WHITE PAPER

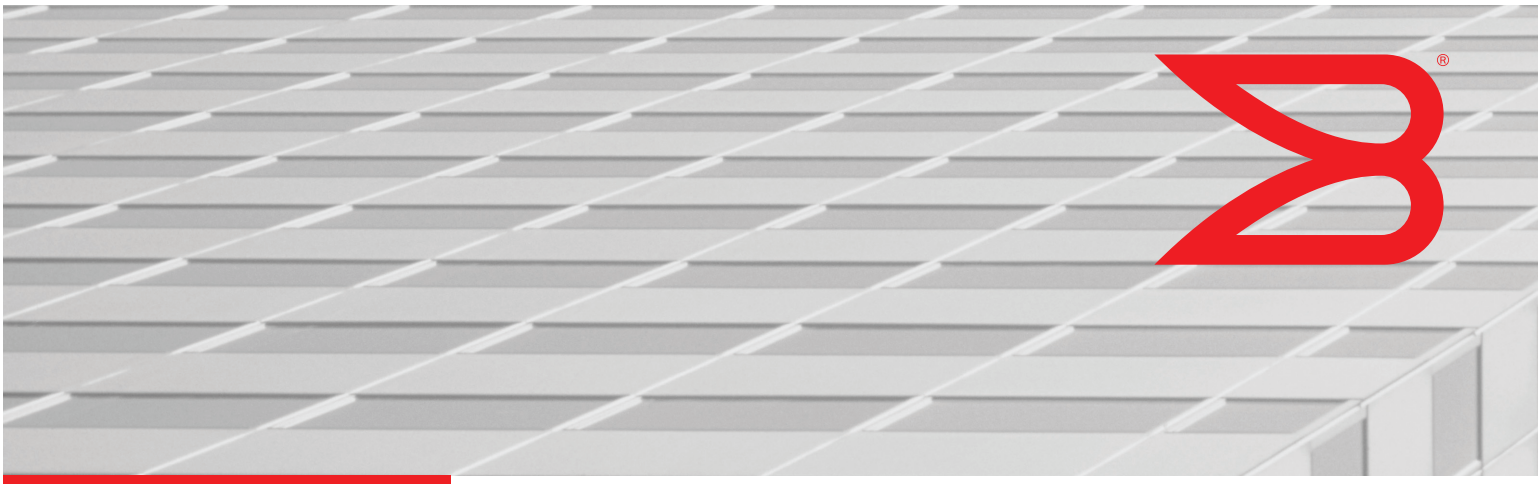

STORAGE AREA NETWORK

Secure SAN Zoning Best Practices

Hosts and storage in a Brocade SAN fabric can be easily secured using the Zoning best practices in this paper

BROCADE

This paper describes and clarifies Zoning, a security feature in Storage Area Network (SAN) fabrics. By understanding the terminology and implementing Zoning best practices, a Brocade® SAN fabric can be easily secured and scaled while maintaining maximum uptime.

The following topics are discussed:

- **Zoning defined and LUN security in the fabric**
- **Identifying hosts and storage members of a zone**
- **How do SAN switches enforce Zoning?**
- **Avoiding Zoning terminology confusion**
- **Approaches to Zoning, how to group hosts and storage in zones**
- **Brocade Zoning recommendations and summary**

WHAT IS ZONING?

Zoning is a fabric-based service in Storage Area Networks that groups host and storage nodes that need to communicate. Zoning creates a situation in which nodes can communicate with each other *only if they are members of the same zone.* Nodes can be members of multiple zones—allowing for a great deal of flexibility when you implement 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. RSCNs are managed by the fabric Name Server and notify end devices of events in the fabric, such as a storage node or a switch going offline. Brocade isolates these notifications to only the zones that require the update, so nodes that are unaffected by the fabric change do not receive the RSCN. This is important for non-disruptive fabric operations, because RSCNs have the potential to disrupt storage traffic. When this disruption was more common, that is, with older Host Bus Adapter (HBA) drivers, RSCNs gained an undeserved negative reputation. However, since that time most HBA vendors have addressed the issues. When nodes are zoned into small, granular groupings, the occurrences of disruptive RSCNs are virtually eliminated. See a discussion of single HBA zoning in the section of this paper entitled, “Approaches to Zoning.”

