

How to Build Multi-Vendor Switched Fabrics

Ed McGlaughlin

Engineering Fellow

QLogic Corporation




Abstract

FC-SW-2 is an open standard for switch-to-switch communication, allowing IT professionals to choose best-in-class products with the assurance that these products can be deployed in multi-vendor storage area networks (SANs). Fibre Channel switches complying with this standard communicate connectivity and configuration information, path selection, and routing, as well as management and event services using the same language. FC-SW-2 also provides standardized mechanisms for SAN management. These applications can configure, manage, and monitor multi-vendor Fibre Channel SANs from any particular point in the fabric. In this session, IT professionals will learn how to configure and deploy multi-vendor switched fabrics. Detailed switch configuration data and step-by-step configuration procedures are provided to integrate the QLogic switches into existing Brocade, Cisco, IBM, INRANGE, and McDATA Fibre Channel switched fabrics that comply with the Fibre Channel switch fabric 2 (FC-SW-2) standard.

Objectives


- Show why multi-vendor fabrics are important
- Show that multi-vendor fabrics are possible
- Highlight capabilities and limitations
- Show how they can be managed
- Future directions in SAN fabrics

Where the industry has been



- Smaller scale SANs
- SAN islands
- Single vendor per device type in SAN islands

Where the industry is going



- Larger scale SANs
- Consolidation of SAN islands
- Multiple vendors per device type in larger SANs
- Role specialization of switch types

Why Multi-Vendor Switch Fabrics are Important

- Consolidation of SAN islands combine multiple vendors
- New switch suppliers are entering the switch market
- New switch types and roles are being established – from embedded blade server switches to core switches and virtualization platforms
- Users are demanding that they have the opportunity to select the best-in-class switch for their applications
- As a result vendors that vigorously support interoperability are gaining market share

Fibre Channel Overview

- Fibre Channel interoperability standards are defined by the T11 technical committee of INCITS (InterNational Committee for Information Technology Standards)
- Within T11 there are two working groups that focus on SAN fabric issues: SW (switch interoperability) and GS (device to fabric generic services)
- Current approved revision levels are SW-2 and GS-3
- Broad vendor participation in current development efforts:

SW-3

Facilitator – Ed McGlaughlin, QLogic
Editor – Steve Wilson, Brocade
Secretary – Scott Kipp, McDATA
Contributor – Claudio DeSanti, Cisco


GS-4

Facilitator – Mike O'Donnel, McDATA
Editor – Craig Carlson, QLogic
Secretary – Bob Nixon, Emulex

Switch Fabric Overview


- SW-2 defines all of the requirements to build reliable interoperable multi-vendor switch fabrics
 - Build flexible topologies
 - Device discovery and event notification
 - Zoning

Building Fabrics




- Build flexible topologies: cascade, mesh and multi-tier interconnects, etc.
- Routes are defined on a shortest path first.
- Path fail-over and load sharing are supported with multiple equal cost paths

Device Discovery




- Hosts can discover all the targets that they are configured for, regardless of where they are attached in the fabric via the GS-3 Name Server and the SW-2 distributed Name Server (dNS)

Event Notification



- Host are notified when targets are added or removed via RSCN (Registered State Change Notice)
- Following an RSCN, hosts perform a new Discovery

Zoning



- Standardized method of limiting the scope of targets that a host can access based on WWN identification.
- Changes made on one switch are propagated to all.
- Zoning is automatically disseminated when a new switch or replacement switch is added to the fabric

Interoperability Compliance Testing

■ SANmark Certification

- HBA and Target
 - SCD-2003 Fabric Attached Devices
- Fabric/Switch Behavior
 - SCD-3001 E_Port Conformance
 - SCD-3002 FL_Port Conformance
 - SCD-3010 Registered State Change Notification
 - SCD-3020 Zoning

