

Advanced SAN Troubleshooting

BRKSAN-3708



Agenda

- **Fibre Channel & MDS Switch Basics**

 - FC Operations Review

 - Addressing, FC Services

 - Domains, Zoning

- **MDS Serviceability Tools**

 - FC analyzer

 - SPAN & PAA (WireShark usage)

 - SAN/OS (Output analysis, debug, logs, Cores)

 - Performance Manager (Licensed part of Fabric Manager)

 - NTOP (Using Netflow and SPAN w/PAA)

- **Troubleshooting**

 - Device connections

 - ISL's

 - Zoning

 - IVR

 - NPV

Fibre Channel Operations-

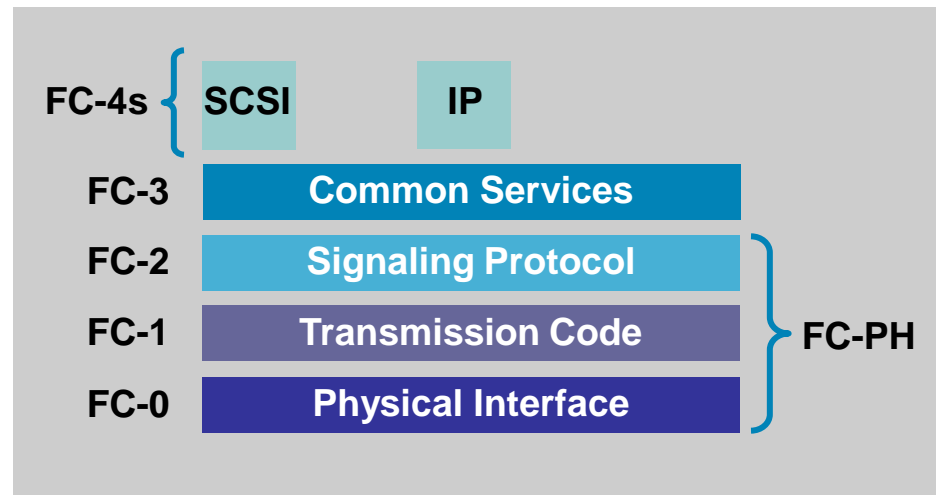
What Are We Looking for When We Troubleshoot?

- **Physical Layer basics**
- **Understanding FC addressing**
 - Needing to live within the limits of the standards
- **Fibre Channel Protocol Services & SAN/OS**
 - Refreshers on FLOGI. PLOGI, standards operations as they relate to SAN/OS
 - The ISL connection
- **Domains** - Operational understanding
- **Zoning** – Basic & Enhanced operation

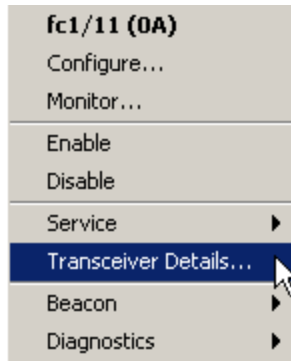
Fibre Channel Layers

Structure Is Divided into Five Levels of Functionality

- **FC-0**—defines the physical interface characteristics
Signaling rates, cables, connectors, distance capabilities, etc.
- **FC-1**—defines how characters are encoded/decoded for transmission
Transmission characters are given desirable characters
- **FC-2**—defines how information is transported
Frames, sequences, exchanges, login sessions
- **FC-3**—place holder for future functions
- **FC-4**—defines how different protocols are mapped to use Fiber Channel
SCSI, IP, Virtual Interface Architecture, others



Detailed SFP Levels



This output of detailed SFP transceiver only available on new 4, 8 and 10 gig qualified Cisco SFP's

10.91.42.189 - Show int fc1/11 Transceiver Details

```
fc1/11 sfp is present
  name is CISCO-FINISAR
  part number is FTLF8524P2BNL-C2
  revision is 0000
  serial number is FMS1021R2PQ
  fc-transmitter type is short wave laser w/o OFC (SN)
  fc-transmitter supports intermediate distance link length
  media type is multi-mode, 62.5m (M6)
  Supported speed is 400 MBytes/sec
  Nominal bit rate is 4300 MBits/sec
  Link length supported for 50/125mm fiber is 150 m(s)
  Link length supported for 62.5/125mm fiber is 70 m(s)
  cisco extended id is unknown (0x0)

no tx fault, no rx loss, in sync state, Diag mon type 104
SFP Detail Diagnostics Information
```

		Alarms		Warnings	
		High	Low	High	Low
Temperature	38.58 C	0.00 C	0.00 C	0.00 C	0.00 C
Voltage	3.31 V	0.00 V	0.00 V	0.00 V	0.00 V
Current	7.77 mA	0.00 mA	0.00 mA	0.00 mA	0.00 mA
Tx Power	-4.33 dBm	N/A	N/A	N/A	N/A
Rx Power	-4.60 dBm	N/A	N/A	N/A	N/A
Transmit Fault Count = 0					

Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning

Authentication successful

Close

Example Transceiver Details Commands

10 Gig at 10KM

```
switch# sh int fc 7/1 transceiver details
fc7/1 sfp is present
  name is CISCO-INTEL
  part number is TXN179092013I21
  revision is E1
  serial number is INT0915M07L
  FC Transceiver Type is X2 Medium
  FC Connector Type is SC
  Bit Encoding is NRZ
Protocol Type is 10GFC
  Standards Compliance Codes :
  10GbE Code Byte 0 : 10GBASE-LR
Fiber type Byte 0 : SM, Generic
  Fiber type Byte 1 : Unspecified
  Transmission Range is 3 (in 10m increments)
  cisco extended id is Unknown (0x0)
  no tx fault, no rx loss, in sync state,
  Diag mon type 209
  SFP Detail Diagnostics Information
```

		Alarms		Warnings	
		High	Low	High	Low

Temperature	34.61 C	70.00 C	0.00 C	70.00 C	0.00 C
Voltage	0.00 V	0.00 V	0.00 V	0.00 V	0.00 V
Current	7.17 mA	6.00 mA	2.00 mA	14.00 mA	2.40 mA
Tx Power	-0.57 dBm	1.00 dBm	-8.21 dBm	0.00 dBm	-7.21 dBm
Rx Power	-4.50 dBm	1.00 dBm	-16.58 dBm	0.00 dBm	-14.44 dBm

Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning					

Primitive Sequences

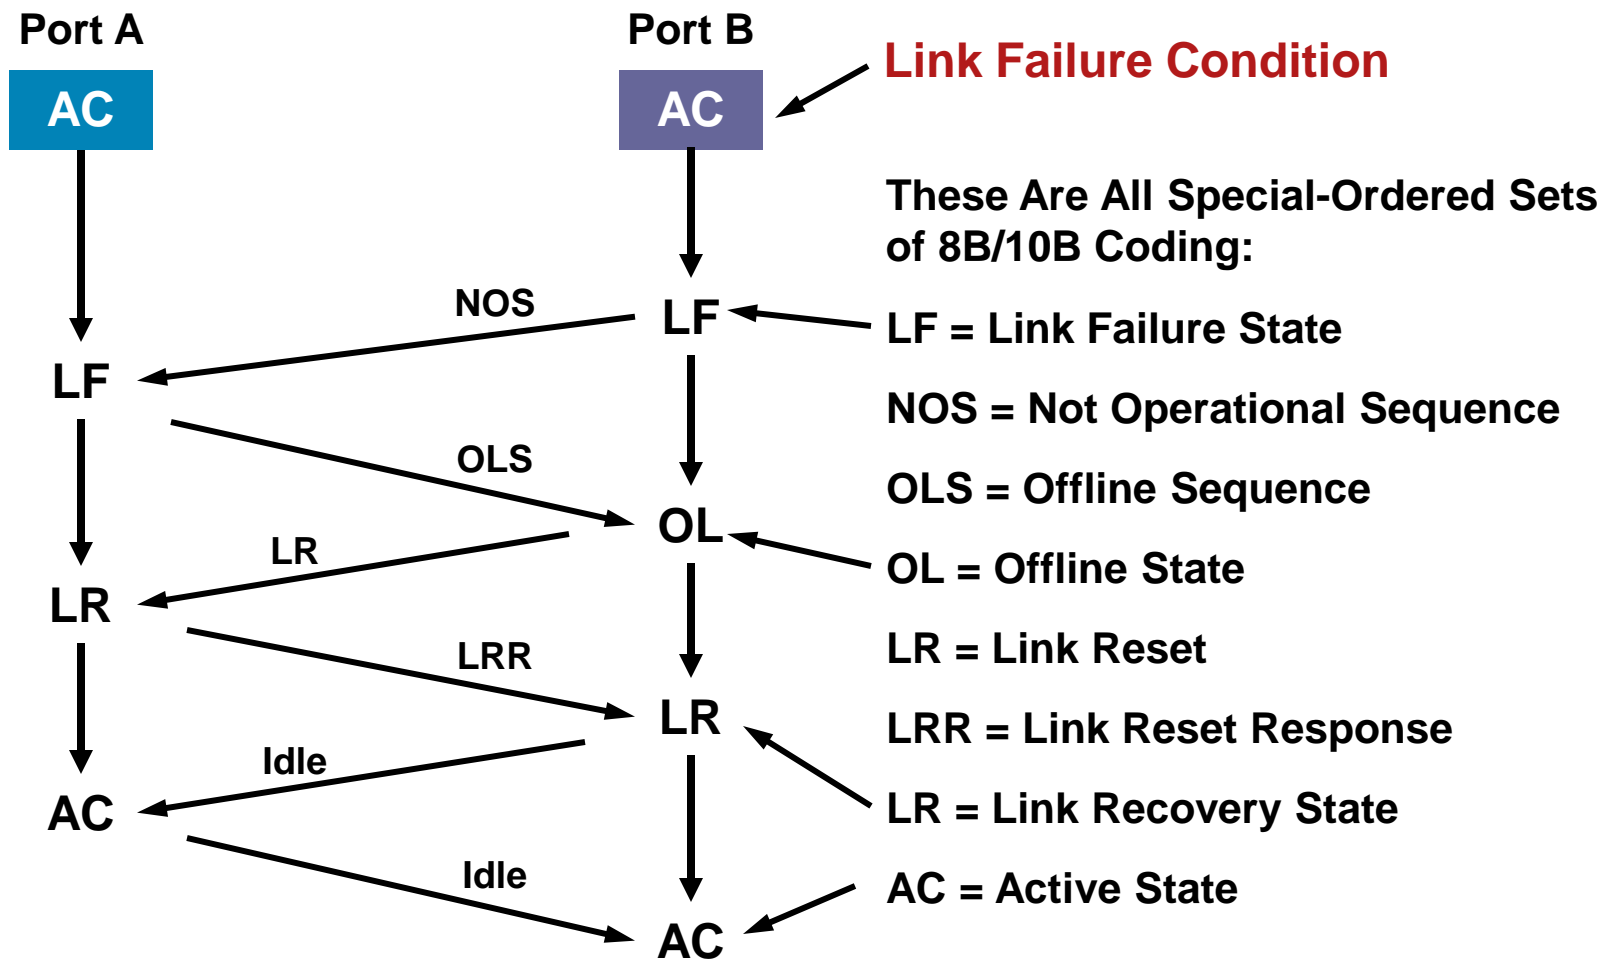
- Ordered set that is transmitted continuously to indicate that specific conditions within the Port are encountered
- Examples:
 - NOS – Not Operational Sequence**
 - OLS – Offline Sequence**
 - LR – Link Reset**
 - LRR - Link Reset Response**
 - LIP – Loop Initialization**
 - LPB – Loop Port Bypass**
- Four primitive sequences which can determine where problem exist

Primitive Sequence Meanings and Response

Currently Transmitting	Meaning	Transmit in Response
NOS	Not Operational <ul style="list-style-type: none"> ▪ Link Failure 	OLS
OLS	Offline State <ul style="list-style-type: none"> ▪ Internal port failure ▪ Transmitter power down, perform diags, or perform initialization ▪ Receiver shall ignore Link error or Link Failure 	LR
LR	Link Reset <ul style="list-style-type: none"> ▪ Remove class_1 Conn ▪ Reset F_Port ▪ OLS recognized 	LRR
LRR	Link Reset Response <ul style="list-style-type: none"> ▪ Link Reset Recognized 	Idles
IDLE	Operational Link <ul style="list-style-type: none"> ▪ Idles and R_RDY recognized 	Idles or R_RDY

Link Failure Flow

Fibre Channel Layer 1 Protocol (FC-1)



Fibre Channel Port Issues

In Order for an F_Port to Come Up on an MDS Switch:

1. The switch port must first acquire bit and word synchronization with the N_Port
2. N_Port must issue a FLOGI to the MDS

**Primitive Sequences
Counters Can Determine
Layer 0–1 Problems**

Tip:
Clear Counters and Monitor to Verify Active Issues, Use Device Manager Monitor Tool to Monitor Live; Set and Activate Threshold Manager to Alert You;

MDS_Switch# clear counters interface fc 1/1

```
MDS9509# show interface fc1/11
fc1/11 is up
  Port description is Smelt
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:0b:00:0d:65:b4:2b:40
  Admin port mode is FX
  sump link state traps are enabled
  Port mode is F, FCID is 0xe30003
  Configured Port vsan is 2
  Port vsan is 2
  Speed is 4 Gbps
  Rate mode is dedicated
  Transmit B2B Credit is 16
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  541397219 frames input, 837031650172 bytes
  0 discards, 0 errors
  0 CRC, 0 unknown class
  0 too long, 0 too short
  593258648 frames output, 820603001440 bytes
  4 discards, 0 errors
  7 input OLS, 7 LRR, 0 NOS, 123 loop inits
  51 output OLS, 38 LRR, 36 NOS, 74 loop inits
  16 receive B2B credit remaining
  16 transmit B2B credit remaining
```

Line Card Basics

- Show module will display slot locations and type of card
- General information in interface and statistics can be gathered from switch main-level prompt
- More detailed gathering of ASIC counters may be required to troubleshoot difficult issues, attach would then be required

Attach to Module with “Attach” Command, All Modules Can Be Attached to, Including Standby Supervisor and IPS Blade

Exit to Detach

```
MDS9509# show module
Mod Ports Module-Type          Model          Status
-----
1  24   1/2/4 Gbps FC Module         DS-X9124       ok
3  32   Storage Services Module      DS-X9032-SSM   ok
4  4    10 Gbps FC Module            DS-X9704       ok
5  0    Supervisor/Fabric-1          DS-X9530-SF1-K9  ha-standby
6  0    Supervisor/Fabric-1          DS-X9530-SF1-K9  active *
8  4    IP Storage Services Module    DS-X9304-SMIP   ok
9  32   1/2 Gbps FC Module           DS-X9032       ok
```

```
Mod Sw      Hw      World-Wide-Name(s) (WWN)
-----
1  3.1(2)    1.4     20:01:00:0d:65:b4:2b:40 to 20:18:00:0d:65:b4:2b:40
3  3.1(2)    5.1     20:81:00:0d:65:b4:2b:40 to 20:a0:00:0d:65:b4:2b:40
4  3.1(2)    0.3     20:c1:00:0d:65:b4:2b:40 to 20:c4:00:0d:65:b4:2b:40
5  3.1(2)    4.0     --
6  3.1(2)    4.0     --
8  3.1(2)    1.0     --
9  3.1(2)    1.0     22:01:00:0d:65:b4:2b:40 to 22:20:00:0d:65:b4:2b:40
```

```
Mod MAC-Address(es)          Serial-Num
-----
1  00-18-18-d1-f7-f4 to 00-18-18-d1-f7-f8  JAB102200H4
3  00-13-19-57-a2-d2 to 00-13-19-57-a2-d6  JAB090403QE
4  00-05-30-01-cb-36 to 00-05-30-01-cb-3a  JAB092906RB
5  00-0c-30-0c-fd-ac to 00-0c-30-0c-fd-b0  JAB073404HS
6  00-0c-30-0c-fl-ec to 00-0c-30-0c-fl-f0  JAB073306ZW
8  00-0e-38-c5-ed-00 to 00-0e-38-c5-ed-08  JAB081904MY
9  00-0b-be-f7-3d-54 to 00-0b-be-f7-3d-58  JAB065204YW
```

* this terminal session

```
MDS9509# attach mod 1
Attaching to module 1 ...
To exit type 'exit', to abort type '$.'
Last login: Mon Apr 23 10:59:17 2007 from sup2 on pts/0

module-1#
```

Monitor Link Init State – Gen 1 LC

(Note: Flow from Bottom to Top)

Attached to Mod 1

FC Port 1/2

module-1# show hardware internal fc-mac port 2 stateinfo

F-Port
Point
to Point
Negotiated

```
LINK: 052 022180c5 LR_RECEIVE(03) =>ACTIVE(01)
LINK: 051 022180c3 OLS_TRANSMIT(07) =>LR_RECEIVE(03)
LINK: 050 022180c1 SENTINAL(00) =>OLS_TRANSMIT(07)
```