LUN SECURITY IN THE FABRIC

The most basic function of a SAN is to connect hosts to storage. The unit of storage for most storage arrays is the Logical Unit Number (LUN), a partition of a single drive or a RAID-protected set of drives. Multiple LUNs are mapped to Fibre Channel (FC) ports on the array. Currently, Zoning is almost always used to group storage ports and hosts. Zoning used by itself allows a host access to all of the LUNs presented to a particular storage port on an array. LUN masking, which can be performed via the HBA or more commonly via the storage array, ensures that the authorized host has access only to a defined set of LUNs presented by each storage port. LUN masking could be used exclusively to map LUNs to hosts, but this would expose every host to every RSCN in the fabric, which compromises stability. Zoning and LUN masking combine to provide two layers of security, to prevent inappropriate access to a LUN even when one layer is misconfigured.

BEST PRACTICE: Always implement Zoning, even if you use LUN masking.

IDENTIFYING HOSTS AND STORAGE ZONE MEMBERS

When configuring Zoning, nodes must be identified and grouped. How the nodes are identified can affect the following:

- How well the Zoning configuration will scale
- How some advanced features are configured
- What processes are required if an HBA needs to be replaced

Types of Zoning

There are two types of Zoning identification: port World Wide Name (pWWN) and Domain,Port (D,P). You can assign aliases to both pWWN and D,P identifiers for easier management. The pWWN, the D,P, or a combination of both can be used in a zone configuration or even in a single zone. pWWN identification uses a globally unique identifier built into storage and host interfaces. Interfaces also have node World Wide Names (nWWNs). As their names imply, pWWN refers to the port on the device, while nWWN refers to the overall device. For example, a dual-port HBA has one nWWN and two pWWNs. Always use pWWN identification instead of nWWN, since a pWWN precisely identifies the host or storage that needs to be zoned.

D,P identification uses the switch domain ID and switch port to identify zone members. Until Fabric OS® 5.2.0, the "P" in D,P was equal to the port area ID and ranged in value from 0 to 255. Since the Brocade 48000 48-port blade requires the value of P to range from 0 to 383, P is now equal to the port index for a port. Since the use of port index instead of area ID is required when using D,P identification on a 48-port blade, all switches in the fabric must be running Fabric OS 5.2.0 or greater. Note that D,P identification is not globally unique (as shown in Figure 1), since every fabric can have duplicate domain IDs and every switch has a port index of 1, 2, 3, and so on. Any device cabled to a zoned D,P can communicate as defined by the zone configuration. The only identification method that is unique for all fabrics is pWWN.

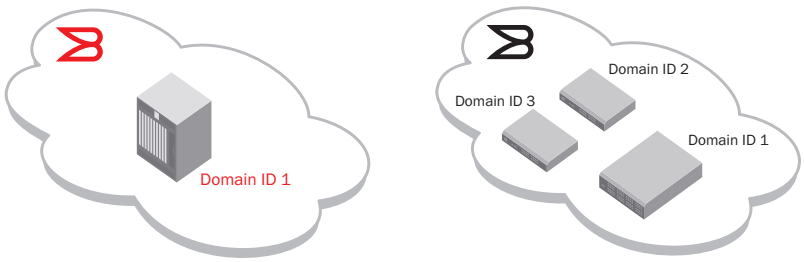

The red fabric has a Brocade 48000 with 32-port blades. The D,P identifier for port 64 is 1,63.

The black fabric has 2 x Brocade 200Es and a 4900. The D,P identifier for port 64 on the 4900 is 1,63.