- Reference: <http://www.sanmark.org/process/>



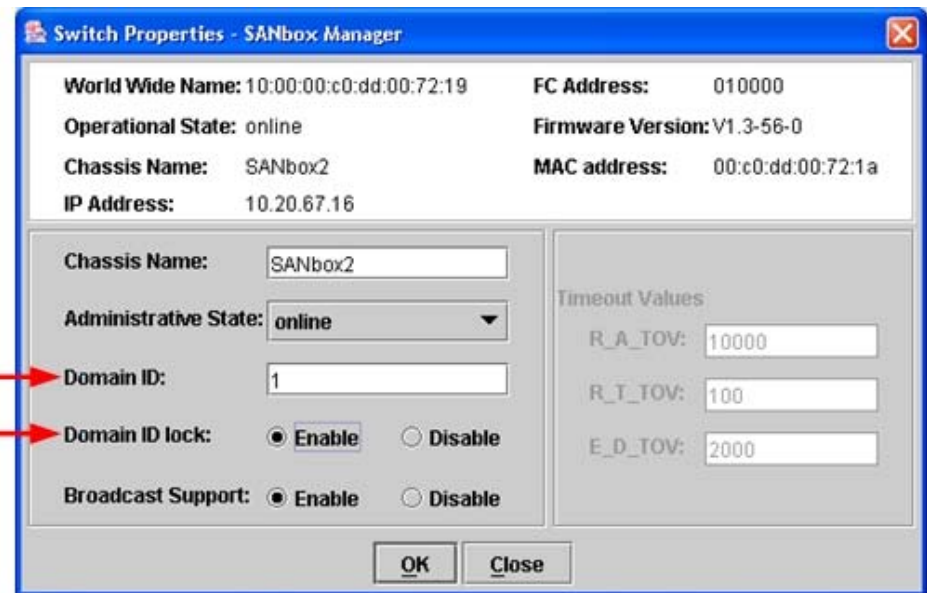
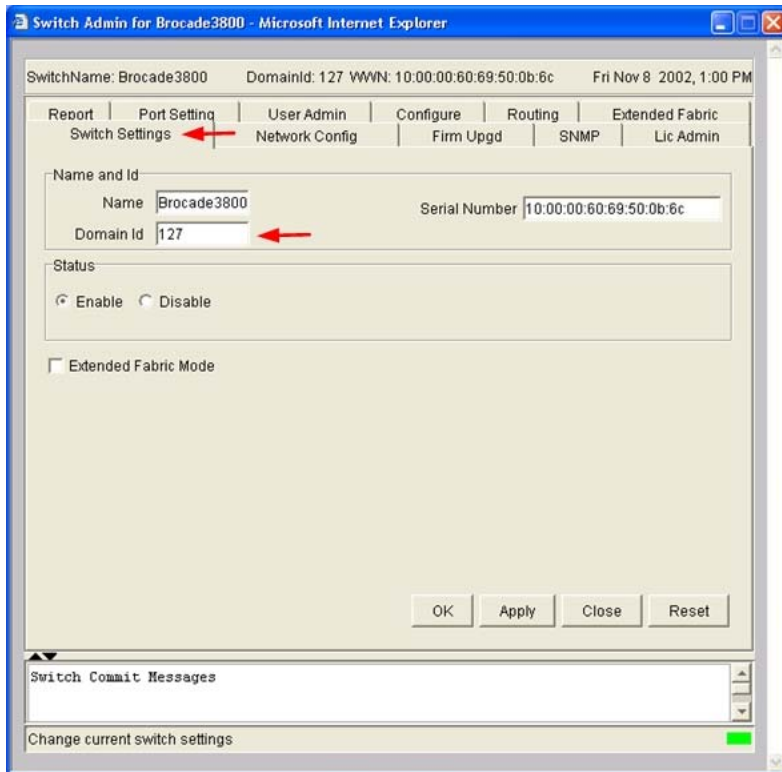
Switch Interoperability Configuration

- Domain ID Configuration
- Principal Switch Configuration
- Timeout Values
- Zone Configuration
- Operating Mode Configuration
- Vendor Specific Configuration
- Refer to this guide for more information:
 - http://www.qlogic.com/interopguide/thanks_pdf.asp

Domain ID Configuration

- A switch is identified by its Domain ID which is the upper 8-bits of the FC address
- T11/SW-2 allows for 239 Domain IDs with dynamic assignment made by the Principal Switch
- In practice, often a subset of Domain IDs are used and they are set statically by configuration
- Some switches add an offset to the internal representation, for example the Domain ID range of 1-31 on a McDATA switch corresponds to the range of 97-127 on QLogic and Brocade switches.

