateway Mode of Operation .....	4
UCS Servers Connecting to the SAN .....	5
Automatic Mapping of Uplinks.....	5
Manual Mapping of Uplinks using PIN Groups .....	5
Port channels .....	6
Recovery after an Uplink Goes Down .....	7
<b>Summary</b> .....	<b>7</b>
<b>Appendix A – Creating UCS SAN PIN Groups with UCS manager</b> .....	<b>8</b>

## INTRODUCTION

Brocade SAN fabrics provide a secure, low-latency, and high-performance infrastructure for mission-critical applications on which Fortune 1000 companies depend, offering industry-leading Gen 5 Fibre Channel backbones, switches, and embedded switches. With more than 80 percent of the SANs deployed and more than 70 percent SAN marketshare shipping every quarter, Brocade is the preeminent leader in Fibre Channel data center fabrics, providing proven hardware and software reliability to support non-stop operations.

Brocade offers best-in-class SAN fabric features and functionality to ensure high performance, reliable, and highly available storage access, regardless of the server connecting to the Brocade SAN. Brocade's broad partner ecosystem and commitment to open standards ensures that applications hosted on servers connected to a Brocade SAN deliver the highest levels of performance without deploying any proprietary connectivity models. Brocade SAN fabrics support the largest number of server solutions in the industry, more than any other SAN vendor. In fact, all major server vendors support open standards-based connectivity to a Brocade SAN, without proprietary features and without sacrificing performance, offering our joint customers maximum benefit and choice.

Cisco UCS, a relatively new entrant to the server market, is also fully supported on a Brocade SAN, excluding Cisco proprietary features such as their non-standard port channeling technology. By using PIN groups on Cisco's UCS, however, a Brocade SAN user can achieve the same levels of performance as with using Cisco proprietary port channels. The rest of this document provides more details on how Cisco UCS and Brocade work together.

## CISCO UCS IN THE MARKET

Cisco's Unified Computing System (UCS) is Cisco's offering in the blade server market, which was first introduced in 2009. It is positioned as a data center platform, combining compute, network connectivity (LAN) and storage connectivity (SAN) in a single chassis<sup>1</sup>. UCS is positioned to consolidate both physical and virtual workloads onto a single, centrally managed and automated system for compute, network, and storage access. Cisco promises reduced IT effort due to unified management of all these components, especially in virtualized environments because VMs can be provisioned with compute, network, and storage resources all in one place. Within a few years, Cisco has captured more than 15% of the blade server market and is currently the third largest blade server vendor.

## UCS CONNECTIVITY TO SAN

Given Brocade's market leadership and ubiquitous presence in Data Center SANs and the growing popularity of Cisco UCS, it is clear that joint customers need interoperability between Cisco UCS servers and Brocade SANs. In fact, considering that more than 80% of the SANs deployed are Brocade SANs, there is a very high probability that the majority of Cisco UCS servers connect to Brocade SANs today without performance concerns.

Cisco uses CNAs within the UCS server architecture. The CNAs use the FCoE protocol as defined by the ANSI T11 Standards Committee to provide SAN connectivity. The FC traffic is encapsulated inside the ethernet frames. Instead of the buffer-to-buffer (B2B) credit system used in native fabric topologies, FCoE relies on a new set of Ethernet standards to enhance the ethernet protocol to provide lossless transport. UCS CNAs are then connected to a fabric interconnect (via fabric extender modules) to provide connectivity to SAN and LAN. Thus every server within Cisco's unified fabric technology is capable of connecting to both the LAN and SAN. Network connectivity to UCS servers is provided through the use of Cisco UCS 6100/6200 series Fabric Interconnects (FIs), which separates the LAN and SAN traffic and directs it to the corresponding upstream LAN and SAN switches.