RSCNs and Zoning

RSCNs do not cause broadcast storms, because they use a direct, unicast address instead of the broadcast address used in the Ethernet world. HBAs discard RSCNs unless the HBA is the specific destination for the RSCN message. In a worst case scenario, a single device could be flooded with RSCNs, but this does not extend to every device in the fabric. In versions of Fabric OS® (FOS) prior to 4.4.0, RSCNs were sent during a fabric reconfiguration, and HBAs would sometimes erroneously abort and then retry I/Os after a 30-second pause from a SCSI timeout. This behavior changed in FOS 4.4.0 (released in September 2004), and Brocade no longer sends an RSCN during a fabric reconfiguration. For more details on RSCNs and other SAN availability characteristics, contact your Brocade representative for papers entitled, "Availability in SANs and LANs" and "Brocade RSCN Events Matrix."

Figure 1.
D,P identification is not globally unique.

If an HBA or storage port fails, the Zoning configuration may have to be adjusted to accommodate a new pWWN when the replacement device does not allow its pWWN to be programmed. With D,P identification, a failed HBA or storage port just needs to be replaced—no Zoning changes are required—since the D,P identification allows any pWWN connected to switch port D,P to communicate with other members of its zones. With pWWN identification, a device can be cabled to a different switch port without impacting the Zoning, since the pWWN follows the device and not the switch port.

pWWN identification can be more secure than D,P identification, because any device physically cabled to a port could inappropriately grant storage access to an unauthorized host. Environments that allow multiple users access to switch ports and that have limited or no physical access monitoring should not use D,P identification because of the risk of unauthorized hosts being cabled to the wrong switch port.

Security

While pWWN identification prevents unauthorized Zoning access when cables are moved, it is possible to spoof a pWWN and thereby gain access to a zone using pWWN access. For this to happen, a device must be attached to the fabric and the spoofing attacker must choose a pWWN that has a zone to the node that is being attacked. Although this process would be very difficult, it would be possible, since fabrics do not reject duplicate logins, in accordance with Fibre Channel standards. This means that the spoofed pWWN would be allowed to join the fabric even if the legitimate pWWN were already logged in to the fabric.

The possibility of spoofing pWWNs can easily be reduced if you define Device Connection Control (DCC) policies (also known as "Port ACLs"). DCC policies map a pWWN to a physical switch port; then, if a pWWN tries to log in to a port not allowed by the DCC policy, the pWWN login is rejected. The DCC policy enforcement occurs before Zoning, since DCC controls device login and not which devices can communicate. To completely eliminate the risk of WWN spoofing, Diffie-Hellman Challenge-Handshake Authentication Protocol (DH-CHAP) authentication can be implemented. DH-CHAP enforces strong authentication between the host and switch when the host logs in to the fabric. The best approach for securing devices communicating in a fabric requires a combination of Zoning using pWWN identification, DCC policies, and DH-CHAP device authentication.

For ease of operation, use a single identification method. While pWWN and D,P identification can be mixed in a Zoning configuration or a zone, this can lead to configuration errors, because different processes are required for each method. Also, with mixed identification methods, a failed HBA may require Zoning and DCC modifications.

BEST PRACTICE: Use pWWN identification for Zoning for both security and operational consistency.

Advanced Features

Some advanced Brocade features, such as Fibre Channel Routing (FCR) and Fibre Channel Write Acceleration, use pWWN identification. For example, FCR allows zones to be created between independent fabrics using the Brocade 7500 SAN Router and/or Brocade FR4-18i Routing Blade. The routers analyze the zones in each fabric and allow nodes to communicate between the fabrics only if zones exist in both fabrics containing the relevant pWWNs. Since Brocade FCR technology uses Zoning, pWWNs must be used to uniquely identify the members of the routed zones. D,P identification cannot be used since non-unique pairs of nodes cannot be zoned, for example, domain ID 1, port ID 1 in fabric 1 cannot be zoned with domain ID 1, port ID 1, in fabric 2. Even if a fabric does not use FCR today, you should still use pWWN identification in the event that FCR is ever introduced to the fabric. In this way, a single, consistent identification method can be used for local, non-routed zones and zones that are routed between fabrics.

