

Cisco Smart Zoning for Cisco NX-OS SAN Fabrics



Overview


Zoning has been the primary method of access control among SAN devices since the early days of Fibre Channel, and there have been few changes in the way zoning has been implemented since then. The Cisco® Smart Zoning feature is a major step forward in simplifying zoning to reduce the burden on administrators and enable automation of SAN provisioning without affecting SAN hardware resource consumption.

The Problem with Traditional Zoning

SAN administrators allow servers (initiators) and storage devices (targets) in a Fibre Channel SAN to talk to each other by adding them to the same zone. In the fabric, permissions defined in this way are converted to access control entries, (ACEs) which are programmed into ternary content-addressable memory (TCAM) hardware in the switches.

Traditionally, zones have members, and all members of a zone can talk to each other. Each pair of members consumes two ACEs in the TCAM: one ACE permits the first member to receive traffic from the second member, and the other ACE permits the second member to receive traffic from the first member. Mathematically, the number of ACEs consumed by a zone with n members would be $n \times (n - 1)$. Since hardware resources are finite, a moderate number of large zones can exceed the TCAM capacity of a switch (Figure 1).

Figure 1. TCAM Consumption Grows Geometrically with Zone Size



The solution to this problem has been to use 1-1 zoning, in which each zone consists of a single initiator and a single target. This solution solves the problem of excessive TCAM consumption, but it imposes a burden on the SAN administrator by requiring the creation and management of a large number of zones. More zones generate more work, and more possibilities for errors. In very large fabrics, this solution may even run up against system software limits on the size of the total zone database.

Solution

Cisco Smart Zoning takes advantage of the fact that storage traffic is not symmetrical or egalitarian like LAN traffic, in which any Ethernet or TCP/IP host may need to talk to any other host. Storage is asymmetrical: zone members are either initiators or targets, and in most cases, initiators do not talk to other initiators, and targets do not talk to other targets. There are exceptions to this generalization, such as array-to-array replication, and any solution must take those into account.

Consider an example in which an application has eight servers, each with dual host bus adapters (HBAs) or converged network adapters (CNAs), talking to eight storage ports. These devices are split among two redundant, disjoint SAN fabrics, so each fabric has eight HBAs and four storage ports for this application. Figure 2 shows the permissions that will be provisioned by traditional zoning. A total of 132 ACEs are created because each of the 12 members of the zone is provisioned with access to all 11 other members.

Figure 2. Permissions with Traditional Zoning






Figure 3. Permissions with Smart Zoning




Figure 3 shows how dramatically Smart Zoning can reduce ACE consumption. With Smart Zoning, there are only 32 initiator and target pairs, and with each pair consuming an ACE at each end, TCAM used only 64 ACEs in the TCAM, reducing consumption by more than 50 percent. For larger Smart Zoning deployments, the savings is even greater.

The Most Important Savings: Administrative Time

Zoning Becomes Intuitive

With Smart Zoning, zones can now be defined as one-to-many, many-to-one, or many-to-many without incurring a penalty in switch resource consumption. Thus, administrators can now define zones to correspond to entities that actually are meaningful in their data center operations. For example, they can define a zone for an application, or for an application cluster, or for a hypervisor cluster without compromising internal resource utilization. Consider a maximum VMware vSphere cluster of 32 servers that uses a total of 8 storage ports. A single zone for this cluster has 40 members. With traditional zoning, it would consume $n \times (n - 1)$ ACL entries, which translates to $40 \times 39 = 1560$ ACL entries. With Smart Zoning, ACL consumption drops to $32 \times 8 \times 2 = 512$.

Zoning Becomes Faster

Smart Zoning also saves time. Consider what you have to do to create a zone. First, you create the zone itself; that takes one command. Then you have to add each member to the zone; that takes one command per member. So with one-to-one zoning, you need three commands per zone, if you're using the command-line interface (CLI). The process using the Cisco Data Center Network Manager for the SAN (DCNM-SAN) GUI is comparable. With Smart Zoning, the number of commands is reduced significantly, as shown in Table 1.

Table 1. Smart Zoning Saves Administrative Time

Operation	Today: One-to-One Zoning			Today: Many-to-Many			Smart Zoning		
	Zones	Commands	ACL Entries	Zones	Commands	ACL Entries	Zones	Commands	ACL Entries
Create zones	32	96	64	1	13	132	1	13	64
Add an initiator	+4	+12	+8		+1	+24		+1	+8
Add a target	+8	+24	+16		+1	+24		+1	+16

Smart Zoning Scales Better

By reducing the number of zones required, Smart Zoning extends the amount of time before extremely large fabrics need to worry about encountering software limits on the number of zones in a fabric. This extension is especially valuable as new technologies such as CNAs on the motherboard reduce the cost of SAN-enabling servers to nearly zero, increasing SAN penetration of new servers to nearly 100 percent. Fabrics with fewer zones are much easier to maintain, because administrators can find a zone that needs modification much more quickly. Because zone names can now correspond to applications, clusters, etc., the likelihood of administrative mistakes that could affect application availability is reduced.

Smart Zoning Paves the Way for Future Automation

When each application or cluster requires the creation of several, or even dozens, of zones, automating this process is very difficult. However, with Smart Zoning, you can consider zone membership as just another attribute of the higher-level system. For example, all server and storage ports involved in a payroll application cluster would be members of the Payroll zone, and adding a new server port or storage port to the payroll cluster would be easy for a higher-level management application to do.



Implementing Smart Zoning

Smart Zoning requires all Cisco MDS 9000 Family switches in the fabric to use Cisco MDS 9000 NX-OS Software Release 5.2(6) or later and have the Smart Zoning feature enabled. Smart Zoning cannot be enabled if the fabric contains third-party switches or switches running earlier versions of Cisco MDS 9000 NX-OS or SAN-OS. After Smart Zoning is enabled, you cannot add a switch to the fabric that is incapable of supporting Smart Zoning. Smart Zoning is available in both basic and enhanced zoning modes.

Smart Zoning is enabled on a per-VSAN basis. After it has been enabled for a VSAN, it can be disabled for individual zones, so if you need to have storage devices talk to each other for replication or backup purposes, for instance, you can simply put those devices in a zone or a VSAN that does not have Smart Zoning enabled.

A utility is provided to assist the administrator in evolving the existing zone set to take advantage of Smart Zoning. The utility interrogates the name server and classifies each member as an initiator (I) or target (T). To zone devices that have not yet been connected to the fabric, the administrator simply specifies I or T in the Cisco MDS 9000 NX-OS CLI or Cisco DCNM-SAN GUI.

Ordering Information

Smart Zoning requires no special or additional software license. It is included in the base operating system in Cisco MDS 9000 NX-OS Software Release 5.2(6) and later, and in Cisco NX-OS Software Release 6.1 and later for Cisco Nexus® 7000 Series Switches.

For More Information

For more information about Cisco Smart Zoning and the Cisco MDS 9000 Family, visit <http://www.cisco.com/go/storage>.