Port Tries
Loop First
When Port Is
Set to Auto

```
LOOP: 049 022180c1 HW_ALPAS(0d) MASTER_LISA_WAIT(1f)=> OLD_PORT(3f)
LOOP: 048 022180c1 HW_ALPAS(0d) MASTER_LISA(1e)=> MASTER_LISA_WAIT(1f)
LOOP: 047 022180c1 HW_ALPAS(0d) MASTER_LIHA_WAIT(1d)=> MASTER_LISA(1e)
LOOP: 046 022180c1 HW_ALPAS(0d) MASTER_LIHA(1c)=> MASTER_LIHA_WAIT(1d)
LOOP: 045 022180c1 HW_ALPAS(0d) MASTER_LIPA_WAIT(1b)=> MASTER_LIHA(1c)
LOOP: 044 022180c1 HW_ALPAS(0d) MASTER_LIPA(1a)=> MASTER_LIPA_WAIT(1b)
LOOP: 043 022180c1 HW_ALPAS(0d) MASTER_LIFA_WAIT(19)=> MASTER_LIPA(1a)
LOOP: 042 022180c1 HW_ALPAS(0d) MASTER_LIFA(18)=> MASTER_LIFA_WAIT(19)
LOOP: 041 022180c1 HW_X_ARB(0c) MASTER_START(17)=> MASTER_LIFA(18)
LOOP: 040 022180c1 HW_LISM0(0a) OPEN_INIT_SELECT_MASTER(06)=> MASTER_START(17)
LOOP: 039 022180bf HW_R_LIP(09) OPEN_INIT_START(05)=> OPEN_INIT_SELECT_MASTER(06)
LOOP: 038 022180bf HW_X_LIP(08) NORMAL_INITIALIZE(04)=> OPEN_INIT_START(05)
LOOP: 037 022180bf HW_R_LIP(09) LPSM_STARTED(01)=> NORMAL_INITIALIZE(04)
LOOP: 036 022180b0 HW_OLDP(07) LPSM_DISABLED(00)=> LPSM_STARTED(01)
LINK: 035 022170e8 ACTIVE(01) =>SENTINAL(00)
```

Interface Shut/
No Shut via
Configuration

Monitor Link Init State - Gen-2 LC

The newer Generation-2 Line Cards (4 & 8 Gig) support a different look to debugging the link events
Still requires attach to the module

module-7# show hardware internal fc-mac port 1 link-event

```
===== FCP Port#1 Link State Machine Event Log =====
```

MMDDYY	HHMMSS	usecs	Event	Current State
032907	014953	369768 (0000)	E_LINK_IDLE	LINK_ACTIVE
032907	014953	368963 (0000)	E_LINK_LR	LINK_LR_RX
032907	014953	365690 (0000)	E_LINK_NOS	LINK_NOS_RX
032907	014953	365593 (0001)	E_LINK_MIN_OLS	LINK_OLS_TX
032907	014953	360463 (016D)	E_LINK_LINK_INIT	LINK_INIT
032907	014949	710690 (413C)	E_LINK_CLEANUP	LINK_DIS
				LINK_ACTIVE

Example - Identifying a NOS Occurred

```
module-4# sh port-config internal link-events
```

```
***** Port Config Link Events Log *****
```

Time	PortNo	Speed	Event	Reason
Apr 18 18:48:56 2006 00631612	fc4/16	2G	UP	
Apr 18 18:48:51 2006 00595873	fc4/16	---	DOWN	NOS Rcvd
Apr 18 18:48:36 2006 00621589	fc4/16	2G	UP	
Apr 18 18:48:31 2006 00592949	fc4/16	---	DOWN	NOS Rcvd
Apr 18 18:48:16 2006 00621613	fc4/16	2G	UP	

Determining Link-Event Reasons Reference

```

module-8# show process link-events
***** Port Config Link Events Log *****
-----
Time                               PortNo                               Speed  Link Event
-----
Jun 25 23:37:42 2004  00216931 fc8/14  ---  DOWN    (7)
Jun 25 23:37:42 2004  00770233 fc8/14  2G   UP
Jun 26 00:40:54 2004  00151363 fc8/14  ---  DOWN    (7)
Jun 26 00:40:54 2004  00710234 fc8/14  2G   UP

```

```

#define LINK_STATUS_REASON_NONE 0
#define LINK_STATUS_REASON_LINK_INIT_SUCCESS 1
#define LINK_STATUS_REASON_LOOP_INIT_SUCCESS 2
#define LINK_STATUS_REASON_LINK_RESET_OCCURRED 3
#define LINK_STATUS_REASON_PORT_INIT_FAILED 4
#define LINK_STATUS_REASON_PORT_UNUSABLE 5
#define LINK_STATUS_REASON_LOSS_OF_SIGNAL 6
#define LINK_STATUS_REASON_LOSS_OF_SYNC 7
#define LINK_STATUS_REASON_PEER_FCIP_CLOSED_TCP_CONNECTION 8
#define LINK_STATUS_REASON_PEER_FCIP_RESET_TCP_CONNECTION 9
#define LINK_STATUS_REASON_TCP_MAX_RETRANSMISSIONS 10
#define LINK_STATUS_REASON_TCP_KEEP_ALIVE_TIMER_EXPIRED 11
#define LINK_STATUS_REASON_TCP_PERSIST_TIMER_EXPIRED 12
#define LINK_STATUS_REASON_ETHERNET_LINK_DOWN 13
#define LINK_STATUS_REASON_ETHERNET_INTERFACE_DOWN 14
#define LINK_STATUS_REASON_ADMIN_CONFIG_CHANGE 15
#define LINK_STATUS_REASON_MOS_RCVD 16
#define LINK_STATUS_REASON_OLS_RCVD 17
#define LINK_STATUS_REASON_DEBOUNCE_TIMEOUT 18

```

```

#define LINK_STATUS_REASON_LR_RCVD_B2B 19
#define LINK_STATUS_REASON_CREDIT_LOSS 20
#define LINK_STATUS_REASON_RX_QUE_OVERFLOW 21
#define LINK_STATUS_REASON_USER_REQ 22
#define LINK_STATUS_REASON_TOO_MANY_INTR 23
#define LINK_STATUS_REASON_BIT_ERR_RT_THRES_EXCEEDED 25
#define LINK_STATUS_REASON_LIP_F8_RCVD 26
#define LINK_STATUS_REASON_LINE_CARD_PORT_SHUTDOWN 27
#define LINK_STATUS_REASON_LIP_RCVD_B2B 28
#define LINK_STATUS_REASON_OPNY_TMO_B2B 29
#define LINK_STATUS_REASON_OPNY_RET_B2B 30
#define LINK_STATUS_REASON_CREDIT_LOSS_B2B 31
#define LINK_STATUS_REASON_FRAME_FLOW_HALTED_TOO_LONG 32
#define LINK_STATUS_REASON_MAC_ERR 33
#define LINK_STATUS_REASON_GBIC_REMOVED 34
#define LINK_STATUS_REASON_GBIC_INSERTED 35
#define LINK_STATUS_REASON_LRR_RCVD_B2B 36

```


Understanding FC Addressing and Principal Switch



FCP Login Fundamentals as Seen w/fcanalyzer

Fabric Login to Login Server

Fabric ID assigned by Login Server

Port Login to Fabric Nameserver

State Change Registration with Fabric Controller

Query Nameserver for FC Type 4 devices (Storage)

PLOGI to storage

```
Small2(config)# fcanalyzer local brief
Warning: Couldn't obtain netmask info (eth2: no IPv4 address assigned).
Capturing on eth2
0.000000 00.00.00 -> ff.ff.fe 0x71 0xffff FC ELS FLOGI
0.020091 ff.ff.fe -> 26.00.00 0x71 0x2f FC ELS ACC (FLOGI)
0.020760 26.00.00 -> ff.ff.fc 0x72 0xffff FC ELS PLOGI
0.020985 ff.ff.fc -> 26.00.00 0x72 0x30 FC ELS ACC (PLOGI)
0.021102 26.00.00 -> ff.ff.fd 0x73 0xffff FC ELS SCR
0.027508 ff.ff.fd -> 26.00.00 0x73 0x31 FC ELS ACC (SCR)
0.030566 26.00.00 -> ff.ff.fc 0x74 0xffff dNS RFT_ID
0.033908 ff.ff.fc -> 26.00.00 0x74 0x32 dNS ACC (RFT_ID)
0.034541 26.00.00 -> ff.ff.fc 0x75 0xffff dNS RFF_ID
0.038120 ff.ff.fc -> 26.00.00 0x75 0x33 dNS ACC (RFF_ID)
0.038766 26.00.00 -> ff.ff.fc 0x76 0xffff dNS RNN_ID
0.045064 ff.ff.fc -> 26.00.00 0x76 0x34 dNS ACC (RNN_ID)
0.045698 26.00.00 -> ff.ff.fc 0x77 0xffff dNS RSNN_NN
0.047067 ff.ff.fc -> 26.00.00 0x77 0x35 dNS ACC (RSNN_NN)
0.047709 26.00.00 -> ff.ff.fc 0x78 0xffff dNS GNN_FT
0.049927 ff.ff.fc -> 26.00.00 0x78 0x36 dNS ACC (GNN_FT)
0.050588 26.00.00 -> ff.ff.fc 0x79 0xffff dNS GPN_ID
0.052039 ff.ff.fc -> 26.00.00 0x79 0x37 dNS ACC (GPN_ID)
0.053892 26.00.00 -> 26.04.00 0x7a 0xffff FC ELS PLOGI
0.058804 26.00.00 -> 26.04.00 0x7a 0xffff FC ELS PLOGI
0.060112 26.04.00 -> 26.00.00 0x7a 0x8475 FC ELS ACC (PLOGI)
0.064908 26.04.00 -> 26.00.00 0x7a 0x8475 FC ELS ACC (PLOGI)
0.065572 26.04.00 -> 26.00.00 0x7a 0x8476 FC ELS ACC (PLOGI)
0.069811 26.04.00 -> 26.00.00 0x7a 0x8476 FC ELS ACC (PLOGI)
```

Reserved Addresses

- FC-PH has defined a block of addresses for special Functions:

High order 16 addresses in the 24 bit address space

Called the Well Known Addresses

Main addresses used today

FF FF FC Directory Server

FF FF FD Fabric Controller

FF FF FE Fabric F_Port which N_Port
 logins to.

FF FF FA Management

Domains



Principal Switch Selection

Who becomes Principal?

1. A switch with a non-empty domain ID list has priority over a switch that has an empty domain ID list, and the principal switch will be the principal switch of the first fabric. In the case of a single switch fabric, it does not contain a domain ID list.
2. If both fabrics have a domain ID list, the priority between the two principal switches is determined by configured switch priority. This is a user-settable parameter - the lower the value the higher the priority.
3. 3. If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the World Wide Names of the two switches. The lower value has the higher priority.


The Domain ID

Domain IDs are assigned by the principal switch based on the non-principal switch's requesting domain ID.

If it is available, the domain ID is assigned to that switch. If not, a domain ID is picked from a "Least Recently Used" free list.

On a fresh switch, the search for the free domain starts from 239 and goes in decreasing order.

Before a switch ever joins a fabric, each switch assigns itself a domain ID based on its configured domain ID. If the configured domain ID type is preferred and configured domain ID is 0, then it assigns itself a random domain ID.



Domain ID's

Running Domains in a VSAN

Switch	State	DomainId	Local WWN	Local Priority	Principal WWN	Principal Priority
Fishfry	stable	0x5(5)	20:64:00:0d:65:b4:2b:41	128	Cisco 20:64:00:0c:85:67:b1:c1	2
BEAR	stable	0xa(10)	20:64:00:0c:85:67:b1:c1	2	Cisco 20:64:00:0c:85:67:b1:c1	2

Configured Domains in a VSAN

Switch	Enable	Running DomainId	Config DomainId	Config Type	FabricName	Priority	Contiguous Allocation	Auto Reconfigure	Restart
Fishfry	<input checked="" type="checkbox"/>	0x5(5)	5	static	20:01:00:05:30:00:28:df	128	<input type="checkbox"/>	<input type="checkbox"/>	nonDisrup...
BEAR	<input checked="" type="checkbox"/>	0xa(10)	0	preferred	20:01:00:05:30:00:28:df	128	<input type="checkbox"/>	<input type="checkbox"/>	nonDisruptive disruptive NoSelection

Choices for action to configure the Domain ID. If not Disruptive, then will Change upon next disruption

BF RCF

Duplicate Domain ID's but different VSAN's

If a requirement to run IVR between VSAN's with same ID's then IVR NAT would be required

```
BEAR# show fcdomain domain-list
```

```
VSAN 1
Number of domains: 3
Domain ID          WWN
-----
0xec(236)         20:01:00:0c:85:67:b1:c1 [Local] [Principal]
0x62(98)          20:01:00:0d:65:b4:2b:41
0x05(5)           20:01:00:0b:be:77:72:41

VSAN 3
Number of domains: 2
Domain ID          WWN
-----
0x03(3)           20:03:00:0c:85:67:b1:c1 [Local] [Principal]
0x80(128)         20:03:00:0d:65:b4:2b:41

VSAN 100
Number of domains: 2
Domain ID          WWN
-----
0x0a(10)          20:64:00:0c:85:67:b1:c1 [Local] [Principal]
0x05(5)           20:64:00:0d:65:b4:2b:41

VSAN 200
Number of domains: 2
Domain ID          WWN
-----
0xd1(209)         20:c8:00:0b:be:77:72:41 [Principal]
0x80(128)         20:c8:00:0c:85:67:b1:c1 [Local]
```

Local is the switch I am connected to

Principal Switch Priority- Operation

```
MDS-switch(config)# fcdomain priority 1 vsan 10  
MDS-switch(config)# do show fcdomain vsan 10  
The local switch is a Subordinated Switch.
```

Local switch run time information:

```
State: Stable  
Local switch WWN: 20:0a:00:05:30:01:97:43  
Running fabric name: 20:0a:00:05:30:00:49:df  
Running priority: 128  
Current domain ID: 0x4d(77)
```

Local switch configuration information:

```
State: Enabled  
FCID persistence: Enabled  
Auto-reconfiguration: Disabled  
Contiguous-allocation: Disabled  
Configured fabric name: 20:01:00:05:30:00:28:df  
Optimize Mode: Disabled  
Configured priority: 1  
Configured domain ID: 0x00(0) (preferred)
```

Principal switch run time information:

```
Running priority: 4
```

```
MDS-switch(config)# fcdomain restart vsan 10
```

Priority set per VSAN, Try to reserve priority 1 for reasons of capabilities to trump any switch

Does not immediately go into effect , Still at default 128

Configured priority

Priority of principal switch

Must do a non-disruptive (BF) restart

Domain ID - Be in the Know!

- **Per VSAN**

 - It is per VSAN configuration

- **Build Fabric or Reconfigure Fabric**

 - 2 choices when defining a Domain ID per FC-SW standards

 - BF= non disruptive to complete Fabric, does force changes that could effect logged in devices to re login.

 - RCF= Disruptive to Fabric, loss of complete path to a switch

- **Interop modes and effects on Domain ID assignments**

 - Depending on interop mode, range of domains may be limited to 97-127 due to McData's 31 DomainID limit

- **Planning, proper administration**

 - Consider no duplicate domain ID's across all VSAN's if plan may include someday IVR, then NAT will not be required

 - Most best practice is to configure core Domain ID's and principle switch primary & secondary

Gotcha - HBAs with Unique Area Requirements

Some HBA vendors require that the target they are logging into be in a different area, for example: Qlogic. Therefore, each HBA needs to be assigned a unique area versus default behavior is to reused the area portion of the FCID.