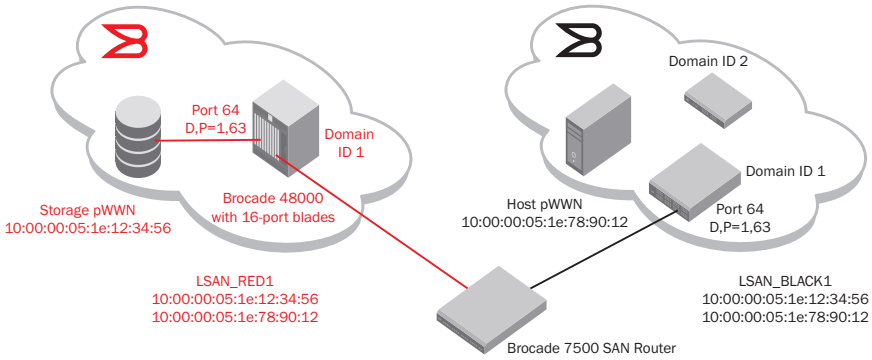

Figure 2.
Fibre Channel Routing with pWWN versus D,P.

As shown in Figure 2, the only identification method that is unique for all fabrics is pWWN. Storage or host pWWN connected to each port 64 in each fabric is unique.

Enabling LSAN_BLACK1 in the black (right) fabric and LSAN_RED1 in the red (left) fabric causes the Brocade 7500 SAN Router to allow those two pWWNs to communicate, while keeping the red and black fabrics separate and unmerged.

NOTE: The unused ports on the Brocade 7500 SAN Router are not part of either fabric.

Virtualization

Using pWWN identification is also important if you are implementing the Brocade Access Gateway feature, which uses the N_Port ID Virtualization (NPIV) protocol to remap server HBA N_Ports to external switch ports that act as N_Ports. Zoning the server HBA N_Ports must be done by pWWN, since a switch in Access Gateway mode, unlike a regular fabric switch, does not have a domain ID for use in D,P identification. Additionally, server Operating System (OS) virtualization software uses the NPIV protocol to allow virtual machines to have virtual HBAs. Zoning to these virtual HBAs must be done using the virtual pWWN for the virtual HBA.

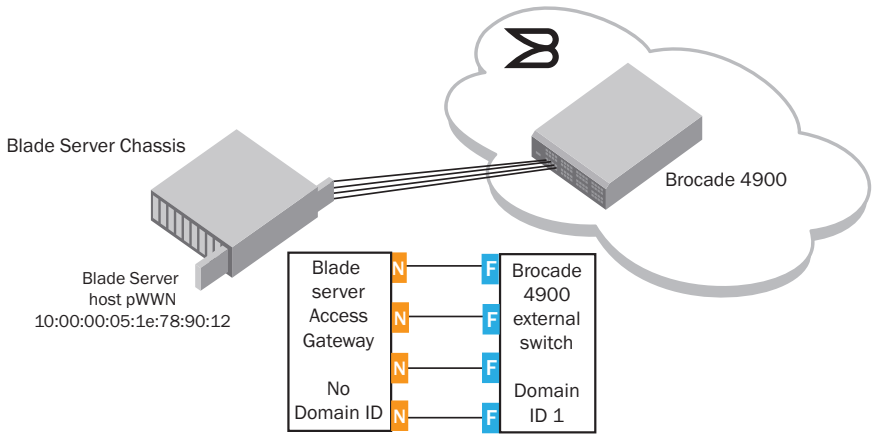

Figure 3.
pWWN identification for Access Gateway and server OS virtualization using NPIV.

