

Zoning: The Key to Flexible SAN Management

***Chris King, System Engineer
QLogic Corporation***

Introduction

Adding Fibre Channel switches to Storage Area Networks (SANs) previously cobbled together with legacy private devices can give storage managers tremendous flexibility to easily create and manage SANs that include both private and public devices. The key to these highly flexible, manageable SANs is the switch's capability to create zones, or partitions within the total SAN fabric. Prior to the FC-SW-2 specification, Fibre Channel switch vendors developed their own naming conventions and features. Now that the FC-SW-2 specification is completed and approved, any switch that is fully FC-SW-2 compliant, regardless of vendor will coexist in a SAN fabric. As storage managers evaluate switches from various vendors, it is important to understand how zoning is accomplished within the fabric. All methods of zoning are not equal: this paper will discuss QLogic's SANbox and SANbox2 zoning capabilities and naming conventions in order to help storage managers make informed decisions regarding their switch purchases.

SANs: Building on Yesterday, Primed for Tomorrow

Early SANs were built with hubs based on private, arbitrated loops. Although these early SANs had a theoretical maximum of 126 devices, the realistic maximum was much lower, and performance varied. Troubleshooting was difficult at best, with many devices sharing a loop. With the introduction of the fully public switch, SAN fabrics can theoretically contain over 7.7 million nodes, and exponential SAN growth is a reality. With newer switch features, particularly zoning, creating and managing these large SANs is now easier than ever before. Best of all, full fabric based switches allow storage managers to maximize the effectiveness of currently installed private loop devices.

Zoning the original SANbox switch

Switch zoning allows the SAN manager to partition the SAN into various groupings to allocate appropriate storage where it makes the most sense. QLogic's original SANbox (1Gb switch) supports four zone types: hard zones; name server zones; broadcast zones, and segmented loop zones. These zone types give the SAN manager the flexibility to partition the SAN into logical groupings of devices that can share information whether these devices use private or public fabric addressing schemes, maximizing the investment in the installed private loop devices.

Defining zones, or adding or changing devices within a zone of the SANbox switch is easily performed via QLogic's SANsurfer™ management software. Using the SANsurfer GUI, the SAN manager can dynamically reconfigure the current fabric zone configuration to add or reallocate devices to existing or new zones to meet the growth needs of the company.

QLogic zoning also can be used to simplify a heterogeneous environment within the same switch fabric. This gives SAN managers the freedom to add any type of device to the fabric, while keeping these devices separated by zones to prevent conflicts between fabric devices.

Hard Zoning

Hard zoning, so named because it is programmed into the hardware, is the most secure of any zone type, preventing communication from any device that is not in the same hard zone. QLogic's hard zones cannot overlap, and require at least one dedicated Inter-Switch Link (ISL) for each hard zone that includes more than one switch. The dedicated ISL guarantees the I/O bandwidth in the hard zone, giving the SAN manager the flexibility to balance bandwidth across all hard zones for maximum overall fabric performance. Designing hard zones for maximum performance eliminates the need to reconfigure the SAN zone to adjust the workload and minimizes I/O bottlenecks.

The QLogic SANbox switch supports a maximum of sixteen hard zones per SAN fabric. Each of these sixteen isolated hard zones can be further sub-divided into other QLogic zones via the name server, segmented loop or broadcast zone features, allowing the SAN manager to create specific sub zones inside the hard zones.

The combination of hard zones with other zone types enables the larger fabric to be carved into separate fabrics for specific uses, maximizing switch port efficiency and reducing the number of switches required.

Name Server Zones

Name server zones are extremely flexible, allowing the SAN manager to create up to 256 zones using either switch ports or world wide names (WWN) to assign zones. Name server zones can overlap, and all ISLs within a hard zone are available to all the name server zones, providing load balancing for maximum data throughput under heavy workloads. While there is no performance difference between a port-based name server zone and a WWN server zone, there are reasons to choose one type over another, as discussed below.

Note — some competitors mistakenly refer to QLogic's SANbox switch port zoning as "hard zoning," and its world wide name zoning as "soft zoning." QLogic's hard zoning is much more secure than any form of name server zoning. QLogic's SANbox name server zoning is considered neither "hard" nor "soft."

Defining zones based on switch ports

Since all devices are connected to a particular port on a specific switch, the easiest way to physically map out all the devices onto a SAN is to define a name server zone that correlates with switch ports. There are two instances when zones based on switch ports can create problems: First, switch port zones are relatively coarse—all devices on the port must be included in the zone; and second, if a device is moved from one port to another it may end up moving to a different zone. This problem is commonly found in cable or GBIC replacements.

Defining zones based on world wide names

A more flexible solution is to define zones based on world wide names (WWN). With the world wide name server zone, a device is assigned to a zone according to its unique name, giving the SAN manager total freedom to install the device, host or storage, anywhere within the SAN fabric. A device assigned by WWN will stay in its assigned zone regardless of the physical port that serves as its connection. This type of zoning also eases troubleshooting by allowing the SAN manager to move the device at the questionable port location to another port location to verify if the problem is with the port/GBIC/cable or if it follows the device to the other port without reconfiguring the zone.