These are the OUI's
that will receive
unique area
assignments

```
rtp9-cae-9124-3a# sh fcid-allocation ?  
area          Show information about fcid-allocation list  
company-id-from-wwn  Company id (or OUI).  
rtp9-cae-9124-3a# sh fcid-allocation area
```

Fcid area allocation company id info:

```
00: 2:6B  
00: 6:2B  
00: 9:6B  
00:11:25  
00:14:5E  
00:50:2E  
00:50:8B  
00:60:B0  
00:A0:B8  
00:D0:60  
00:D0:B2  
00:E0:69  
00:E0:8B
```

Use this command
to add an OUI if not
already seen above

```
rtp9-cae-9124-3a(config)# fcid-allocation area company-id ?  
<0x0-0xfffff> Enter company id(or OUI) (0xhhhhhh)
```

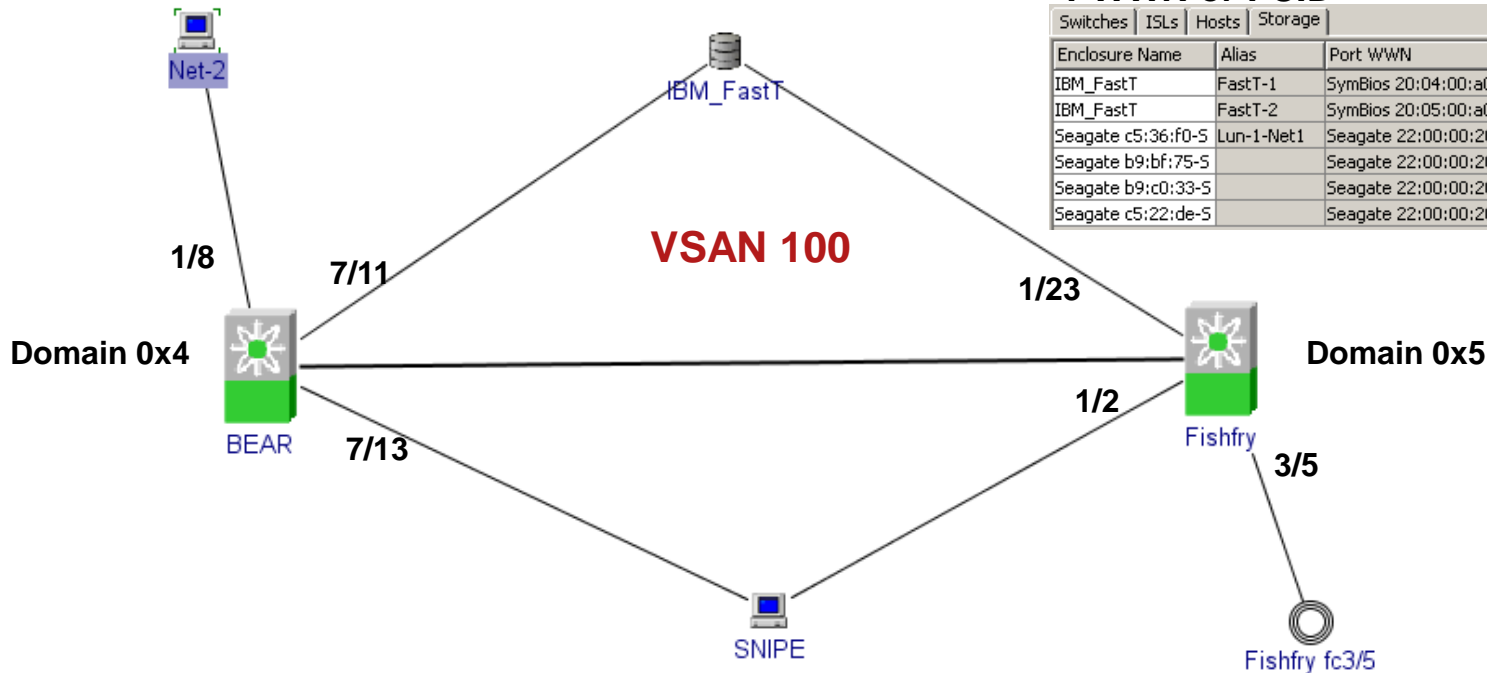
Zoning



Zoning – Basic Operational Understanding

- What can we Zone
- RSCN's and Zones
- Zoning Standards – Basic vs. Enhanced
- Zone Distribution, Export, Import, Merge

Physical Points for Zone Members



PWWN or FCID

Switches	ISLs	Hosts	Storage		
Enclosure Name	Alias	Port WWN	FcId	Switch Interface	
IBM_FastT	FastT-1	SymBios 20:04:00:a0:b8:0c:64:51	0x050300	Fishfry fc1/23	
IBM_FastT	FastT-2	SymBios 20:05:00:a0:b8:0c:64:51	0x040200	BEAR fc7/11	
Seagate c5:36:f0-5	Lun-1-Net1	Seagate 22:00:00:20:37:c5:36:f0	0x0501e1	Fishfry fc3/5	
Seagate b9:bf:75-5		Seagate 22:00:00:20:37:b9:bf:75	0x0501e4	Fishfry fc3/5	
Seagate b9:c0:33-5		Seagate 22:00:00:20:37:b9:c0:33	0x0501e8	Fishfry fc3/5	
Seagate c5:22:de-5		Seagate 22:00:00:20:37:c5:22:de	0x0501e2	Fishfry fc3/5	

Switch FC Interface or Fabric Port WWN (FWWN)

BEAR	fc1/8	none	<input type="checkbox"/>	20:08:00:0c:85:67:b1:c0
BEAR	fc7/13	none	<input type="checkbox"/>	21:8d:00:0c:85:67:b1:c0
BEAR	fc7/11	none	<input type="checkbox"/>	21:8b:00:0c:85:67:b1:c0
BEAR	fc8/1	channel1	<input type="checkbox"/>	21:c1:00:0c:85:67:b1:c0
BEAR	fc8/2	channel1	<input type="checkbox"/>	21:c2:00:0c:85:67:b1:c0

Zoning Choices

Zone By:

- WWN
- Switch & Port
- Domain & Port
- iSCSI IP Address/Subnet
- iSNS Host
- deviceAlias
- FcId
- Switch Port WWN
- iSCSI Name
- iSCSI Proxy
- fc-Alias

RSCNs and Zoning Smooth Operation

- Devices must register with MDS if they would like to receive RSCN.
- MDS generates local RSCN to devices within the affected zone when a zoned member logs in or out of the fabric.
- Devices that log in or out of the fabric and are not part of the zone, will not generate RSCNs to devices in the zone.
- SW_RSCNs will be sent to all switches, and they in turn will decide if a local RSCN needs to be sent (based on zoning and affected devices)

Zone Server Modes

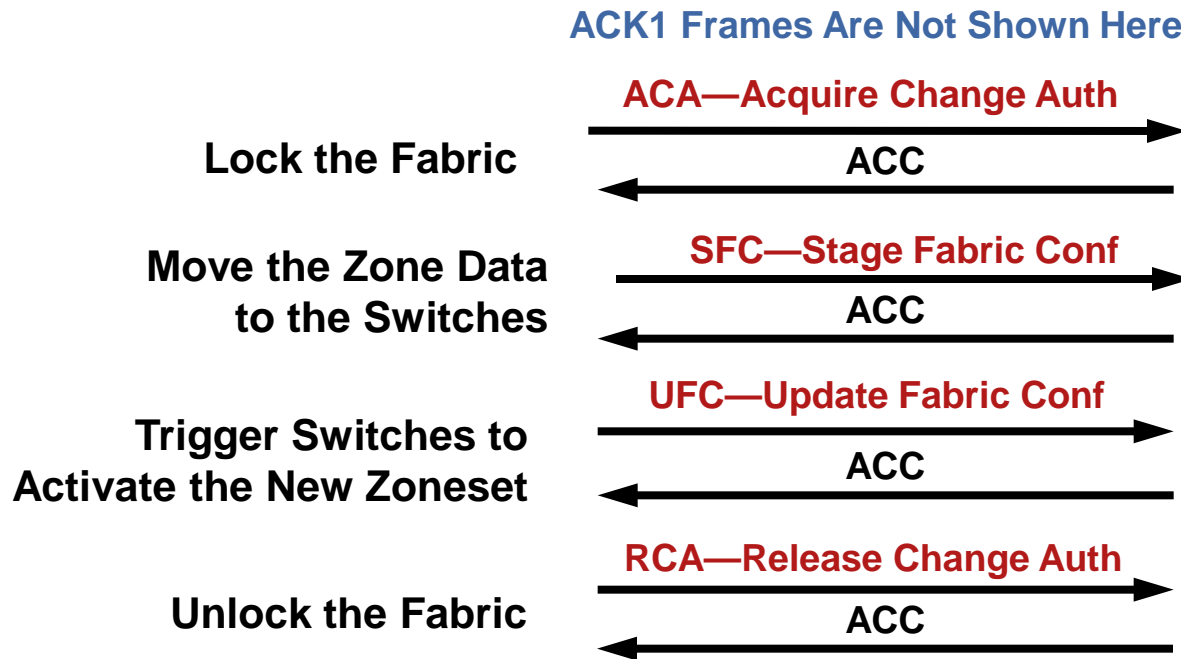
- Zone server supports 2 different modes

Basic mode – represents the zone server behavior of FC-GS3/FC-SW2 standard. All SAN-OS support basic mode.

Enhanced mode – represents the zone server behavior of FC-GS4/FC-SW3 standard. SAN-OS 2.0 and later are required for enhanced mode.

Activate Zoneset Flow Across ISLs

Zoneset Distribution Would Go to Every Domain Within the VSAN if successful

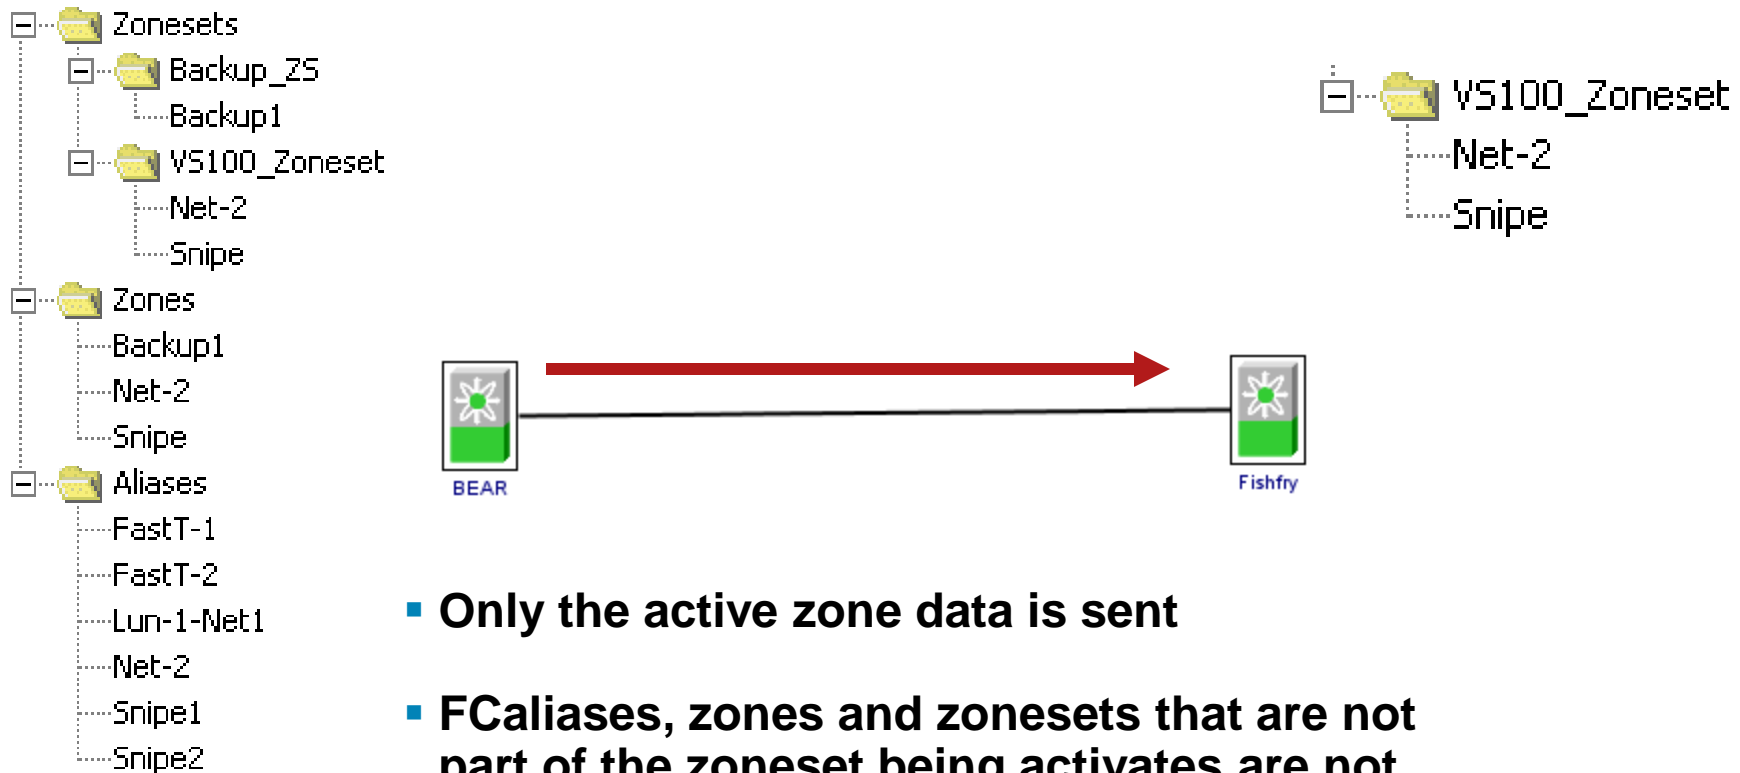


This flow can be seen with the fcanalyzer. If zoneset activation is failing, this flow above will show where the process is failing

Enhanced Zoning

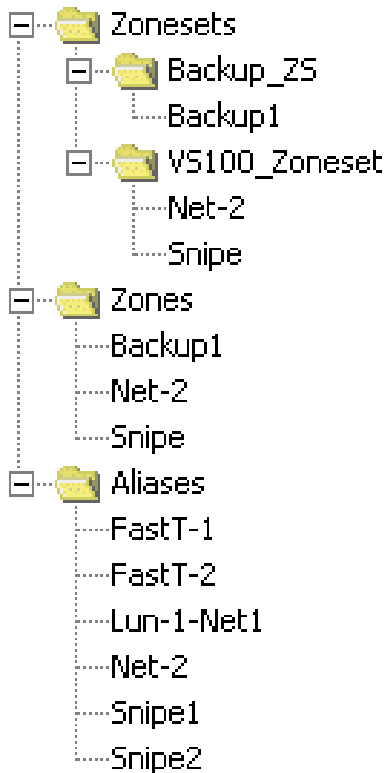
- Enhanced zoning provides the following advantages
 - Disallow parallel configuration attempts
 - Standardized generation of RSCN
 - Reduced payload size of the SFC frame
 - Fabric-wide policy enforcement (default zone, merge control)
 - Enhanced error reporting
 - Distributing zonesets without activation
 - Unique Vendor Types
 - FWWN based member type standardized
 - Enhanced interop thru ESS (Exchange Switch Support) defined in SW-3

Active Only Zoneset Activation



- Only the active zone data is sent
- FCaliases, zones and zonesets that are not part of the zoneset being activated are not sent
- The running configuration on Switch Fishfry will not show the active zoning information.

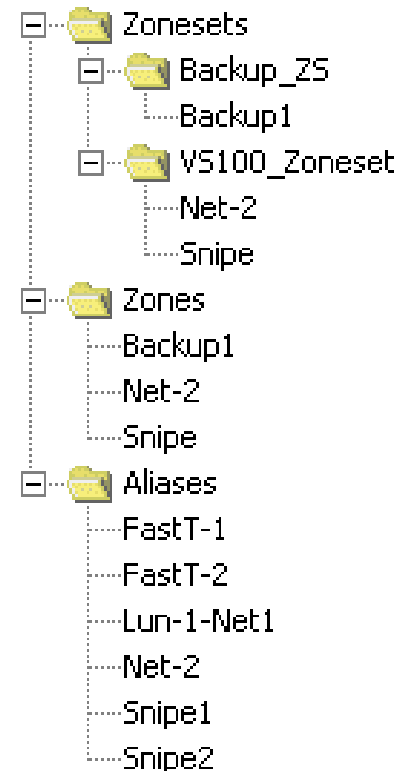
Full Zoneset Activation



BEAR



Fishfry



- **The complete zone data is sent**
- **FCaliases, zones and zonesets that are not part of the zoneset being activated are sent**
- **The running configuration on Fishfry will show the active zoning information.**

Recommendations

- If the SAN administrators wish to be able to manage zones from any switch in the fabric, then configure all switches/VSANs for 'distribute full zone database'.
- If the SAN administrator wishes to manage zone changes from only 1 switch in the fabric, then they can leave the default configuration of 'distribute active zone database only'.
- Inconsistent zone distribution policies can cause problems when a zoneset is modified on a switch that may not have the most current information in its configuration when the change was made.