There is no Domain ID associated with a blade switch in Access Gateway mode, so D,P identification cannot be used. pWWN identification must be used when zoning hosts switched by a device in Access Gateway mode.

Regardless of the Zoning identification, unused ports should be persistently disabled. Also, ports that are connected to storage and host devices should have their E_PORT functionality persistently disabled. See the Brocade white paper, "The Growing Need For Security in Storage Area Networks," for more information (available on www.brocade.com).

BEST PRACTICE: Use pWWN identification for all Zoning configuration unless special circumstances require D,P identification (for example, FICON).

NAMING CONVENTIONS

Naming conventions are very important to simplify Zoning configuration management for SAN administrators. Aliases for pWWNs, zone names, and zone configuration names can be very detailed, but can be designed in a user-friendly fashion. User-friendly alias names ensure that zone members can be understood at a glance and configuration errors minimized. Keep in mind also that the size of the Zoning database is not unlimited. In very large SANs, keep alias and zone names short to maximize the number of Zoning database entries. In Fabric OS 5.2.0 or later, a maximum Zoning database size of 1 MB is supported (four times larger than the previous Fabric OS maximum). You can check the size of the database by running the cfgSize command.

NOTE: If an older Fabric OS switch is present in a fabric, the maximum database size is reduced to the smallest common denominator for all switches in the fabric.

BEST PRACTICE: Make Zoning aliases and names only as long as they need to be—which enables maximum scaling (for very large fabrics, such as a FOS 5.2.0 and later fabric with 5000+ ports).

HOW DO SAN SWITCHES ENFORCE ZONING?

Now that node identification has been reviewed and zones have been created and activated—how does the switch actually enforce the defined zones? Brocade implements two methods:

- Frame-based hardware enforcement
- Session-based hardware enforcement

Advanced Features

Although software enforcement has been used by Brocade in the past, Brocade Fabric OS supports only hardware-based enforcement of Zoning. These scenarios are covered in detail below.

Software enforcement occurs when the Name Server service in the fabric masks the Name Server entries that a host should not access. When the host logs in to the fabric, it discovers only the unmasked Name Server entries. This "security through obscurity" technique depends on hiding the existence of certain storage targets from certain hosts. There is no mechanism with software-enforced Zoning that prevents a host from accessing storage.

For example, if a host was zoned to storage ports 1, 7, and 10 and the SAN administrator removed storage port 7 from the host's zone, the host would still know that storage port 7 was on the fabric and it would know the route to the port. Since software enforcement depends on hosts logging in to the name server to learn about storage targets, a host that does not log back in after a Zoning change can still access storage targets known from the initial login. Some HBAs will continue to access a storage port with software-enforced Zoning even though access was removed. (This should be much less common as HBA drivers have been improved over the years.) Note again that *Brocade Fabric OS supports only hardware enforcement of zones*.

Hardware enforcement is performed by the Application-Specific Integrated Circuits (ASICs) in fabric switches. Unlike software enforcement, hardware enforcement is a proactive security mechanism. Every port has a filter that allows only the traffic defined by the Zoning configuration to pass through. If traffic disallowed by the Zoning configuration is initiated, the ASIC will discard the traffic. Hardware enforcement is much more secure than software enforcement, because it does not depend on the "good citizen" behavior of the server HBA. Even if the host has knowledge of a storage port, hardware enforcement can proactively block any attempt to access that port.

The type of Zoning enforcement is determined on a zone-by-zone basis and depends on how the Zoning identification is configured. A zone can contain all pWWNs, all D,Ps, or a combination of pWWNs and D,Ps. An alias can be created for a pWWN or D,P and will act just like the aliased item for enforcement purposes.

In addition, an overlapping zone occurs when, for example, a pWWN is in Zone1 and, the D,P that this pWWN is cabled to, is in Zone2. Both zones are considered to overlap, as shown in Figure 4.