A secondary advantage in using WWN zoning within a public fabric is the ability to troubleshoot down to the device level on a loop. For example, “just a bunch of disks” (JBODs) are often attached to a single port. If this loop appears as one device on a SAN, it is hard to troubleshoot if a single device misbehaves. However, if the WWN registers problems at the device level, troubleshooting is much easier. Additionally, the WWN naming feature can take advantage of RAID controllers that have the ability to present multiple Fibre Channel devices on a single arbitrated loop.

Unfortunately, there are some legacy devices currently installed that do not report a WWN to the name server—obviously switch port zoning is the only name server option in these instances.

Broadcast Zoning

Broadcast zones are assigned to separate TCP/IP, or network traffic from that of SCSI, or storage traffic in a SAN environment. Broadcast zones can be set up to send broadcast messages to only IP devices that need to receive them. Broadcast zones reduce traffic on the fabric by eliminating unnecessary message processing by host and storage connections that don't process IP traffic. Broadcast zones can be flexibly applied by creating up to 16 overlapping zones. Broadcast zones can overlap name server zones; however, they cannot overlap hard zones.


Segmented Loop Zoning

Segmented Loop Zoning is a feature unique to QLogic that enables private devices to be zoned much as if they were part of the public fabric. To address private Fibre Channel Arbitrated Loop (FC-AL) devices, QLogic designed a port feature in the switch called SL_Port (Segmented Loop Port). SL_Ports extend access to private devices in the fabric by automatically recognizing all devices connected to that assigned port as private FC-AL devices.

The SAN manager can create up to 256 non-overlapping Segmented Loop Zones in a fabric, with the freedom to locate the ports in each zone on any switch in the fabric. These zones can extend across and up to six cascaded or mesh switches in a fabric, with all ISLs shared between switches.

Note — SL_Ports are unique to QLogic. While a competitor's QuickLoop feature has been compared, it is strictly limited to a pair of switches sharing private devices, a restrictive feature when designing large fabrics with private devices. For more information on integrating legacy devices, see QLogic's white paper "Reducing the Complexities of Managing Private Loops".

Diagram 1. Illustrates QLogic's SANbox switched fabric with four zone types



Designing the Truly Flexible, Manageable SAN Fabric

The real power of QLogic zoning lies in its ability to be used in a combination of zone types. Hard zones, the most secure of all zoning, allow the fabric to be partitioned into multiple, independent, virtual fabrics. If no hard zones are defined, the entire fabric is transparently treated as a single hard zone.

Within hard zones, other zones can be overlapped, allowing some ports to be dedicated to private legacy devices using SL Zones (or Segmented Loop Zones). All other ports can be zoned using the port and WWN name server zones to allow all devices to be connected for efficient SAN utilization and ease of management. Finally, broadcast zones can be used if IP traffic is present to limit the impact of IP broadcasts on SCSI devices. For maximum ease in growing large fabrics, all QLogic zone types can be configured on fabrics of any size, with no limitations as to which ports or devices in a fabric can be placed into a zone.

Zoning in Combination

Below are a few real-world examples of the benefits of creating sub-zones within a hard zone:

- Dedicate an ISL from one switch to another within a zone
- Dedicate known I/O bandwidth within a zone
- Segregate specific company departments
- Overlap specific port or WWN groups dynamically on the fly for data backup—then reconfigure the zones to original configuration
- Limit IP broadcasts to specific devices in multiple zones within the hard zone and overlap those broadcast ports with name server zone ports to communicate to other devices in the SAN
- Allow private loop devices to share the same switch with public devices. Both private and public devices will operate independently using the same switch, reducing the number of switches necessary to build the fabric.

Zoning with SANbox2

The zoning features and naming conventions of the SANbox2 (2Gb fibre channel switch) differ from that of the original SANbox switch because the SANbox2 was developed to conform to the recently ratified T11 ANSI standards, FC-GS-3 and FC-SW-2. This means the SANbox2 is completely interoperable with other vendors' switches that also conform to these new standards. Zoning is administered through the QLogic SANbox Manager application. Additionally, because the SANbox2 is the first 2-gigabit switch to market, the IT storage manager now has increased flexibility to expand an existing SAN by integrating higher speed industry-compliant switches into the SAN without being limited by another vendors slower non-conforming switches.

The New Hierarchy in Zoning

The latest ANSI standards define a new hierarchy within a zone set, described below.

Zone Sets

A zone set is the highest level of the zoning hierarchy (see Diagram 2 below). The zone set contains assigned zones; the zones contain assigned members. There can be several zone sets in a single fabric; however, only one zone set can be active at a time. Multiple zone sets are especially helpful for dynamic reconfiguration of the SAN to perform backups, for example, when moving a tape library from one server to another.