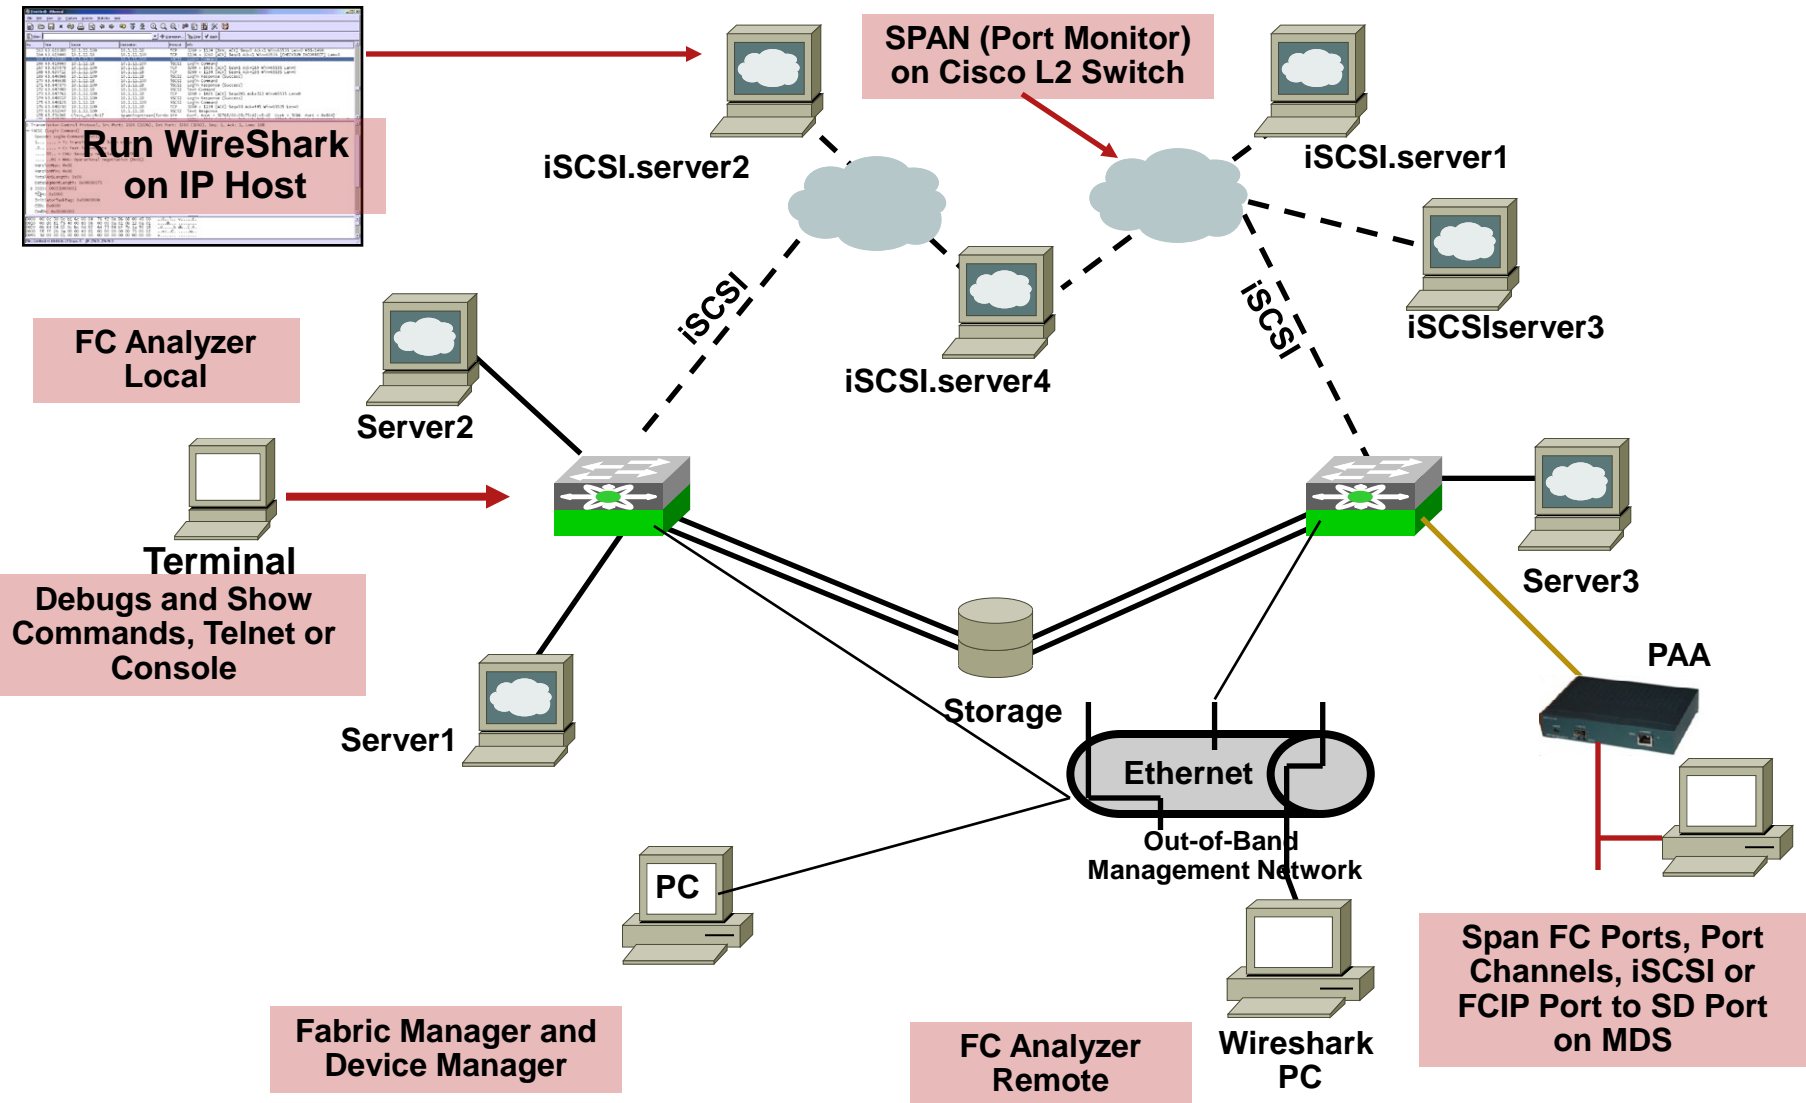
Questions So Far?



MDS9000 SAN Tools

- **FCanalyzer**
- **SPAN & PAA** (WireShark usage)
- **SAN/OS** (Output analysis, debug, logs, Cores)
- **Performance Manager** (Licensed part of Fabric Manager)
- **NTOP** (Using Netflow and SPAN w/PAA)

How Do We Troubleshoot the Network?



Wireshark Is Most Used Troubleshooting Tool

Wireshark (Once known as Ethereal) is part of the SAN/OS system image and can be run directly on the switch via ssh/telnet. (FCAnalyzer command)

Wireshark on a PC with use of Cisco SPAN and a Port Analyzer Adaptor can be used as a inline tool with no disruption to traffic.

The combination of Wireshark on a PC with a PAA can give complete look at the flow beyond the FLOGI/PLOGI process

We will look more at Wireshark best practice later in this session

The screenshot displays the Wireshark interface with a list of captured packets and a detailed view of a selected frame. The packet list shows various protocols including FC ELS, DNS, and FCP. The selected frame (No. 1) is an FC ELS FLOGI packet. The packet details pane shows the following information:

- Frame 1 (164 bytes on wire, 164 bytes captured)
- Ethernet II, Src: Andiamo_00:a2:9d (00:05:30:00:a2:9d), Dst: Broadcast (ff:ff:ff:ff:ff:ff)
- Boardwalk
- Fibre Channel
 - [Exchange Last In: 0]
 - R_CTL: 0x22(Extended Link Services/Request)
 - Dest Addr: ff.ff.fe
 - CS_CTL: 0x00
 - Src Addr: 00.00.00
 - Type: Ext Link Svc (0x01)
 - F_CTL: 0x290000 Exchange Originator, Seq Initiator, Exchg First, Seq Last, CS_CTL, Transfer Seq Initiative
 - SEQ_ID: 0x00
 - DF_CTL: 0x00
 - SEQ_CNT: 0
 - OX_ID: 0x033d
 - RX_ID: 0xffff
 - Parameter: 0x00000000
 - FC ELS
 - Cmd Code: FLOGI (0x04)
 - Common Svc Parameters
 - B2B Credit: 3
 - Common Svc Parameters: 0x0000 Normal B2B Credit Mgmt Payload Len=116 bytes

The packet bytes pane shows the raw data of the frame, including the Ethernet II header and the FC ELS payload.

Gathering Protocol Traces for Analysis

- **Using built-in FC Analyzer (CLI)**

Wireshark.com

- **Using Wireshark on PC** (local and remote)

- **Using the MDS port analyzer adapter w/SPAN**

- **Using an external FC Analyzer tester in line or with SPAN**

Non-Disruptive to Switch Operations and Traffic on the SAN

MDS FC Analyzer (SAN/OS Embedded)

- Output is displayed to the console in readable sniffer-like format
- Is only used to monitor Fibre Channel traffic to and from supervisor on the MDS9000
- Traffic-like fabric logins, FSPF routing, switch-to-switch control traffic
- Output can go direct to your console screen or to a workstation running WireShark program

Note: SPAN Is Used for FC Port to FC Port Monitoring

What Tool Captures What?

- fcanalyzer is a tool that will collect frames that are destined to, or originate from the MDS supervisor. Node to switch, or switch to switch traffic can be seen with this tool.
- SPAN is a feature whereby frames that are transient to the switch, will be copied to a second port for analysis. Node to switch or node to node traffic can be seen via this method.

FC Analyzer Options

- **Local** or **remote**—Where to send the trace. Can be to local devices or remote PAA attached to different MDS switch.
- **Brief** or **detailed**—Header information vs. full output of frame including hex. Detail is default.
- **Limit-captured-frames**—Number of frames to capture. Default is FC analyzer will trace 100 frames. Specifying zero is unlimited frame capture.
- **Limit-capture-size**—Allows to capture N number of bytes of frame. Useful for not capturing frame data when it is not relevant to troubleshooting.

Display-Filter Options

- When not specified, FC analyzer will capture all traffic on VSAN 1

Example: `fcanalyzer local brief`

- To specify a different VSAN, use a display-filter and specify the VSAN to be captured in hex or decimal

Example for 100 (note the two equal signs):

```
fcanalyzer local display-filter mdshdr.vsan==0x64 or  
mdshdr.vsan==100
```

- To capture for a specific address in VSAN 100:

Example:

```
fcanalyzer local brief display-filter ((fc.d_id==64.01.00)or(fc.s_id==64.01.00))
```

- Recommend to use remote capture method for ease of filter capabilities on the WireShark GUI

Using Write Option for fcanalyzer

- Using write option sends output of the fcanalyzer to a file on the switch, directory called volatile. This trace file can be then copied off MDS switch and viewed with WireShark app on PC

Example: Capture 250 frames of all traffic on VSAN 200

```
fcanalyzer local display-filter mdshdr.vsan==0xc8 write  
volatile:capture.trc limit-captured-frames 250
```

- The file name on volatile: filesystem will have extra characters appended. Issue following command to see contents of filesystem and then copy the file via tftp/ftp

dir volatile:

```
copy volatile:capture_00001_20090321172628.trc tftp://<tftp  
server ip addr>/capture.trc
```

**Added Characters
When File Is Written
show date**

FCAnalyzer Local Brief

Capture Is Done in Configuration Mode

Brief Option Used to get Single Line Caption Along with a Display Filter to Narrow Output to only VSAN 2

Control-C Stops Trace Capture if number of frames are not specified

- Display filters are a must for narrowing output on a busy network. See MDS Config Guide for other filter types.
- Capture is done by default to console screen, so make sure you are able to save output to large capture buffer or log with your telnet application

```
TOP-9216i# conf t
Enter configuration commands, one per line. End with CNTL/Z.
TOP-9216i(config)# fcanalyzer local brief display-filter mdshdr.vsan==2
Warning: Couldn't obtain netmask info (eth2: no IPv4 address assigned).
Capturing on eth2
```

```
2.829871 00.00.00 -> ff.ff.fe 0x2288 0xffff FC ELS FLOGI
2.853261 ff.ff.fe -> 7e.01.00 0x2288 0xc728 FC ELS ACC (FLOGI)
2.853422 7e.01.00 -> ff.ff.fc 0x22a0 0xffff FC ELS PLOGI
2.853592 ff.ff.fc -> 7e.01.00 0x22a0 0xc729 FC ELS ACC (PLOGI)
2.853565 7e.01.00 -> ff.ff.fd 0x22b8 0xffff FC ELS SCR
2.859648 ff.fc.7e -> 7e.01.00 0xc72c 0xffff FC ELS PLOGI
2.860885 7e.01.00 -> ff.fc.7e 0xc72c 0x22d0 FC ELS ACC (PLOGI)
2.861007 ff.fc.7e -> 7e.01.00 0xc72b 0xffff FC ELS PRLI
2.861175 7e.01.00 -> ff.fc.7e 0xc72b 0x22e8 FC ELS ACC (PRLI)
2.862053 7e.01.00 -> ff.ff.fc 0x2300 0xffff dNS RFT_ID
2.865904 ff.fc.7e -> ff.fc.0a 0xc72e 0xffff SW_ILS SW_RSCN
2.865981 ff.fc.7e -> ff.fc.cb 0xc72f 0xffff SW_ILS SW_RSCN
2.866153 ff.fc.0a -> ff.fc.7e 0xc72e 0x77f9 FC Link Ctl, ACK1
2.866297 ff.fc.cb -> ff.fc.7e 0xc72f 0x77fa FC Link Ctl, ACK1
2.866445 ff.fc.0a -> ff.fc.7e 0xc72e 0x77f9 SW_ILS SW_ACC (SW_RSCN)
2.866496 ff.fc.7e -> ff.fc.0a 0xc72e 0x77f9 FC Link Ctl, ACK1
2.866615 ff.ff.fd -> 7e.01.00 0x22b8 0xc72a FC ELS ACC (SCR)
2.868792 ff.fc.cb -> ff.fc.7e 0xc72f 0x77fa SW_ILS SW_ACC (SW_RSCN)
2.868857 ff.fc.7e -> ff.fc.cb 0xc72f 0x77fa FC Link Ctl, ACK1
2.871132 ff.fc.7e -> 7e.01.00 0xc730 0xffff FC ELS LOGO
2.872013 7e.01.00 -> ff.fc.7e 0xc730 0x2318 FC ELS ACC (LOGO)
2.872021 7e.01.00 -> ff.fc.7e 0x2318 0xffff FC ELS PLOGI
2.872139 ff.fc.7e -> 7e.01.00 0x2318 0xc731 FC ELS LS_RJT (PLOGI)
2.872163 ff.fc.cb -> ff.fc.7e 0x77fb 0xffff dNS GE_ID
2.872234 ff.fc.7e -> ff.fc.cb 0x77fb 0xc732 FC Link Ctl, ACK1
2.891239 ff.fc.7e -> ff.fc.cb 0x77fb 0xc732 dNS ACC (GE_ID)
2.891359 ff.ff.fc -> 7e.01.00 0x2300 0xc72d dNS ACC (RFT_ID)
2.891469 7e.01.00 -> ff.ff.fc 0x2330 0xffff dNS RFF_ID
2.891613 ff.fc.cb -> ff.fc.7e 0x77fb 0xc732 FC Link Ctl, ACK1
2.900160 ff.ff.fc -> 7e.01.00 0x2330 0xc733 dNS ACC (RFF_ID)
2.900394 7e.01.00 -> ff.ff.fc 0x2378 0xffff dNS RNN_ID
2.901916 ff.ff.fc -> 7e.01.00 0x2378 0xc734 dNS ACC (RNN_ID)
2.902151 7e.01.00 -> ff.ff.fc 0x23a8 0xffff dNS RSNN_NN
2.908296 ff.ff.fc -> 7e.01.00 0x23a8 0xc735 dNS ACC (RSNN_NN)
2.908444 7e.01.00 -> ff.ff.fc 0x23c0 0xffff dNS GNN_FT
2.919880 ff.ff.fc -> 7e.01.00 0x23c0 0xc736 dNS ACC (GNN_FT)
```

SPAN and PAA



Use of SPAN Feature

- Used for **FC port to FC port** analyzing
- Same type of tool as used on Cisco Catalyst® products. Cisco Catalyst uses port monitor.
- Can be left configured on switch
- Ingress and egress ports are sent to an FC-port setup as a SPAN Destination (SD-port type)
- No limits to where the ports are located on the MDS switch network
- Used to output to third-party test equipment or to Cisco Port Analyzer Adapter

Best Practices Using SPAN

- MDS9500/9200
 - 16 SPAN sessions, Multi Source & Destinations
- MDS9124
 - Only 1 SPAN session, 1 Direction at a time
- MDS9020 no SPAN capabilities
- Try to design into the Fabric solution a SD port dedicated to for SPAN usage. 1 per fabric (w/ remote SPAN) or 1 per switch (w/o remote SPAN)
- Bladeswitch ports can also be configured for SPAN
- Use filters mnemonics with FCAnalyzer to limit capture
- Learn to use WireShark