- pWWN1 is cabled to D,P1
- Zone1 = (pWWN1; pWWN2; and so on)
- Zone2 = (D,P1; D,P2; and so on)
- Both Zone1 and Zone2 overlap

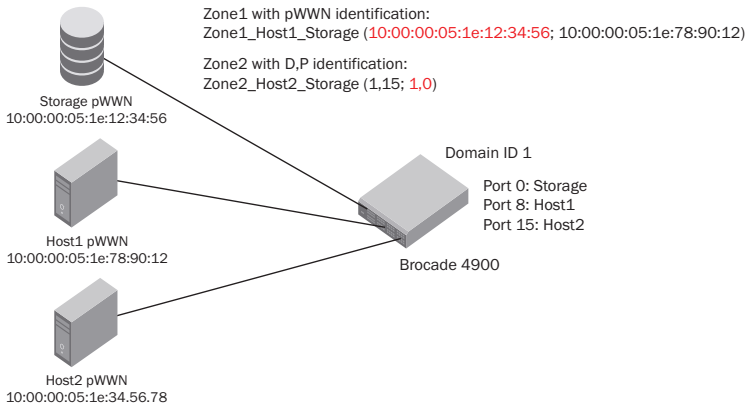

Figure 4.
If a Zoning configuration containing both Zone1 and Zone2 is enabled, the zones will overlap since the pWWN and D,Ps of the storage are used in separate zones.

Brocade still performs hardware enforcement on overlapping zones but enforcement is session-based and not frame-based.

Frame-based hardware enforcement occurs in the following cases on a per-zone basis:

- All zone members in a configuration use pWWN identification
- All zone members in a configuration use D,P identification
- Zones use either pWWN or D,P per zone with no mixed or overlapping zones

Session-based hardware enforcement occurs in the following cases on a per-zone basis:

- A mixed zone with both pWWN and D,P members
- Any overlapping zones

To ensure that all Zoning implements frame-based hardware enforcement, use pWWN or D,P identification exclusively. pWWN is more secure than D,P because of physical security issues and it enables the use of FCR, FC FastWrite, Access Gateway, and other features.

BEST PRACTICE: All zones should use frame-based hardware enforcement; the best way to do this is to use pWWN identification exclusively for all Zoning configurations.

AVOIDING ZONING TERMINOLOGY CONFUSION

Over the years, many terms have evolved to describe Zoning, some of which confuse the true nature of the Zoning method in question and lead to incorrect beliefs about how Zoning actually works. The most significant error made in describing Zoning is associating the identification method with how Zoning is enforced. With Brocade 1 Gbit/sec Fabric OS switches, only D,P identification was hardware enforced. This led to the term "hard zoning" being associated with D,P identification; D,P is also called "port zoning." pWWN identification on Brocade 1 Gbit/sec Fabric OS switches was software enforced, so the term "soft zoning" came into the SAN lexicon in addition to "WWN zoning."

Once Brocade released 2 Gbit/sec FOS switches, pWWN identification became hardware enforced, so the older association of pWWN identification with software enforcement and D,P with hardware enforcement became technically obsolete. But the terms "hard" and "soft" persisted with the incorrect belief that using the D,P identification was more secure than using the pWWN identification. Today, Zoning should be viewed as a security mechanism for SANs with two identification options and three enforcement methods. The relationship to identification and enforcement was outlined earlier, but f*or recent switches all Zoning is hardware enforced*. This is true for switches and directors running Brocade Fabric OS and Brocade M-EOS operating systems. The old terminology of "hard" and "soft" should be abandoned in favor of specifying enforcement and identification. For example, use statements such as "... the Zoning for this fabric is all frame based and hardware enforced using pWWN identification." or "... the Zoning is frame based and hardware enforced using pWWN except for some session-based, hardware-enforced overlapping zones."