---

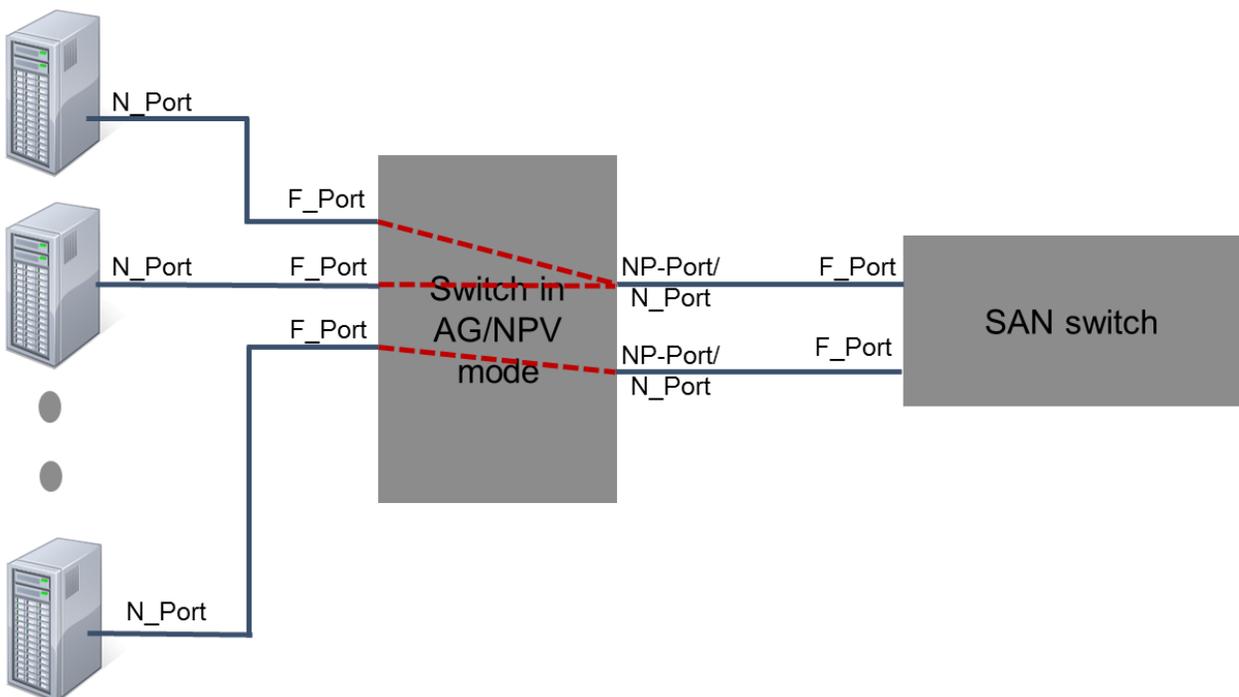
<sup>1</sup> <http://www.cisco.com/en/US/products/ps10265/index.html>

## NPV/Access Gateway Mode of Operation

The Fabric Interconnect operates in N-Port Virtualization (NPV) mode and not as a FC switch in the fabric for SAN connectivity (although switch mode can be used for specific deployments like direct attach or an all-Cisco SAN). Cisco's NPV offering is similar to Brocade Access Gateway (AG) offering. In the NPV/AG mode, the switch does not require a FC domain ID and keeps the number of domain IDs in the SAN fabric the same. To allow Cisco NPV or Brocade AG connectivity to the SAN, N\_Port ID Virtualization (NPIV) should be enabled on the ports in the core SAN switch.

NPV/AG multiplexes host connections to the fabric. The switch presents its downlink as an F\_Port to the host and an uplink N\_Port to an Edge fabric switch. Using NPIV, NPV/AG allows multiple FC initiators to access the SAN on the same physical port. Therefore, you can increase the number of hosts that have access to the fabric without increasing the number of switch domains. This simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports. This reduces the hardware requirements and management overhead of hosts to the SAN connections. The size of the fabric is limited to a maximum of 239 Domain IDs as specified in the ANSI T-11 standards; Brocade Access Gateway and Cisco's NPV mode of operation help scale beyond these limits. Operationally, the actual number of domain IDs (switches) in the fabric may be a lot smaller than the maximum limit.