In order to change zone sets, the administrator simply deactivates the old zone set and activates the new. While this can be done without bringing down the SAN, it should not be performed when I/O activity is present in the SAN. Instead, active zone sets should be changed when the SAN is idle.

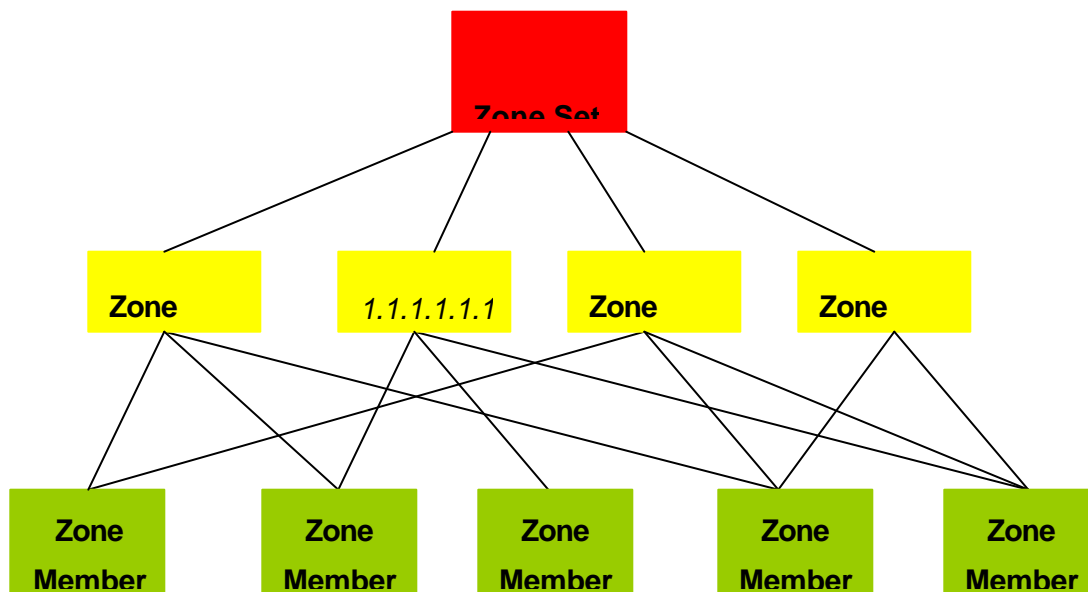
Zones

As with the original SANbox switch, zones are made up of a group of assigned devices, now called “members.” Additionally, all device members assigned in the zone can belong to one or multiple zones within the zone set. This capability helps facilitate the sharing of backup devices with multiple storage devices.

Members

Members are simply devices within a zone. With the SANbox2 and other compliant switches, member devices can be assigned to a zone via world wide name (WWN), port number, or Fibre Channel Address (FCA). Any member can be assigned to multiple zones within a SAN.

Diagram 2. The Zone Hierarchy



Hard Zoning

According to the FC-GS-3 standard, a hard zone is defined as “A zone that is enforced by the fabric, often as a hardware function. The fabric will forward frames among zone members within a hard zone. However, the fabric prohibits frames from being forwarded to members not within a hard zone.”

As with the original SANbox switch, hard zoning is the most secure zoning type. Because the hard zone device table is created and stored at the ASIC level of the switch, the fabric enforces the hard zones and will only forward frames among other hard zone members. The SANbox2 can support a maximum of sixty-three (63) hard zones within a fabric. Additionally, hard zones can overlap and contain both name server and broadcast zones. All hard zone members are defined by domain/port number.

Name Server Zones (Soft Zones)

According to the FC-GS-3 standard, a soft zone “consists of zone members who are made visible to each other through client service requests. Typically, soft zones contain zone members that are visible to devices via the name server exposure of zone members. The fabric does not enforce a soft zone.”

As with the original SANbox switch, name server zones are extremely flexible. Name server zones can overlap broadcast zones, and all ISLs within a hard zone are available to the name server zones. With the SANbox2, there is **no** maximum to the number of name server zones that can be created. The newest ANSI standards identify three member-address schemes used in name server zoning: port, world wide name (WWN); and Fibre Channel Address (FCA). Each address member can coexist in the same zone. The SANbox Manager application’s zoning window displays all three member types per device, giving the IT storage manager the flexibility to choose the member-address scheme that best fits their zoning needs.

©2002 QLogic Corporation. All rights reserved. The QLogic logo is a trademark of QLogic Corporation, which may be registered in some jurisdictions. All other brands and product names are trademarks or registered trademarks of their respective holders. Information supplied by QLogic Corporation is believed to be accurate and reliable. QLogic Corporation assumes no responsibility for any errors that appear in this brochure. QLogic Corporation reserves the right, without notice, to make changes in product design or specifications. Brocade, and QuickLoop are trademarks or registered trademarks of Brocade Communications Systems Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.