BEST PRACTICE: Describe Zoning by enforcement method and identification type, that is, pWWN or D,P with hardware (frame- or session-based) or software enforcement. Abandon terms such as "hard" and "soft" or "port" and "WWN" Zoning. Brocade performs only frame- or session-based hardware enforcement.

APPROACHES TO ZONING

There are many ways to group SAN host and storage nodes for a particular Zoning configuration. Zone membership is primarily based on the need for a host to access a storage port. Hosts rarely need to interact directly with each other and storage ports never initiate SAN traffic by virtue of their nature as targets. Zones can be grouped by array, by host operating system, by application, or by location within the data center. In most cases, none of these methods is recommended for the reasons outlined earlier.

The recommended grouping method for Zoning is Single Initiator Zoning (SIZ), sometimes called "Single HBA Zoning." With SIZ, each zone has only a single HBA and one or more storage ports. If the HBA has both disk and tape storage devices, then you need to create two zones: one zone with the HBA and the disk devices and a second zone with the HBA and the tape devices. SIZ is optimal because it prevents any host-to-host interaction and limits RSCNs to just the zones that need the information within the RSCN.

Separating the disk and tape devices into separate zones prevents disk RSCNs from impacting tape devices, which tend to be more sensitive to RSCNs. While this level of Zoning might seem to be more labor-intensive compared to less granular grouping methods, it lays the best foundation for SANs, and will ultimately increase uptime and reduce the time required to troubleshoot problems.BEST PRACTICE: Use single HBA Zoning with separate zones for tape and disk traffic when an HBA is carrying both types of traffic.

BEST PRACTICE: Use Single Initiator Zoning with separate zones for tape and disk traffic when an HBA is carrying both types of traffic.

When no Zoning configuration is enabled, the default zoning access level can be either open, with all nodes seeing all other nodes, or closed, with all nodes isolated. This state is called the default zone and it describes the state of a fabric if no Zoning configuration is enabled.

For Brocade M-EOS fabrics, the default zone is set to disallow any communication in the fabric if no Zoning configuration has been performed. This very secure default setting limits unintentional node communication on an unconfigured SAN or during the time between disabling and enabling Zoning configurations.

For Brocade Fabric OS fabrics, the default zone can be optionally set to behave in the same way as the default zone in M-EOS fabrics. By default, Fabric OS fabrics have no default zone, which means that all nodes can communicate if no Zoning configuration has been performed or no Zoning configuration is enabled. To prevent this and make the Fabric OS SAN behave like an M-EOS SAN with default no access, enable the default zone with the -- noaccess setting. Additionally, there is a small period of time that occurs when disabling one zone configuration and enabling another during which all nodes can communicate by default in a Fabric OS fabric. This small period of time can, however, be disruptive in large fabrics, so it is important to implement a default zone with a --noaccess setting in large Fabric OS environments.

BEST PRACTICE: Implement default zone --noaccess in Fabric OS fabrics.

SUMMARY

Zoning is the most common management activity in a SAN. To create a solid foundation for a new SAN, adopt a set of best practices to ensure that the SAN is secure, stable, and easy to manage.

The following recommendations comprise the Zoning best practices that SAN administrators should consider when implementing Zoning.

- Always implement Zoning, even if LUN Masking is being used.
- Always persistently disable all unused ports to increase security and avoid potential problems.
- Use pWWN identification for all Zoning configuration unless special circumstances require D,P identification (for example, FICON).
- Make Zoning aliases and names only as long as required to allow maximum scaling (in very large fabrics of 5000+ ports for Fabric OS 5.2.0+).
- All Zones should use frame-based hardware enforcement.
- Use Single Initiator Zoning with separate zones for tape and disk traffic if an HBA is carrying both types of traffic.
- Implement default zone --noaccess for FOS fabrics.
- Abandon inaccurate Zoning terminology and describe Zoning by enforcement method and identification type.
- Use the free Brocade SAN Health™ software and the Fabric OS command zone -validate to validate the Zoning configurations.

For more information, 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