The Fabric Interconnect in NPV mode joins the fabric through a normal FLOGI (Fabric Login) on the SAN uplink ports with NPIV. Subsequent logins from the server blades (FLOGIs originated from the server blade adapters) are translated by the NPV/AG switch into FDISC (Fabric Discovery) into the fabric. FDISC response from the fabric switch is then relayed back to the server as FLOGI response completing the server login process. If the login is successful, a data path is established between the downlink and uplink of the switch in NPV/AG mode. Switches in NPV/AG mode are logically transparent to the host and the fabric and must be connected to an upstream fully functional SAN switch that can provide full Fibre Channel services.



## Figure 2: High-Level Connectivity Diagram in NPV Mode<sup>2</sup>

As shown in Figure 2, an *NP port* (proxy N port) is a port on the NPV device and connected to the NPIV capable core SAN switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports that are connected to the downlinks (F\_ports) of the NPV device. (Per Brocade terminology, these ports are referred to as N\_Ports, the same as N\_Ports from servers). In this mode, the F\_Ports are assigned to NP ports. The NPV device can support more than one NP port for SAN connectivity and one or more of the F\_Ports are assigned to an NP port or port channel. (A port channel is a group of several physical Fibre Channel ports to create one logical Fibre Channel link as discussed in detail below).

All Brocade SAN switches support NPIV and allow SAN connectivity to Cisco switches in NPV mode<sup>3</sup>. In the context of Cisco UCS, vHBAs correspond to the server or F\_port.

## UCS Servers Connecting to SAN

The Fabric Interconnect (FI) operates in NPV mode and not as a FC switch in the fabric for connectivity to a Brocade SAN. In this mode, the downlink ports of the FI that are connected to servers are mapped to uplink ports of the FI that are connected to upstream SAN switches. The servers then use the mapped uplink ports to login to the fabric to access its storage. This mapping could be done automatically or manually by using SAN PIN groups depending upon the requirements. Although automatic configuration makes deployment easy, manual configuration using PIN groups provides more control in assigning certain servers to specific uplinks.

### Automatic Mapping of Uplinks

In automatic mapping, the UCS servers will be uniformly assigned using round robin to available uplinks so that the logins on uplink ports are evenly distributed. With this mapping in place, as soon as a new server comes online, it is automatically mapped to the uplink with the least number of logins. If the link fails, the servers on the failed link experiences traffic disruption and with round robin, the server is mapped to the next available uplink in the group and is brought back online. This mapping is intended for minimal management.

**NOTE:** Servers are mapped purely on login basis; the traffic pattern originating from the servers is not considered. For example, intense I/O generating servers can be mapped to the same uplink and cause congestion, even though additional uplinks are available that maybe underutilized.

### Manual Mapping of Uplinks using PIN Groups

In cases where more control is required to rate limit bandwidth and/or different servers need to access storage in different fabrics, PIN groups should be used. Pinning allows SAN traffic from a server (vHBA) to be manually mapped/pinned to an uplink Fibre Channel port on the Fabric Interconnect. Manual mapping enables an administrator to provide Quality of Service (QoS) and rate limit bandwidth by dedicating an uplink to a server or group of servers. This approach allows configuring certain uplink bandwidth for a group of servers and helps manage the distribution of traffic from the servers. The UCS server can be pinned only to one PIN? group at a time.

Once a Cisco UCS server (through its vHBA) is pinned to use a certain pin group for SAN connectivity, all traffic from the vHBA will be switched through the Fabric Interconnect on the specified uplink Fibre Channel port. If multiple servers are pinned to the same pin group, they are load-balanced among the uplink ports in the pin

---

<sup>2</sup> [http://www.cisco.com/en/US/docs/storage/san\\_switches/mds9000/sw/rel\\_3\\_x/configuration/guides/cli\\_3\\_4\\_x/npv.html](http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/sw/rel_3_x/configuration/guides/cli_3_4_x/npv.html)

<sup>3</sup> <http://www.brocade.com/solutions-technology/technology/compatibility/index.page>

group, just as with automatic mapping. If there is a link failure in this configuration, bringing the server back online will be limited to the number of available uplinks in that particular pin group (however only one uplink FC port is supported per pin group at this time).