Domain ID Configuration




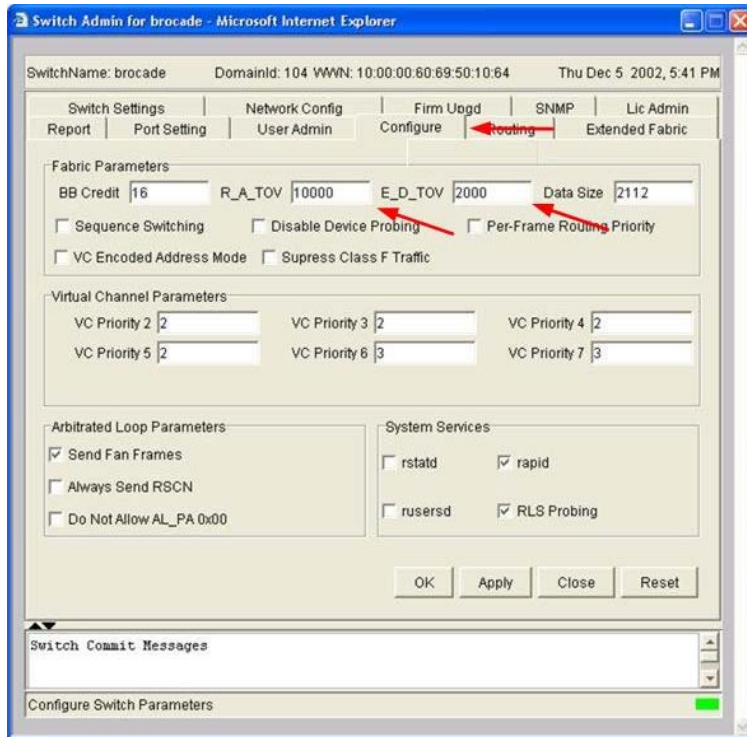
Principal Switch Selection

- The role of the Principal Switch is to provide Domain ID assignments to the other switches in the SAN
- T11/SW-2 define the process of selecting the Principal Switch and the method of assigning Domain IDs
- All switch vendors perform this negotiation for principal switch automatically and therefore no additional configuration is required

Timeout Values

- As per FC-SW-2 Fibre Channel standards, set all switches to the following timeout values (TOV) in order to successfully establish an E-port connection:
 - R_A_TOV = 10 seconds
 - E_D_TOV = 2 seconds
- These are the default values set by all vendors but should be verified as part of the setup