FC Linecard



HP Blade



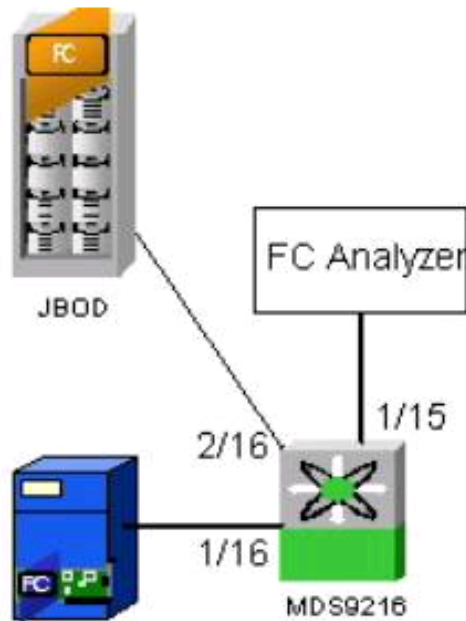
MDS9124



IBM Blade



Local Switch Port Analyzer (SPAN)



The intent is to collect FC frames to/from the host on the 9216 port 1/16 on the FC Analyzer on port 1/15.

```
MDS9216# sh run
vsan 13
vsan 13 interface fc1/16
vsan 13 interface fc2/16
```

```
interface fc1/15
switchport mode SD
switchport speed 2000
no shutdown
```

```
interface fc1/16
no shutdown
```

```
span session 1
destination interface fc1/15
source interface fc1/16 rx
source interface fc1/16 tx
```

DS-PAA Settings Reference

Switch Settings

```
0001 1G NTM
1001 1G ETM
0101 1G STM
0011 1G DTM
0000 2G NTM
1000 2G ETM
0100 2G STM
0010 2G DTM
1111 1G MNM (Used for diagnostics only)
```

MODES

- **NTM** (No Truncate Mode) - FC Frames are passed without any modifications.
- **ETM** (Ethernet Truncate Mode) - Truncates FC frame to max 1496B by reducing payload size from 528 lines to 368 lines.
- **STM** (Shallow Truncate Mode) - Truncates FC frame to max 256B by reducing payload size from 528 lines to 58 lines.
- **DTM** (Deep Truncate Mode) - Truncates FC frame to max 64B by reducing payload size from 528 lines to 10 lines.

Any changes require a power off/on.

SAN/OS Tools



Command Line Debugging

- Available debugs depend on features enabled in SAN/OS
- Many, many different options to select when turning on debugs
- Where is it output going?
 - Logfile—Data file in switch memory
 - Capture to direct to screen via console, telnet or SSH
- Requires admin privileges to run debugs
- Debugs can only be run from CLI
- No debugging available from Fabric Manager or Device Manager

Debug Logging

1. **TOP-9216i# debug logfile networkers_debugs size 5000**

2. Display captured debug to screen

Tip: use show debug to see name of debug file

TOP-9216i# show debug logfile networkers_debug

3. Copy debug file off MDS to a server

TOP-9216i# copy log:networkers_debugs ftp:

Enter hostname for the ftp server: 10.91.42.166

Enter username: networkers

Password: networkers

To delete the debug logfile

TOP-9216i# clear debug-logfile networkers_debugs

Or

TOP-9216i# undebug all

Or

debug logfile will be cleared and over written when next debug logfile is created, only one debug logfile is allowed by system

Debugs to Direct Telnet Window

- Use a telnet/SSH or console application that will capture the expected output to buffer or file
- **Undebug all** or **no debug** of specific debug command is required to turn trace off
- The debugs are not persistent across reboots
- Most debugs are very readable and sensible to understand, some not

Design for Troubleshooting

- Leverage VSAN design to support troubleshooting methodology
- Have a SPAN port allocated in port count needs
- Integrate a Cisco Port Analyzer Adapter (PAA) into SAN design - Cheap
- Provision an Analyzer (*Finisar, Agilent, other*) in the network, keep it operational with FTP access to extract traces - Expensive
- Design for Syslog servers and scheduled configuration saves

Performance Manager



Fabric Manager Performance Manager

- Licensed part of Fabric Manager
- Requirement for proactive SCSI I/O performance across the MDS Fabric
- Accounting logs to Monitor Fabric Changes
- Net Flow Tool
 - Create flows
 - Use for Setting Benchmarks for I/O use
- Use to monitor Port Groups
- Very powerful in examining performance issue macro or singular in the fabric

Accounting

Fabric BEAR > BEAR:Accounting

Showing 1-10 of 1850 records

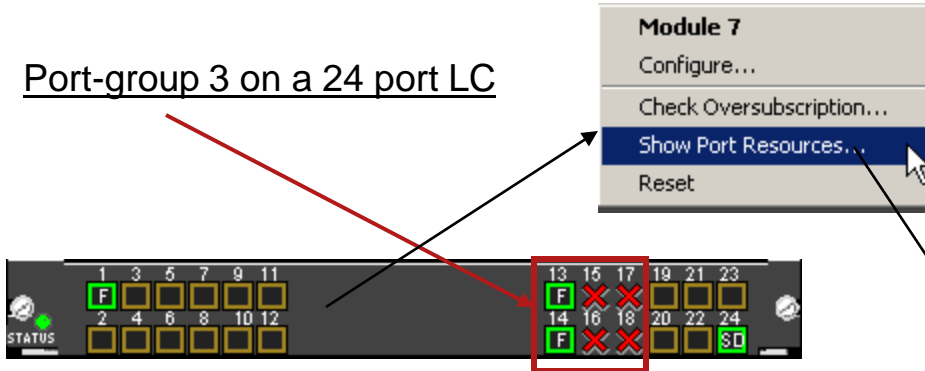
	Switch Time	Action	Protocol & Source	User	Description
1.	2007 May 01 14:51:15 LKT	stop	snmp_2107_64.101.213.216	admin	
2.	2007 May 01 14:51:15 LKT	update	snmp_2107_64.101.213.216	admin	target (name:64.101.213.216/1164/0 address:64.101.213.216:1164) deleted
3.	2007 May 01 13:12:31 LKT	update	snmp_1865_64.101.213.216	admin	Zone commit attempted on VSAN 100
4.	2007 May 01 13:12:27 LKT	update	snmp_1865_64.101.213.216	admin	Activation of zoneset VS100_Zoneset attempted on VSAN 100
5.	2007 May 01 12:53:32 LKT	update	snmp_2107_64.101.213.216	admin	Interface fc1/2 state updated to up
6.	2007 May 01 12:51:50 LKT	update	snmp_2107_64.101.213.216	admin	Interface fc1/2 state updated to down
7.	2007 May 01 12:48:58 LKT	update			Interface fc1/2 mode updated to E
8.	2007 May 01 12:48:49 LKT	update			Interface fc1/2 speed updated to 2 Gbps
9.	2007 May 01 12:48:38 LKT	update			Interface fc1/2 trunk mode updated to auto
10.	2007 May 01 12:45:38 LKT	update	snmp_1865_64.101.213.216	admin	Interface fc1/2 state updated to up

Rows per page: Go to page: of 185 Pages

- Web interface allows for simple accounting views

PM Output – Monitor Port Groups

Port-group 3 on a 24 port LC



Port-Group 3
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 4.8 Gbps
 Allocated dedicated bandwidth is 8.0 Gbps

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate Mode
fc7/13	16	4.0	dedicated
fc7/14	16	4.0	shared
fc7/15	16	4.0	dedicated
fc7/16	16	4.0	shared
fc7/17	16	4.0	shared
fc7/18	16	1.0	shared

Interface	Description	VSAN Id Port	VSAN Id Dynamic	Mode Admin	Mode Oper	Speed Admin	Speed Oper	Rate Mode	Status Service	Status Admin	Status Oper
fc7/13		100	n/a	FX	F	auto	1 Gb	dedicated	in	up	up
fc7/14		100	n/a	FX	F	auto	1 Gb	shared	in	up	up
fc7/15		3	n/a	FX	auto	4Gb	n/a	dedicated	in	up	down
fc7/16		1	n/a	FX	auto	4Gb	n/a	shared	in	up	down
fc7/17		1	n/a	FX	auto	4Gb	n/a	shared	in	up	down
fc7/18		1	n/a	FX	auto	1Gb	n/a	shared	in	up	down

Fabric BEAR > Core_SAN(100)

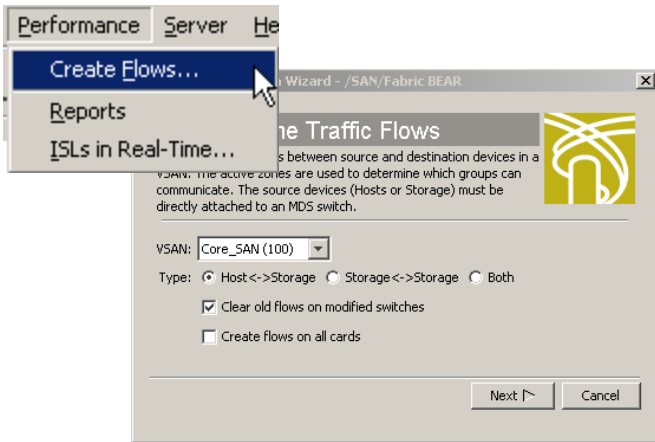
Type: Port Groups Last: 24 Hours Filter

Showing 1-4 of 4 records

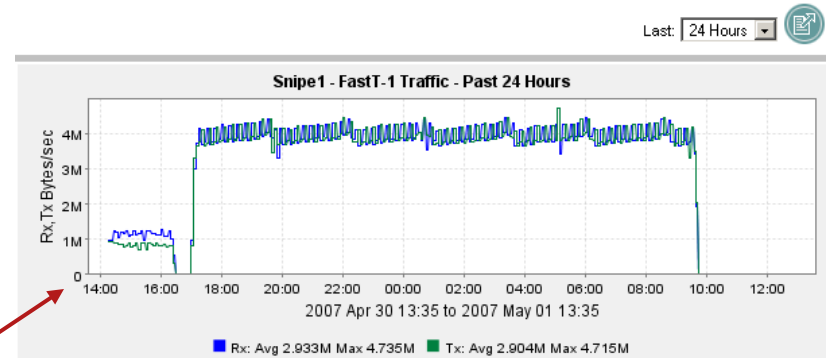
Name	Speed	Avg. Rx/sec	(Rx+Tx)/sec	Avg. Tx/sec	Peak Rx/sec	Peak Tx/sec	Errors	Discards	Last Updated
1. FastT-2_Snipe1	200.000MB	49.584KB	99.217KB	49.632KB	14.280MB	14.294MB	0	0	2007/05/16-13:54

PM- Monitor Netflows

Flow creation done with Flow Wizard



Historical look at Initiator to Target flow



Last: 24 Hours [Filter](#)

Showing 1-6 of 6 records

	Name	Avg. Rx/sec	Avg. Tx/sec	(Rx+Tx)/sec	Peak Rx/sec	Peak Tx/sec	Last Updated
1.	Snipe2<->FastT-1	3.011MB	2.980MB	5.991MB	4.738MB	4.717MB	2007/05/01-12:58
2.	Snipe1<->FastT-1	3.010MB	2.980MB	5.990MB	4.735MB	4.715MB	2007/05/01-12:58
3.	Net-2<->Lun-1-Net1	2.959MB	2.607MB	5.566MB	3.936MB	3.481MB	2007/05/01-12:58

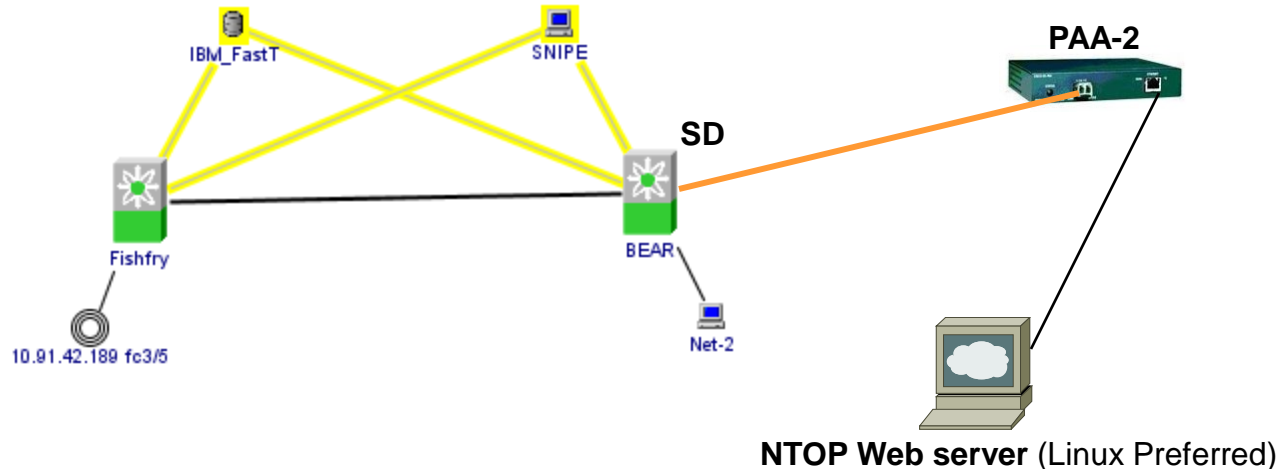
NTOP



NTOP

Network Traffic Probe <http://www.ntop.org/>

- Use PAA-2 to SPAN selected critical paths or devices.
- Set encapsulation on SD port to eisl to capture all VSAN traffic
- SPAN'ed Traffic can be dynamically edited and changed as needed based on issue at hand

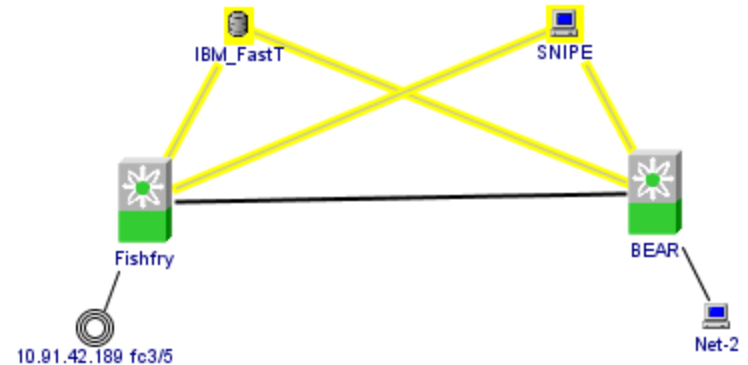
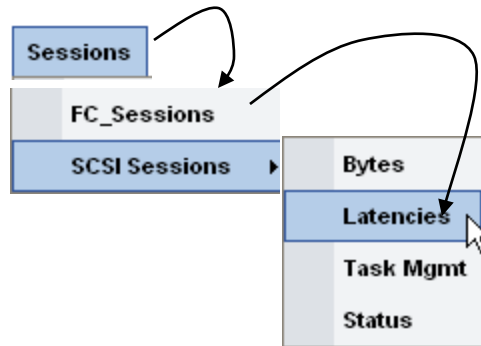


SPAN

Sessions		Source Interfaces	
Switch	Id	Interface	
BEAR	1	sup-fc0	Rx
BEAR	1	sup-fc0	Tx
BEAR	2	fc7/11	Rx
BEAR	2	fc7/11	Tx
BEAR	3	fc7/13	Rx
BEAR	3	fc7/13	Tx
BEAR	4	channel1	Rx

NTOP

Determining SCSI Latencies on the SAN



SCSI Sessions: Latencies

VSAN	Initiator	Target	LUN	Cmd-Status RTT		Cmd-XFR_RDY RTT		Cmd-Data RTT(Rd)		Cmd-Data RTT(Wr)	
				Min	Max	Min	Max	Min	Max	Min	Max
100	04.01.00	05.03.00	0	0.0 ms	225.6 ms	0.0 ms	9.3 ms	0.0 ms	224.7 ms	0.0 ms	9.3 ms


Members						
Zone	Type	Switch Interface	Name	WWN	FcId	
Snipe	WWN	Fishfry fc1/23	FastT-1	20:04:00:a0:b8:0c:64:51	0x050300	
Snipe	WWN	BEAR fc7/11	FastT-2	20:05:00:a0:b8:0c:64:51	0x040200	
Snipe	WWN	Fishfry fc1/2	Snipe2	21:00:00:e0:8b:05:a2:8f	0x050200	
Snipe	WWN	BEAR fc7/13	Snipe1	21:01:00:e0:8b:25:a2:8f	0x040100	

Performance of I/O Across Fabric


Info about 05.03.00

[[Main Page](#)] [[LUN Stats](#)] [[SCSI Session\(Bytes\)](#)] [[SCSI Session\(Times\)](#)] [[SCSI Session\(Status\)](#)] [[SCSI Session\(Task Mgmt\)](#)] [[FC Sessions](#)]

LUN Statistics

LUN 	Total Bytes		Data Bytes			Read Size		Write Size		Xfer Rdy Size		# Failed Cmds	Active Since	Last Seen
	Sent	Rcvd	Read	Write	Other	Min	Max	Min	Max	Min	Max			
0	239.1 MB	8.6 GB	223.2 MB	8.4 GB	2.7 KB	1	257	1	257	512	131584	0	04/27/07 11:48:51	04/30/07 15:21:51
1	7.7 MB	676.9 KB	0	0	8.2 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
2	7.2 MB	1.2 MB	0	0	8.2 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
3	7.8 MB	557.1 KB	0	0	8.1 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
4	7.2 MB	411.3 KB	0	0	7.3 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
5	7.3 MB	714.3 KB	0	0	7.7 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
6	7.7 MB	1.2 MB	0	0	8.6 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47
7	7.4 MB	1.1 MB	0	0	8.3 MB	0	0	0	0	512	131584	0	04/27/07 11:48:51	04/30/07 12:58:47

LUN 0 Traffic



LUN Traffic (Total Bytes)

Information in NTOP is referenced to FCID's. Many places in FM, DM and CLI to reference to what FCID's belong to what.