### Port channels

A Fibre Channel port channel allows the administrator to group several physical Fibre Channel ports to create one logical Fibre Channel link to provide fault-tolerance and high-speed connectivity. You can create up to four Fibre Channel port channels in each Cisco UCS domain. Each Fibre Channel port channel can include a maximum of 16 uplink Fibre Channel ports. A port channel performs the following functions:

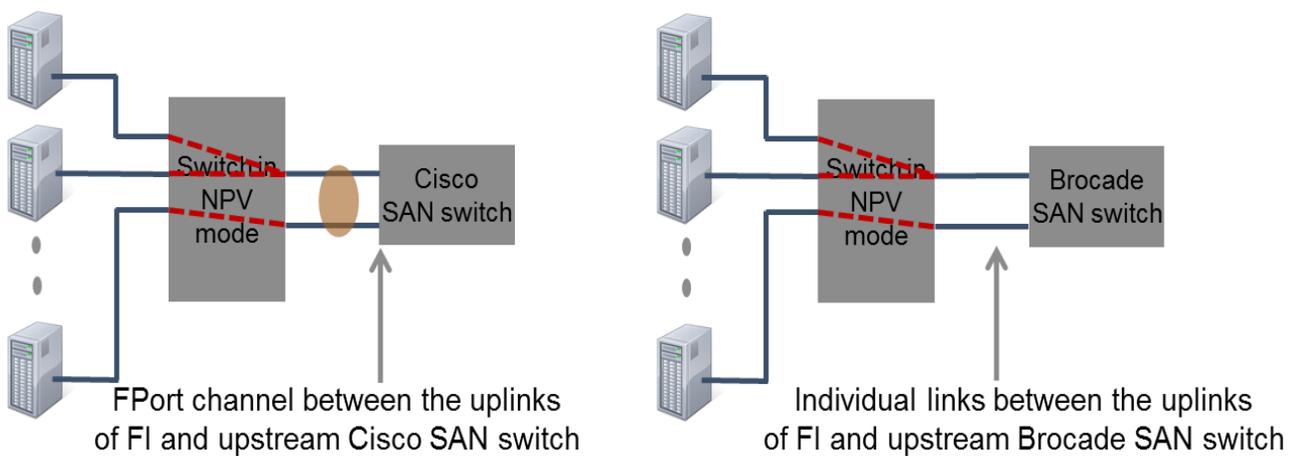
- Provides link redundancy and availability. For instance, if one of the links in the port channel fails, the traffic previously carried on this link is switched to the remaining links. The application is not impacted because of the link failure.
- Increases the aggregate bandwidth on a link by distributing traffic among all functional links in the port channel.

Cisco's Fibre Channel port channeling is non-standard and proprietary technology to Cisco, which can create vendor lock-in and is an interoperability concern for non-Cisco products. As shown in Figure 3, when uplinks of Fabric Interconnect are connected to Cisco SAN switches, the uplinks can form F port channels and will provide benefits of port channels such as link redundancy at the Fabric Interconnect level (multi pathing software can also be used to provide application redundancy, and is discussed later in this paper) and bandwidth utilization.

When uplinks are connected to Brocade SAN switches, however, port channels are not available. In the absence of port channels, one or more of the uplinks from the UCS server to the Brocade SAN may be overloaded and cause congestion due to uneven distribution of traffic on the uplinks. (As previously noted, uplinks are balanced purely on the basis of logins and traffic pattern is ignored).

This situation not only hurts customers, it hurts the Fibre Channel industry. Servers from all major server vendors connect to a Brocade SAN seamlessly without the use of any proprietary connectivity models. Cisco UCS should follow the rest of the industry, and adjust its proprietary port channel connectivity model to support a Brocade SAN.

**NOTE:** With automatic mappings in the Fabric Interconnect, if the link fails, the servers on the failed link will experience traffic disruption, however round robin kicks in and uses the next available uplink in the group to bring the server back online. As a result, link redundancy is still available with automatic mappings even in the absence of port channels.