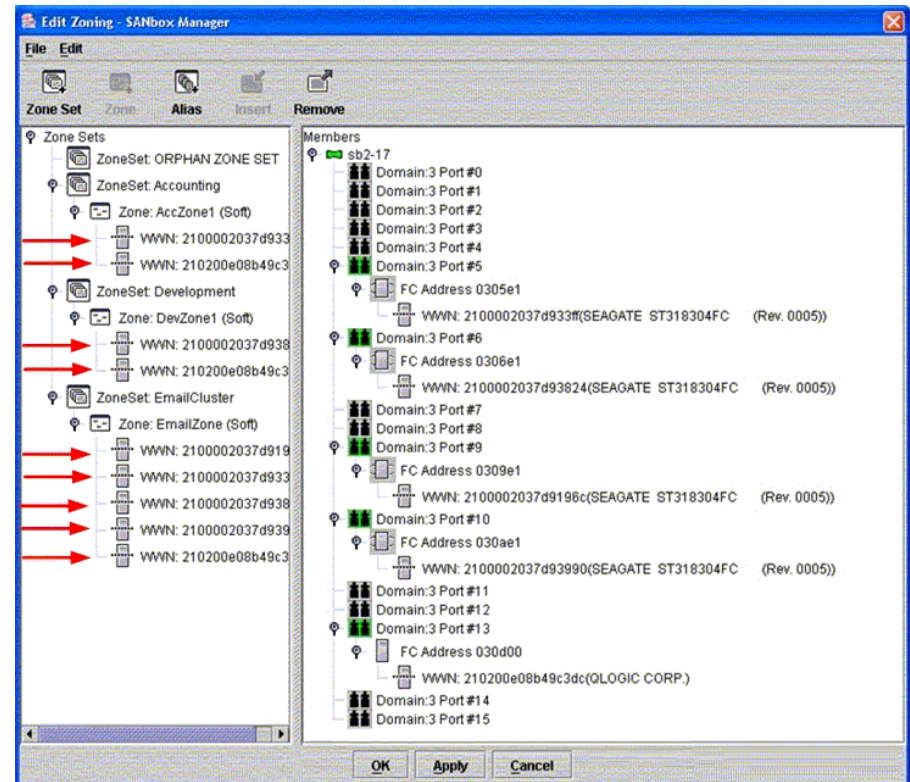
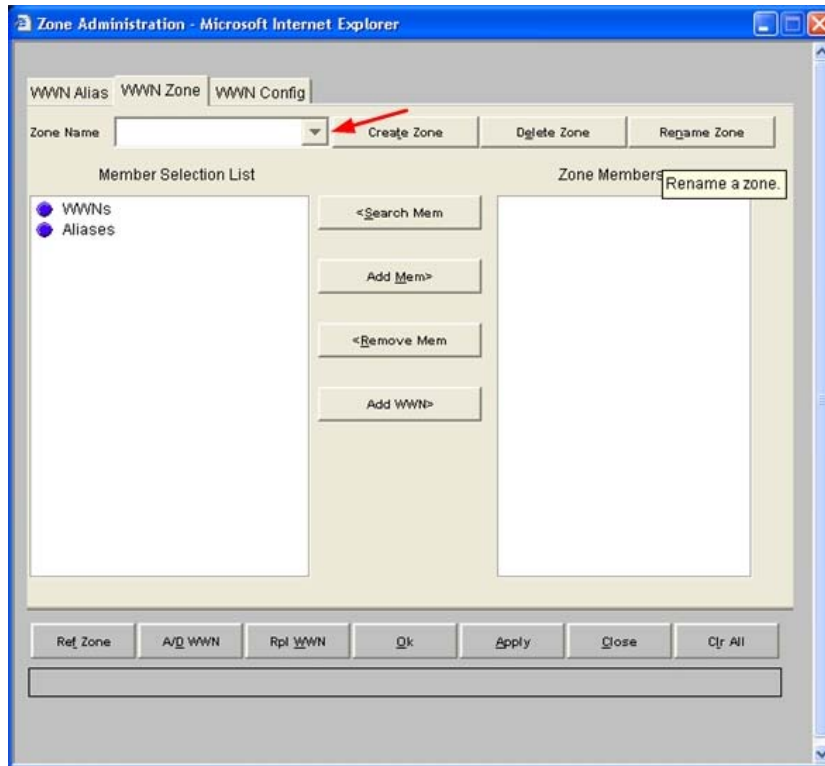
Timeout Value Configuration



Zone Configuration

- Zoning provides a method of limiting visibility of which hosts can access which storage device
- Interoperable zoning is provided via world wide port name (WWPN) zone members
- Verify there is a consistent setting for visibility for un-zoned devices
- Zone naming requirements
 - Must be 1–64 characters in length.
 - All characters are ASCII, template: [a-zA-Z][a-zA-Z0-9_]*

Zoning Configuration



Operating Mode Configuration

- Some switches have two modes of operation – SW-2 interoperability and legacy/native
- The information presented so far is based on SW-2 interoperability mode

Brocade Interoperability

- A significant portion of the Brocade install base operates in native, non-SW-2 mode
- QLogic and other vendor's switches offer a compatibility mode to connect to these switches
- This offers the following advantages:
 - Non-disruptive to existing fabrics
 - Supports Domain/Port zoning as well as WWPN
 - Allows Brocade switches to maintain specific features
- The installed base is starting to transition over to SW-2 compliant mode to support the advantages of open systems

Fabric Management

- Standardized Management is available and is used by industry leading 3rd party management products
- T11/GS-3 defines Management Server, Un-zoned Name Server, Zone Server and Switch Configuration Server for comprehensive in-band management
- The Fibre Alliance MIB provides out-of-band discovery and event reporting.
- SNIA Storage Management Initiative – future path for multi-vendor SAN management

What's Coming

- GS-4 and SW-3 are nearing completion. Most significant changes are in zoning
 - Hard Zoning support
 - Fabric Port Name zoning
 - Broadcast zoning
 - Zoning change session control
 - Zoning activation results diagnostic
- Fabric Security – FC-SP
 - Requires DH-CHAP authentication
 - Switch-to-switch and device level authentication

Summary

- Multi-vendor fabrics:
 - can be set up with a few easy steps
 - are now on-line at several high profile sites
 - are supported by a broad set of 3rd party application software
 - provide IT professionals with ability to select the switch with the best cost/benefit match to the requirements of the application



HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