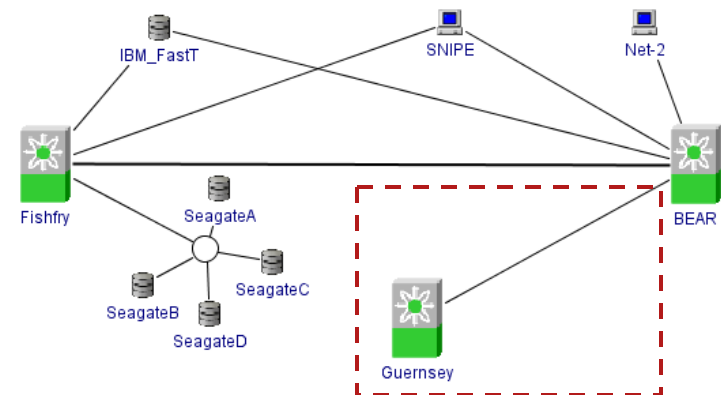
Troubleshooting with the Tools

- ISL's
- Zoning
- IVR
- NPV

Inter Switch Link

- VSAN 1 & 100 already configured up and running in core production SAN
- Add new 9509 switch that has VSAN 1 & 100 configured on it
- New switch using default Domain settings
- No known conflicts with zonesets

Picture of VSAN 100



Add new Switch and ISL to Fabric.

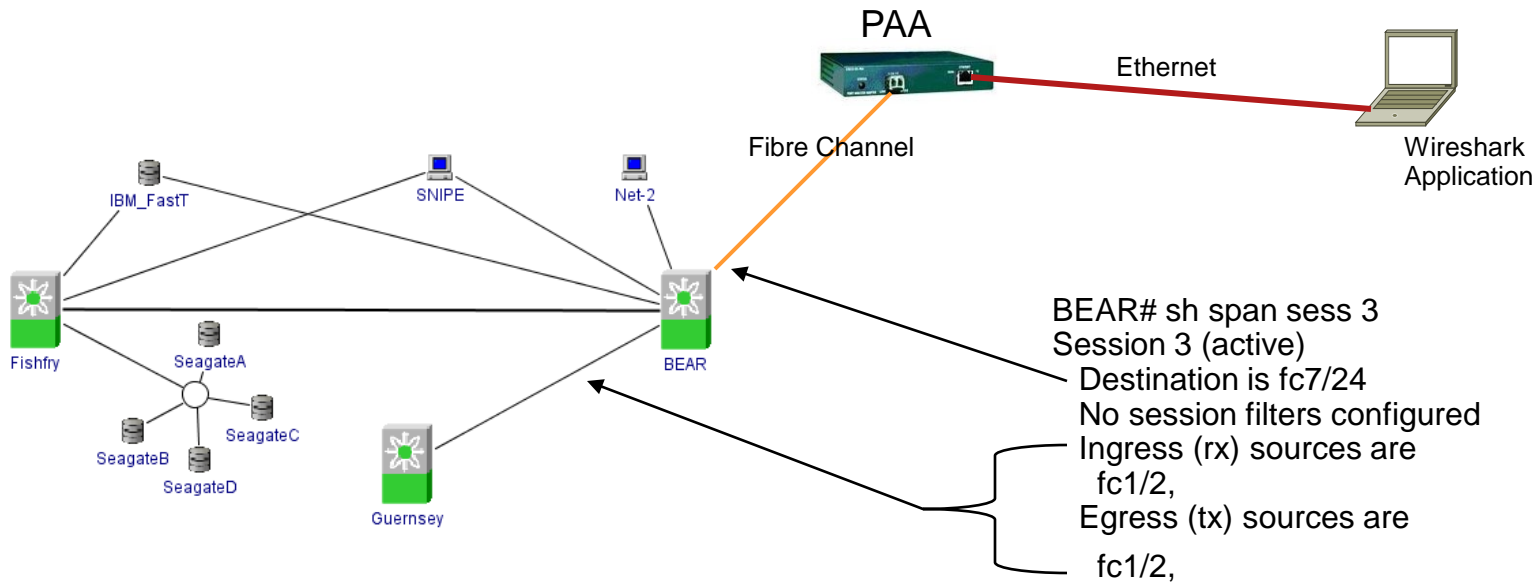
VSAN 100 Domain ID's

Switch	State	DomainId	Local WWN	Local Priority	Principal WWN	Principal Priority
Fishfry	stable	0x5(5)	20:64:00:0d:65:b4:2b:41	128	Cisco 20:64:00:0c:85:67:b1:c1	2
BEAR	stable	0x4(4)	20:64:00:0c:85:67:b1:c1	2	Cisco 20:64:00:0c:85:67:b1:c1	2

VSAN 1 Domain ID's

Switch	State	DomainId	Local WWN	Local Priority	Principal WWN	Principal Priority
Fishfry	stable	0x62(98)	20:01:00:0d:65:b4:2b:41	128	Cisco 20:01:00:0b:be:77:72:41	2
Sunny	stable	0x5(5)	20:01:00:0b:be:77:72:41	2	Cisco 20:01:00:0b:be:77:72:41	2
BEAR	stable	0xec(236)	20:01:00:0c:85:67:b1:c1	128	Cisco 20:01:00:0b:be:77:72:41	2

ISL Trace – TE Port SPAN to PAA



Switch	Id, Interface
Sunny	1, fc1/4 Rx
BEAR	1, sup-fc0 Rx
Sunny	1, fc1/4 Tx
BEAR	1, sup-fc0 Tx
BEAR	2, fc7/11 Rx
BEAR	2, fc7/11 Tx
BEAR	3, fc1/2 Rx
BEAR	3, fc1/2 Tx

SPAN view from Fabric Manager

Switch	Id	Dest Interface	VSAN Filter	Admin	Oper	Inactive Reason	VSAN Source List
BEAR	1	fc1/16		active	active	Session is active	
Sunny	1	fc-tunnel1		active	active	Session is active	
BEAR	2	fc1/16		active	active	Session is active	
BEAR	3	fc7/24		active	active	Session is active	

Id	Dest Interface	Filter VSAN List	Status Admin	Status Oper	Inactive Reason
1	fc1/16		active	active	Session is active
2	fc1/16		active	active	Session is active
3	fc7/24		active	active	Session is active

SPAN view from Device Manager

Session	VSAN List	Or Interface (Direction) List
1		sup-fc0 (Rx), sup-fc0 (Tx)
2		fc7/11 (Rx), fc7/11 (Tx)
3		fc1/2 (Rx), fc1/2 (Tx)

Filter: ((!(fc_r_ctl == 0xc0)) && !(swils_opcode == 0x14))

No. -	Time	Source	Destination	Protocol	Info
1	0.000000	ff.ff.fd	ff.ff.fd	SW_ILS	ELP
3	0.001578	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (ELP)
5	0.034187	ff.ff.fd	ff.ff.fd	SW_ILS	ESC
7	0.034397	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (ESC)
9	0.035247	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
10	0.035394	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
13	0.037132	ff.ff.fd	ff.ff.fd	SW_ILS	SW_RJT (0x71)
15	0.037290	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (0x71)
17	0.071893	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
19	0.076310	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (0x71)
21	0.085301	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
23	0.086075	ff.ff.fd	ff.ff.fd	SW_ILS	SW_RJT (0x71)
25	0.112896	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
27	0.114252	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (0x71)
29	0.117890	ff.ff.fd	ff.ff.fd	SW_ILS	0x71
31	0.147676	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (0x71)
33	0.152583	ff.ff.fd	ff.ff.fd	SW_ILS	ELP
35	0.152778	ff.ff.fd	ff.ff.fd	SW_ILS	ELP
37	0.153590	ff.ff.fd	ff.ff.fd	SW_ILS	ELP
38	0.153824	ff.ff.fd	ff.ff.fd	SW_ILS	ELP
41	0.154432	ff.ff.fd	ff.ff.fd	SW_ILS	SW_RJT (ELP)
42	0.155027	ff.ff.fd	ff.ff.fd	SW_ILS	SW_RJT (ELP)
43	0.155604	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (ELP)
44	0.156336	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (ELP)
49	0.162239	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
51	0.163433	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
53	0.163959	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
55	0.164516	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
56	0.164566	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
59	0.164988	ff.ff.fd	ff.ff.fd	SW_ILS	BF
61	0.165273	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
62	0.165498	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
65	0.166389	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
67	0.166614	ff.ff.fd	ff.ff.fd	SW_ILS	BF
69	0.167737	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (BF)
71	0.170235	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (BF)
73	5.164363	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
75	5.164797	ff.ff.fd	ff.ff.fd	SW_ILS	SW_RJT (EFF)
77	5.165237	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
79	5.165464	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
80	5.165640	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
83	5.165936	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
85	5.166513	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
86	5.166720	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
88	5.166866	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
89	5.166957	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
93	6.164005	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
95	6.164442	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
97	15.173808	ff.ff.fd	ff.ff.fd	SW_ILS	DIA
99	15.174263	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (DIA)
101	15.174771	ff.ff.fd	ff.ff.fd	SW_ILS	RDI
103	15.181814	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (RDI)
105	15.182194	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
107	15.187714	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
109	15.191119	ff.ff.fd	ff.ff.fd	SW_ILS	DIA
111	15.192303	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (DIA)
113	15.192955	ff.ff.fd	ff.ff.fd	SW_ILS	RDI
115	15.195298	ff.ff.fd	ff.ff.fd	SW_ILS	MRRA
117	15.195757	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MRRA)
119	15.196354	ff.ff.fd	ff.ff.fd	SW_ILS	MR[Malformed Packet]
121	15.197555	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MR)
123	15.203309	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
125	15.205308	ff.ff.fd	ff.ff.fd	SW_ILS	MRRA
127	15.208392	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)
129	15.208911	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MRRA)
131	15.209637	ff.ff.fd	ff.ff.fd	SW_ILS	MR[Malformed Packet]
133	15.210572	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (RDI)
135	15.211227	ff.ff.fd	ff.ff.fd	SW_ILS	EFF
139	15.212859	ff.ff.fd	ff.ff.fd	SW_ILS	MRRA
141	15.214333	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MR)
143	15.226231	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (EFF)

ISL Trace

ACK1 Filter Applied

Exchange Link Parameters

Exchange Switch Capabilities

Exchange Peer Parameters
(This Is Proprietary to MDS Only to Determine if Connecting Switch Is Another MDS)

Exchange Fabric Parameters

Build Fabric

Domain ID Assign by Existing Principal Switch

Request Domain ID from New Switch

Enhanced Zoning Merge Request Resource Allocation (New in FC-SW3 Standard)

Zone Merge Request

Break down of how the Domains were selected and distributed on next slide

Domain ID Assignments and Distribution

VSAN 1

```
113 15.192955 ff.ff.fd ff.ff.fd SW_ILS RDI
  Frame 113 (64 bytes on wire, 64 bytes captured)
  Ethernet II, Src: Andiamo_00:a2:9d (00:05:30:00:a2:9d)
  Boardwalk
  Fibre Channel
  SW_ILS
    Cmd Code: RDI (0x13)
    Payload Len: 16
    Req Switch Name: 20:01:00:0d:ec:01:40:81 (Cisco)
    Requested Domain ID: 104
```

Request for ID 104 Decimal

```
135 15.211227 ff.ff.fd ff.ff.fd SW_ILS EFP
  Frame 135 (128 bytes on wire, 128 bytes captured)
  Ethernet II, Src: Andiamo_00:a2:9d (00:05:30:00:a2:9d),
  Boardwalk
  Fibre Channel
  SW_ILS
    Cmd Code: EFP (0x11)
    Record Len: 16
    Payload Len: 80
    Principal Switch Priority: 2
    Principal Switch Name: 20:01:00:0b:be:77:72:41 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x05
    Switch Name: 20:01:00:0b:be:77:72:41 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x62
    Switch Name: 20:01:00:0d:65:b4:2b:41 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0xec
    Switch Name: 20:01:00:0c:85:67:b1:c1 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x68
    Switch Name: 20:01:00:0d:ec:01:40:81 (Cisco)
```

EFP distributes to new switch the list of where other domains are found

VSAN 100

```
101 15.174771 ff.ff.fd ff.ff.fd SW_ILS RDI
  Frame 101 (64 bytes on wire, 64 bytes captured)
  Ethernet II, Src: Andiamo_00:a2:9d (00:05:30:00:a2:9d)
  Boardwalk
  Fibre Channel
  SW_ILS
    Cmd Code: RDI (0x13)
    Payload Len: 16
    Req Switch Name: 20:64:00:0d:ec:01:40:81 (Cisco)
    Requested Domain ID: 113
```

Request for ID 113 Decimal

```
123 15.203309 ff.ff.fd ff.ff.fd SW_ILS EFP
  Frame 123 (112 bytes on wire, 112 bytes captured)
  Ethernet II, Src: Andiamo_00:a2:9d (00:05:30:00:a2:9d),
  Boardwalk
  Fibre Channel
  SW_ILS
    Cmd Code: EFP (0x11)
    Record Len: 16
    Payload Len: 64
    Principal Switch Priority: 2
    Principal Switch Name: 20:64:00:0c:85:67:b1:c1 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x04
    Switch Name: 20:64:00:0c:85:67:b1:c1 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x71
    Switch Name: 20:64:00:0d:ec:01:40:81 (Cisco)
  Domain ID Record
    Record Type: Domain ID Record (0x01)
    Domain ID: 0x05
    Switch Name: 20:64:00:0d:65:b4:2b:41 (Cisco)
```

Debugging Zoning

- Understanding what is Active
- Fabric Manager / CLI
- Basic/Enhanced
- Zoneset activate failure?
- Zoneset merge failure?
- Members not able to communicate?
- Host can not see storage?

Show Zoneset Active

- Shows the zoneset activated in the fabric
- An **asterisk** to left of device indicates that device has registered with the nameserver
- Will not show if a zone has been configured but changes have not been activated
- Zoneset and zone names are case sensitive
- Inter-VSAN routing zone added to ZoneSet via IVR wizard if this VSAN has member being zoned to another VSAN
- Notice different zonesets per VSAN**

```
BEAR# show zoneset active
zoneset name VSAN3 vsan 3
zone name Wolf vsan 3
* fcid 0x030100 [pwwn 21:00:00:e0:8b:04:2f:35]
* fcid 0x030201 [pwwn 50:01:04:f0:00:47:34:e6]
* fcid 0x800100 [pwwn 21:00:00:e0:8b:04:69:3c]

zoneset name VS100_Zoneset vsan 100
zone name Net-2 vsan 100
* fcid 0x0501e1 [pwwn 22:00:00:20:37:c5:36:f0]
* fcid 0x040001 [pwwn 10:00:00:00:c9:2f:99:3d] [DA-Net-2]

zone name Snipe vsan 100
* fcid 0x040100 [pwwn 21:01:00:e0:8b:25:a2:8f]
* fcid 0x050200 [pwwn 21:00:00:e0:8b:05:a2:8f] [Snipe]
* fcid 0x050300 [pwwn 20:04:00:a0:b8:0c:64:51]
* fcid 0x040200 [pwwn 20:05:00:a0:b8:0c:64:51]

zoneset name ZS200_Zoneset vsan 200
zone name Eagle vsan 200
* fcid 0xd10001 [pwwn 10:00:00:00:c9:24:d2:d3] [Eagle]
* fcid 0x8001e1 [pwwn 22:00:00:20:37:e6:8e:58]

zone name Curley vsan 200
* fcid 0x800300 [pwwn 21:00:00:e0:8b:03:98:16] [Curly]
* fcid 0x8001dc [pwwn 22:00:00:20:37:e6:8d:93]

zone name IVRZ_IvrZone1 vsan 200
* fcid 0x8001e1 [pwwn 22:00:00:20:37:e6:8e:58]
* fcid 0x0b0001 [pwwn 10:00:00:00:c9:2b:4f:81]
```


Host Can't See Storage?