**Figure 3: Connectivity between UCS Servers and Upstream Cisco vs. Brocade SAN Switch**

To avoid congestion that could occur when automatic uplink mapping is used, administrators can leverage manual mapping using SAN pin groups. A pin group could be created with a particular uplink and the server(s) with higher I/O requirements can be mapped to the uplink in the pin group. Pinning makes dedicated bandwidth available to the server(s) mapped to the pin group and only servers mapped to the pin group can use the uplinks in the pin group. This technique ensures QoS in terms of bandwidth for the servers that are pinned. One or more servers can be pinned to a given uplink in the pin group. The drawback of pinning is increased management overhead and server link recovery in case the uplink in the pin group goes down. The link will stay down until restored. Absence of port channels might elevate the concern for link redundancies between Fabric Interconnects and upstream SAN switches if a SAN pin group is used.

## Recovery after an Uplink Goes Down

As per storage best practices design, servers are dual homed with multi-pathing software to allow failover to an alternate path in case of link failure and loss of storage access. Multi-path I/O provides storage with improved fault-tolerance and performance by providing more than one physical path between the server and storage elements. A simple example would be a server with multiple HBA's attached via dual SAN fabrics to storage. If the path or storage element from one HBA fails, multi-path I/O software would allow seamless connection via the other path to the storage devices. Failure of a link will cause multi-pathing software to kick in and start using the secondary path from server to storage. UCS servers are also dual-homed onto two parallel and physically separate SANs: SAN A and SAN B. As per recommended practices, UCS will be deployed with two Fabric Interconnects and two separate SAN fabrics (Fabric A and Fabric B). Fabric Interconnect A will connect to Brocade SAN A, while Fabric Interconnect B connects to Brocade SAN B. Neither the Fabric Interconnects nor the fabrics are interconnected, providing two separate physical paths from the server to storage for redundancy. From the UCS server, in each Service Profile in UCS Manager you can then configure a vHBA for each fabric.

With a SAN pin group, multi-pathing software helps bridge the gap in case the uplink in the pin group goes down. For instance, an uplink that is going down will cause the server vHBA to go down. Multi-pathing software will then failover to the alternate vHBA, which will use the redundant path to access storage.

## SUMMARY

As the preeminent leader in Fibre Channel data center SANs, Brocade SAN switches are supported by all major server and storage vendors. Our commitment to driving open standards and developing break-through products based on those open standards ensures maximum interoperability, flexibility, and choice for our customers. As such, Brocade Fibre Channel switches support the NPIV standard to enable interoperability with third party SAN and Universal switches, including Cisco UCS Fabric Interconnects. Brocade has a large customer base that is successfully interoperating a Cisco UCS solution with a Brocade SAN infrastructure.

All server vendors, with the exception of Cisco, support an open protocol connectivity method that maximizes performance on a Brocade SAN, providing customers the maximum benefit from their server, storage, and network investments. Given the popularity and large installed base of Brocade SANs, customers have a right to expect that Cisco UCS offers a standards-based connectivity method that maximizes UCS performance on a Brocade SAN infrastructure. Anything less is simply an attempt at vendor lock-in, limiting customer flexibility and choice.

To find out more about the Cisco UCS interoperability with Brocade SANs, contact your sales representative or visit <http://www.brocade.com/solutions-technology/technology/compatibility/index.page> for more resources.

## APPENDIX A – CREATING UCS SAN PIN GROUPS WITH UCS MANAGER

To configure pinning for a server, SAN pin group must be included in a vHBA policy. The vHBA policy is then included in the service profile assigned to that server. All traffic from the vHBA will be switched through the I/O module to the specified Fibre Channel uplink port.

The steps for creating SAN pin groups within Cisco UCS Manager are outlined below:

Step 1: Create a SAN PIN Group from the SAN tab

The screenshot displays the Cisco UCS Manager interface. On the left, the navigation tree is visible, with 'SAN Pin Groups' highlighted under the 'SAN' tab. A red arrow points to this menu item. The main area shows the 'SAN Pin Groups' configuration page. A table lists the following:

Name	Port
SAN Pin Group UCS-Brocade-1	
Target A	fabric/san/A/phys-slot-2-port-1
Target B	fabric/san/B/phys-slot-2-port-1

A 'Create SAN Pin Group' dialog box is open, showing the following configuration:

- Name: UCS-Brocade-1
- Description: UCS Brocade Traffic Engineering Group 1
- Targets:
  - Fabric A: Interface: FC Interface 2/1
  - Fabric B: Interface: FC Interface 2/1

Step 2: Assign the SAN Pin Group to one of the vHBA's from Servers tab as shown below:

The screenshot displays the Cisco UCS Manager interface. On the left, the 'Servers' tab is selected, and the tree view shows a hierarchy of service profiles and vHBAs. The 'vHBA fc0' is highlighted with a red arrow. The right pane shows the configuration for 'vHBA fc0' under the 'Properties' tab. The 'Pin Group' dropdown menu is open, and a red arrow points to it, indicating the step of assigning a SAN Pin Group to the vHBA.

**Properties**

- Name: **fc0**
- WWPN: **20:00:00:25:B5:00:00:0F**
- WWPN Pool: **UCS\_WWP-1**
- WWPN Pool Instance: **org-root/wwn-pool-UCS\_WWP-1**
- Fabric ID:  A  B
- VSAN: **default**
- Owner: **Logical**
- Type: **Fc**
- Equipment: **sys/chassis-1/blade-1/adaptor-2/host-fc-2**
- Persistent Binding:  Disabled  Enabled
- Boot Device: **Disabled**
- Max Data Field Size: **2048**

**States**

- Operational Speed: **Line Rate**
- State: **Applied**

**Policies**

- Adapter Policy: **<not set>**
- Adapter Policy Instance: **org-root/fc-profile-default**
- QoS Policy: **<not set>**
- QoS Policy Instance: **<not set>**
- Pin Group: **<not set>**
- Stats Threshold Policy: **default**
- Threshold Policy Instance: **org-root/thr-policy-default**

**Virtual Host Interface Placement**

- Desired Placement: **Any**
- Actual Assignment: **1**

**Order**

- Desired Order: **4**
- Actual Order: **4**

Step 3: From the Pin Group pull down menu, select the Pin Group we created earlier:

The screenshot displays the configuration page for a vHBA interface named **fc0**. The interface is divided into several sections:

- Fault Summary:** Shows four status icons (Error, Warning, Info, Success) with a count of 0 for each.
- Actions:** A list of actions including "Change World Wide Port Name", "Bind to a Template", "Unbind from a Template", "Clear Persistent Binding", and "Reset WWPN Address".
- Properties:** The main configuration area containing:
  - Name: **fc0**
  - WWPN: **20:00:00:25:B5:00:00:0F**
  - WWPN Pool: **UCS\_WWP-1**
  - WWPN Pool Instance: **org-root/wwn-pool-UCS\_WWP-1**
  - Fabric ID: Radio buttons for **A** and **B**.
  - VSAN: **default** (dropdown)
  - Owner: **Logical**
  - Type: **Fc**
  - Equipment: **sys/chassis-1/blade-1/adaptor-2/host-fc-2**
  - Persistent Binding: Radio buttons for **Disabled** and **Enabled**.
  - Boot Device: **Disabled**
  - Max Data Field Size: **2048**
  - States: Operational Speed: **Line Rate**; State: **Applied**
  - Policies: Adapter Policy: **<not set>**; Adapter Policy Instance: **org-root/fc-profile-default**; QoS Policy: **<not set>**; QoS Policy Instance: **<not set>**; Pin Group: **SAN Pin Group UCS-Broca...** (highlighted with a red arrow); Stats Threshold Policy: **default**
  - Threshold Policy Instance: **org-root/thr-policy-default**
  - Virtual Host Interface Placement: Desired Placement: **Any**; Actual Assignment: **1**
  - Order: Desired Order: **4**; Actual Order: **4**

Step 4: Click the **Save Changes** in the bottom right hand corner.

© 2013 Brocade Communications Systems, Inc. All Rights Reserved. 05/13

ADX, AnyIO, Brocade, Brocade Assurance, the B-wing symbol, DCX, Fabric OS, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, and Vyatta are registered trademarks, and HyperEdge, The Effortless Network, and The On-Demand Data Center are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks 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.