- View 'show zoneset active vsan x' on each switch.
- Look for * next to the affected device. The * indicates the device is in the zone, and is active in the name server.
- If the * is present, in each switch, then the zoning displays look as good as they get. Verify the zone members are correct for the devices that are affected. Correct PWWN and FCID displayed?

```
BEAR# show zoneset active vsan 100
zoneset name VS100_Zoneset vsan 100
zone name Net-2 vsan 100
* fcid 0x0501e1 [pwwn 22:00:00:20:37:c5:36:f0]
* fcid 0x040001 [pwwn 10:00:00:00:c9:2f:99:3d] [DA-Net-2]
```

Device is not active in the fabric
but is in the zone

```
zone name Snipe vsan 100
* fcid 0x040100 [pwwn 21:01:00:e0:8b:25:a2:8f]
* fcid 0x050200 [pwwn 21:00:00:e0:8b:05:a2:8f] [Snipe]
* fcid 0x050300 [pwwn 20:04:00:a0:b8:0e:64:51]
  pwwn 20:05:00:a0:b8:0c:64:51
```

Device is active in the fabric
and in the zone

```
zone name Net-2-extrastor vsan 100
* fcid 0x0501e2 [pwwn 22:00:00:20:37:c5:22:de]
* fcid 0x040001 [pwwn 10:00:00:00:c9:2f:99:3d] [DA-Net-2]
```

Zoneset Activate Failure

- Look at the messages on the seed switch to determine what caused the activate to fail. **‘show logging log’** (seed is the switch where the change was initiated from)
- For multi switch fabrics, check that the ISL or TE-ISL is operational.
- Show interface fcx/x or show interface port-channel x
- Show interface trunk vsan x
- Show fcdomain domain-list vsan x
- Show zone internal

Change Event History

```
BEAR# show zone internal change event-history vsan 100
Change Protocol Event Log For VSAN: 100
>>>>FSM has 50 logged transitions<<<<
```

```
46) Transition at Tue May 1 10:37:34 2007
   Prev State: [ACA Sent]
   Trig event: [RCVD_ACC] (Dom:5)
   Next State: [ACA Sent]
47) Transition at Tue May 1 10:37:34 2007
   Prev State: [ACA Sent]
   Trig event: [RCVD_ALL_ACC]
   Next State: [ACA Complete]
48) Transition at Tue May 1 10:37:34 2007
   Prev State: [ACA Complete]
   Trig event: [SEND_RCA]
   Next State: [RCA Sent]
49) Transition at Tue May 1 10:37:34 2007
   Prev State: [RCA Sent]
   Trig event: [RCVD_ACC] (Dom:5)
   Next State: [RCA Sent]
50) Transition at Tue May 1 10:37:34 2007
   Prev State: [RCA Sent]
   Trig event: [RCVD_ALL_ACC]
   Next State: [Idle]
   Curr state: [Idle]
```

Command to view
change history for
VSAN 100

One event.

Time of event is noted
previous state, event, and
next state are shown.

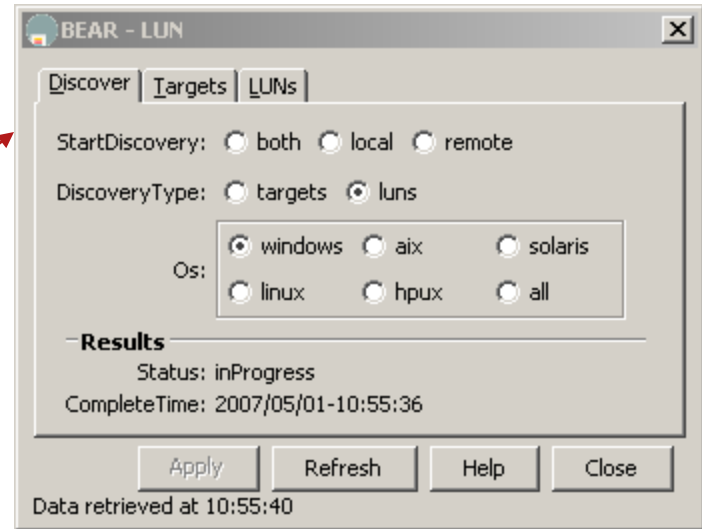
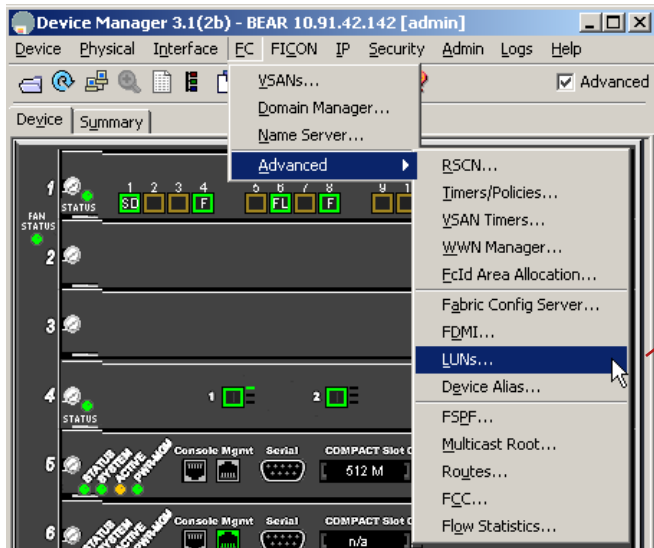
Domain that initiated the
event is also seen

Zone Merge Failure Options

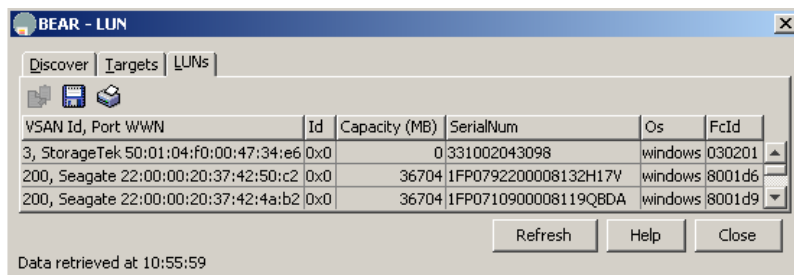
- Confirm that there is a discrepancy in the zonesets on opposite sides of the ISL or E-ISL that is isolated.
- Determine which of the 2 fabrics contain the desired active zoneset, then use the zoneset import/export command. This command only works if the ISL/E-ISL is isolated.
- Prune the VSAN from the TE port, and add it back.
- Edit one or the other zonesets and then shut/no shut the ISL or E-ISL. **This action will impact all VSANs on the E-ISL even those that are not isolated.**

See if the Switch Can See the Storage But Still Not the Host

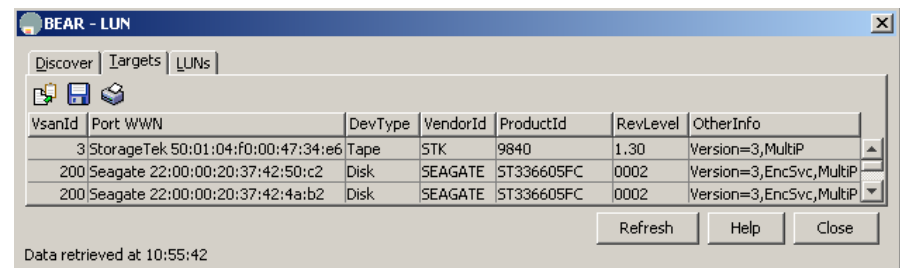
MDS has a pseudo initiator device that can log into targets (if target accepts PLOGI from the switch PWWN). Seen here with the Fabric Manager



LUNs



Targets



Zoneset Activation - Trace

Activate zone in VSAN 100, using VSAN filter option on the SPAN session

BEAR# show span session 3
 Session 3 (active)
 Destination is fc7/24
 Session filter vsans are 100
 Ingress (rx) sources are port-channel 1,
 Egress (tx) sources are port-channel 1,

Success

No. -	Time	Source	Destination	Protocol	Info
1	0.000000	ff.ff.fd	ff.ff.fd	SW_ILS	HLO
2	0.000084	ff.ff.fd	ff.ff.fd	SW_ILS	HLO
3	0.000153	ff.ff.fd	ff.ff.fd	SW_ILS	HLO
4	0.000233	ff.ff.fd	ff.ff.fd	FC	Link Ctl, ACK1
5	0.000721	ff.ff.fd	ff.ff.fd	SW_ILS	HLO
6	0.000937	ff.ff.fd	ff.ff.fd	FC	Link Ctl, ACK1
7	0.001063	ff.ff.fd	ff.ff.fd	FC	Link Ctl, ACK1
8	0.241821	ff.ff.fd	ff.ff.fd	FC	Link Ctl, ACK1
9	3.598559	ff.fc.04	ff.fc.05	SW_ILS	ACA
10	3.599083	ff.fc.05	ff.fc.04	FC	Link Ctl, ACK1
11	3.599533	ff.fc.05	ff.fc.04	SW_ILS	SW_ACC (ACA)
12	3.599689	ff.fc.04	ff.fc.05	FC	Link Ctl, ACK1
13	3.603425	ff.fc.04	ff.fc.05	SW_ILS	SFC
14	3.603834	ff.fc.05	ff.fc.04	FC	Link Ctl, ACK1
15	3.609451	ff.fc.05	ff.fc.04	SW_ILS	SW_ACC (SFC)
16	3.609628	ff.fc.04	ff.fc.05	FC	Link Ctl, ACK1
17	3.610906	ff.fc.04	ff.fc.05	SW_ILS	UFC
18	3.611297	ff.fc.05	ff.fc.04	FC	Link Ctl, ACK1
19	3.747253	ff.fc.05	ff.fc.04	SW_ILS	SW_ACC (UFC)
20	3.747426	ff.fc.04	ff.fc.05	FC	Link Ctl, ACK1
21	3.747719	ff.fc.04	ff.fc.05	SW_ILS	RCA
22	3.748071	ff.fc.05	ff.fc.04	FC	Link Ctl, ACK1
23	3.750139	04.01.00	05.03.00	FC ELS	PLOGI
24	3.750891	ff.fc.05	ff.fc.04	SW_ILS	SW_ACC (RCA)
25	3.751077	ff.fc.04	ff.fc.05	FC	Link Ctl, ACK1
26	3.756635	05.02.00	04.02.00	FC ELS	PLOGI
27	3.758050	05.03.00	04.01.00	FC ELS	ACC (PLOGI)
28	3.758070	04.01.00	05.03.00	FC ELS	PRLI
29	3.758132	05.03.00	04.01.00	FC ELS	ACC (PRLI)
30	3.762430	04.02.00	05.02.00	FC ELS	ACC (PLOGI)
31	3.762470	05.02.00	04.02.00	FC ELS	PRLI
32	3.762529	04.02.00	05.02.00	FC ELS	ACC (PRLI)
33	4.409947	04.00.01	05.01.e1	FC ELS	PDISC
34	4.410487	05.01.e1	04.00.01	FC ELS	ACC (PDISC)

Recovering from Zone Issues

Zoneset Import/Export Command

MDS_Switch# **show int fc1/1**

FC1/1 is trunking (Not all VSANs UP on the trunk)

Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)

Port WWN is 20:c5:00:05:30:00:49:1e

Peer port WWN is 20:81:00:0d:ec:0f:b4:c0

Admin port mode is E, trunk mode is on

snmp traps are enabled

Port mode is TE

Port vsan is 1

Speed is 2 Gbps

Transmit B2B Credit is 255

Receive B2B Credit is 12

Receive data field Size is 2112

Beacon is turned off

Trunk vsans (admin allowed and active) (1,10)

Trunk vsans (up) (1)

Trunk vsans (isolated) (10)

VSAN 10 is isolated

Show Zoneset that is deemed undesirable

MDS_Switch# **show zoneset act v 10**

zoneset name z10 vsan 10

zone name duplicate vsan 10

pwwn 10:10:10:10:10:10:10:10

Command to bypass merge checking and force import of zone data

MDS_Switch# **zoneset import interface fc1/1 vsan 10**

<show port internal info> will give greater detail on merger failure reason

Recovering from Zone Issues

Zoneset Import/Export Command - After the Import

```
MDS_Switch# show int fc1/1
```

```
Fc1/1 is trunking
```

```
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
```

```
Port WWN is 20:c5:00:05:30:00:49:1e
```

```
Peer port WWN is 20:81:00:0d:ec:0f:b4:c0
```

```
Admin port mode is E, trunk mode is on
```

```
snmp traps are enabled
```

```
Port mode is TE
```

```
Port vsan is 1
```

```
Speed is 2 Gbps
```

```
Transmit B2B Credit is 255
```

```
Receive B2B Credit is 12
```

```
Receive data field Size is 2112
```

```
Beacon is turned off
```

```
Trunk vsans (admin allowed and active) (1,10)
```

```
Trunk vsans (up) (1,10)
```

```
Trunk vsans (isolated) ( )
```

```
Trunk vsans (initializing) ( )
```

VSAN 10 is no longer isolated

VSAN 10 zoneset has been imported

Entry in the accounting log reflecting the import action

```
MDS_Switch# show zoneset act vsan 10
```

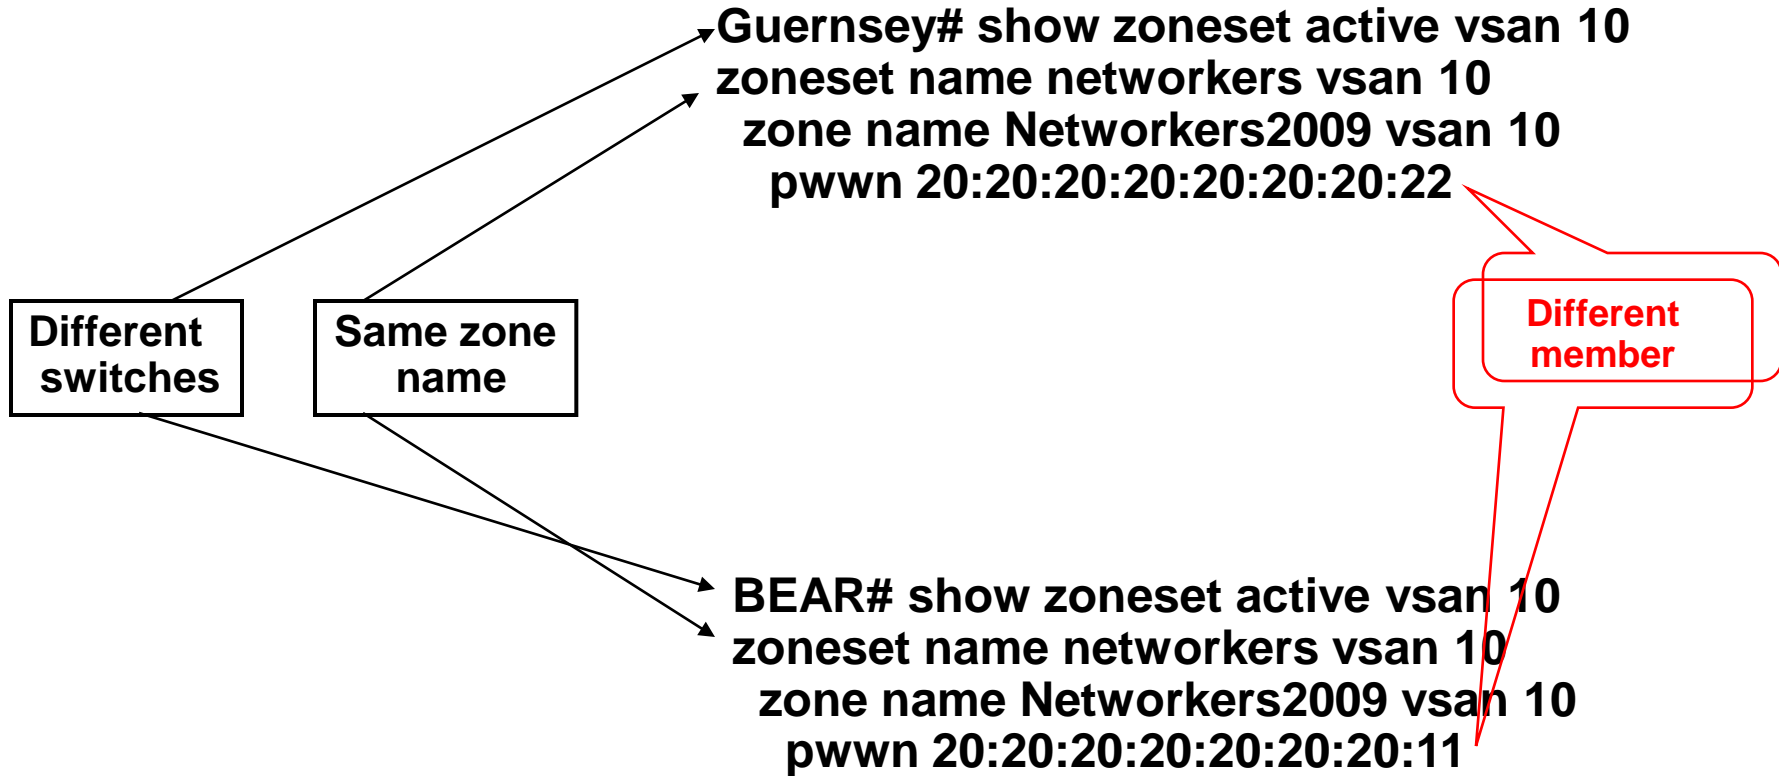
```
zoneset name v10 vsan 10
```

```
zone name duplicate vsan 10
```

```
pwwn 21:00:00:20:37:a9:cd:6e
```

```
admin:Import option is set on interface fc1/1 on VSAN 10
```

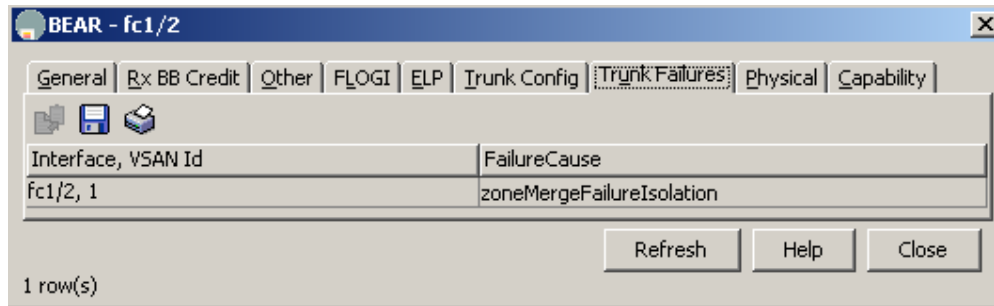

Zone Merge Failure Observations



BEAR %ZONE-2-ZS_MERGE_FAILED: %\$VSAN 10%\$ Zone merge failure, isolating interface fc1/2 error:
Member mismatch

Guernsey %ZONE-2-ZS_MERGE_FAILED: %\$VSAN 10%\$ Zone merge failure, isolating interface fc2/1 error:
Received rjt from adjacent switch

Zone Merge Failure as Seen on the ISL Interface



```

BEAR# show interface fc1/2
fc1/2 is trunking (Not all VSANs UP on the trunk)
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:02:00:0c:85:67:b1:c0
  Peer port WWN is 20:41:00:0d:ec:01:40:80
  Admin port mode is E, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
  Transmit B2B Credit is 12
  Receive B2B Credit is 255
  Receive data field Size is 2112
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1,100)
  Trunk vsans (up) (100)
  Trunk vsans (isolated) (1)
  Trunk vsans (initializing) ( )
  
```

VSANs with no merge conflicts will show here.

VSAN 1 is isolated in this example due to zoneMerge failure

Show port internal info interface fc1/2

Trace of Zone Failure

BEAR# show span session 3
Session 3 (active)
Destination is fc7/24
Session filter vsans are 1
Ingress (rx) sources are
fc1/2,
Egress (tx) sources are
fc1/2,

Filter: `!(fc_r_ctl == 0xc0)` Expression... Clear Apply

No. -	Time	Source	Destination	Protocol	Info
77	15.144990	tt.tt.td	tt.tt.td	SW_ILS	SW_ACC (MRRA)
78	15.145235	ff.ff.fd	ff.ff.fd	FC	Basic Link Svc
80	15.145422	tt.tt.td	tt.tt.td	SW_ILS	MR
83	15.145610	ff.ff.fd	ff.ff.fd	SW_ILS	MRRA
84	15.145718	ff.ff.fd	ff.ff.fd	BLS	BA_ACC
87	15.146023	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MRRA)
89	15.146776	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MR)
91	25.141506	ff.ff.fd	ff.ff.fd	FC ELS	RRQ
93	25.141891	ff.ff.fd	ff.ff.fd	FC ELS	ACC (RRQ)

SW_ILS
 Cmd Code: MR (0x22)
 Active Zoneset Length: 60
 Active Zoneset Name: netw
 Active Zone Set
 Number of zoning object
 Zone object 0
 Zone object Type: Zon
 Zone Protocol: 0x00
 Zone Object Name: Net
 Number of Zone Member
 Zone Member 0
 Zone Member Type: wwn
 Flags: 0x0
 Identifier Length: 8
 Member Identifier: 20
 Full Zone Set Length: 0

No. -	Time	Source	Destination	Protocol	Info
77	15.144990	tt.tt.td	tt.tt.td	SW_ILS	SW_ACC (MRRA)
78	15.145235	ff.ff.fd	ff.ff.fd	FC	Basic Link Svc
80	15.145422	ff.ff.fd	ff.ff.fd	SW_ILS	MR
83	15.145610	ff.ff.fd	ff.ff.fd	SW_ILS	MRRA
84	15.145718	ff.ff.fd	ff.ff.fd	BLS	BA_ACC
87	15.146023	ff.ff.fd	ff.ff.fd	SW_ILS	SW_ACC (MRRA)
89	15.146776	tt.tt.td	tt.tt.td	SW_ILS	SW_ACC (MR)
91	25.141506	ff.ff.fd	ff.ff.fd	FC ELS	RRQ
93	25.141891	ff.ff.fd	ff.ff.fd	FC ELS	ACC (RRQ)

OX_ID: 0x97d2
 RX_ID: 0x0ec0
 Parameter: 0x00000000
 SW_ILS
 Cmd Code: SW_ACC (0x02)
 Zone Command Status: Failed (0x02)
 Zone Command Reason Code: Cannot Merge (0x0a)
 vendor unique: 0x0

MR (Merge Request) Frame From

< FCAnalyzer detail >

```

Fibre Channel
Exchange Last In: 67
R_CTL: 0x2(Device_Data/Unsolicited Control)
Dest Adct: ff,ff,fd
CS_CTL: 0x00
Src Adct: ff,ff,fd
Type: SW_ILS (0x22)
F_CTL: 0x290000 Exchange Originator, Seq Initiator, Exchg First, Seq Last, CS_CTL, Transfer Seq Initiative
SEQ_ID: 0x0d
DF_CTL: 0x00
SEQ_CNT: 0
OX_ID: 0x025f
RX_ID: 0xffff
Parameter: 0x00000000
SW_ILS
Cmd Code: MR (0x22)
Active ZoneSet Length: 0
Full Zone Set Length: 56
Full Zone Set
  Number of zoning objects: 212
  Zone Object 0
    Zone Object Type: Reserved (0x00)
    Zone Protocol: 0x00
    Zone Object Name: \340
    Number of Zone Members: 16
    Zone Member 0
      Zone Member Type: Reserved (0x00)
      Flags: 0x0
      Identifier Length: 86
      Member Identifier: Unknown member type format
    Zone Member 1
      Zone Member Type: Reserved (0x00)
      Flags: 0x0
      Identifier Length: 0
      Member Identifier: Unknown member type format
    Zone Member 2
      Zone Member Type: Zone Alias (0x04)
      Flags: 0x0
      Identifier Length: 0
      Member Identifier: \340\213%\242\217\001
  .....SNIP.....
  
```

**FCID's are the
Fabric
Controller**

**MR
command
code**

... PASTE.....

```

Zone Member 15
  Zone Member Type: Unknown (0x69)
  Flags: 0x65
  Identifier Length: 50
  Member Identifier: Unknown member type format
Zone Object 1
  Zone Object Type: Zone Alias (0x03)
  Zone Protocol: 0x00
  Zone Object Name: FastT-2
  Number of Zone Members: 1
  Zone Member 0
    Zone Member Type: WWN (0x01)
    Flags: 0x0
    Identifier Length: 8
    Member Identifier: 20:05:00:a0:b8:0c:64:51 (SymbiosL)
Zone Object 2
  Zone Object Type: Unknown (0x05)
  Zone Protocol: 0x00
  Zone Object Name: $default_zone_attr_group$
  Number of Zone Members: 0
  
```

Trace taken from
config prompt on
switch

Capture to file if
need be

Can use Wireshark
to also view and
analyze the captured
file

Enhanced Zone Analysis

BEAR# show zone analysis ?

active Show active zoneset analysis *analyze the active zoneset on a vsan*
vsan Show analysis in the specified VSAN *analyze the full database for a vsan*
zoneset Show zoneset analysis *analyze a specific zoneset for a vsan*

BEAR# show zone analysis zoneset

VS100_Zoneset vsan 100

Zoning database analysis vsan 100

Zoneset analysis: VS100_Zoneset

Num zonesets: 1

Num zones: 3

Num aliases: 6

Num attribute groups: 1

Formattted size: 580 bytes / 2048 Kb

SFC will be 580 bytes of the max 2048Kb



Troubleshooting Zone/ACL Issues

What to collect

On the Supervisor:

```
switch# show tech-support details
switch# show tech-support acltcam-soc
switch# show tech-support zbm // This is Zone Block Manager
switch# show tech-support zone
```

On the linecard (attach module 1):

```
module-1# show hardware internal packet-flow dropped
    •check the port-stats for port P & acl-stats on the port-grp N where you see drops

module-1# show hardware internal fwd port <p> port-stats
module-1# show hardware internal fwd 0 port-group <n> acl-stats
module-1# show hardware internal errors
module-1# show hardware internal fwd 0 error-statistics
```

Best Practices to Head Off Having to Troubleshoot

- Make periodic backups of zoning database
- Prior to any changes, make a backup of the current zoning
- Single initiator zones
- Meaningful zone names
- Default-zone set to deny
- Distribute full zoneset
- Use alias, device-alias preferred
- Manage from designated seed switch or switches

IVR

Inter VSAN Routing



Reading IVR Where, What

Host on 9216i switch needs to get at Disk on different VSAN on different switch

IVR only enabled where needed - only 9216i

Control	CFS	Local Topology	Fabric Id	Default Fabric Id	Action	RDI VSANs	Active Topology
Switch	Fabric Id, VSAN Route Switch: WWN, Name			VSAN List			
Sunny	1, 20:00:00:0b:be:77:72:40, Sunny			10,200			

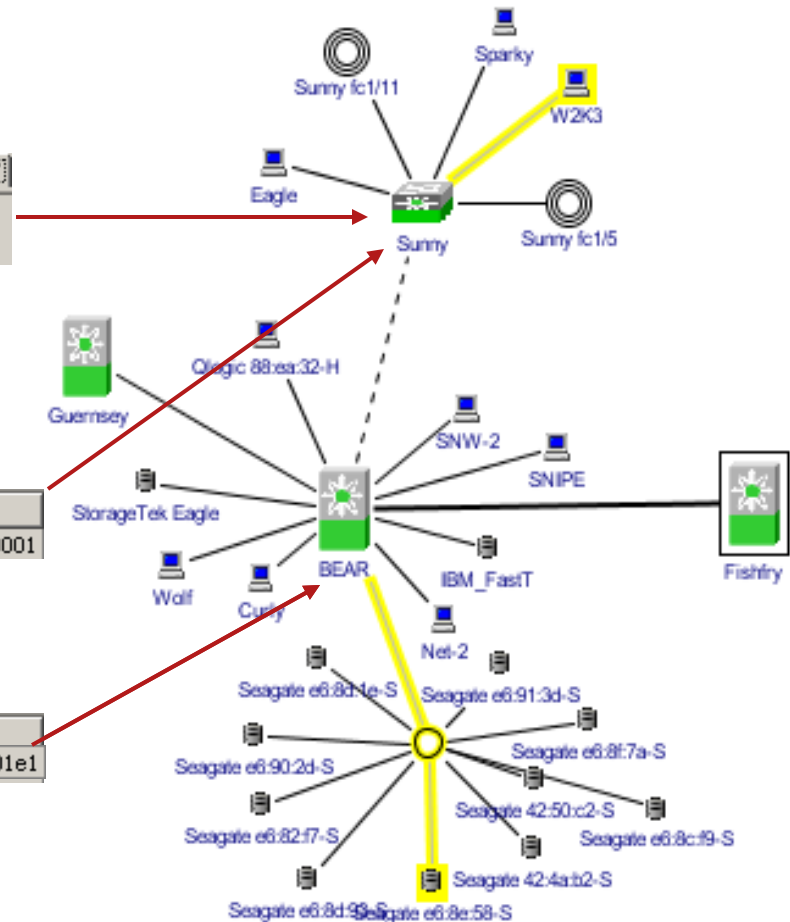
VSAN10

VSAN Id	Zone	Type	Fabric Id	Switch Interface	Name	WWN	FcId
10	IvrZone1	WWN	1	Sunny fc1/3	w2k_serverA	10:00:00:00:c9:2b:4f:81	0x0b0001

No NAT required because Domain ID's are unique

VSAN 200

VSAN Id	Zone	Type	Fabric Id	Switch Interface	Name	WWN	FcId
200	IvrZone1	WWN	1	BEAR fc1/6	Eagle-Storage	22:00:00:20:37:e6:8e:58	0x8001e1



CLI for IVR – Show Configuration

Showing IVR-VSAN Topology Configuration (should match in both switches)

```
switch# show ivr vsan-topology
```

AFID	SWITCH WWN	Active	Cfg.	VSANS
1	20:00:00:05:30:00:3c:5e	yes	yes	1,4
1	20:00:00:05:30:00:58:de	yes	yes	2,4

Showing IVR-Zone Configuration (should match in both switches)

```
switch# show ivr zoneset active
```

```
zoneset name ZoneSet1
```

```
zone name Zone_VSAN1-2
```

```
* pwwn 21:00:00:e0:8b:02:ca:4a vsan 1
```

```
* pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

IVR Initiated Distribution of Information Via CFS (Cisco Fabric Services)

- Information distribution initiated by IVR process, not by user
- Events that alter the IVR topology database
- VSAN creation
- Link Shutdown
- Must do IVR COMMIT to initiate distribution for config changes

Show IVR, What Can I look at in 3.0?

Sunny# show ivr ?

fcdomain	Display IVR persistent fcdomain database
internal	Show ivr internal information
merge	Show ivr merge status
pending	Show ivr pending configuration
pending-diff	Show ivr pending-diff
service-group	Show IVR service groups
session	Show ivr session status
tech-support	Show information for IVR technical support staff
virtual-domains	Show IVR virtual domains for all local VSANs
virtual-fcdomain-add-status	Show IVR-virtual fcdomain status
virtual-switch-wwn	Show IVR-virtual switch WWN
vsan-topology	Show IVR VSAN topology
zone	Inter vsan zone show commands
zoneset	Inter vsan routing zoneset show commands
	Output modifiers.
>	Output Redirection.
<cr>	Carriage return.

Best Practices for IVR

- **Management**

- Configure IVR only in needed border switches

- Configure/manage from Fabric Manager

- Do not use IVR topology auto discovery in production (pre 3.0 SAN-OS, it adds every VSAN to IVR)

- Use transit VSANs for FCIP links or FCIP port channels

- **Domain ID's & VSAN's**

- Plan out your VSAN numbers and domain IDs

- Use static domain IDs

- Use RDI mode to reserve domain IDs

- **Zoning**

- Keep default zone policy at deny

- Manage local zones from IVR enabled switch

- **Keep all IVR enabled switches at the same SAN OS version**

- **Do not mix IVR NAT with FCIP write acceleration**

IVR Troubleshooting - Reference

Database checks

Are devices logged into their local VSAN? (*show flogi database*)

Are devices exported into remote FCNS in both directions? (*show fcns database*)

Are FC Devices in the native VSAN FCNS in all switches in that VSAN? (*show fcns database*)

Do FC devices show correctly in the Transit VSAN, if one is in use?

Does the device have a valid PWWN and NWWN? *IVR checks before exporting.*

Zone checks

Use command line to view active local zoneset (*show zoneset active*)

Does the same IVR zone show up in both local and remote VSAN's active zoneset?

Did IVR zoneset activation succeed in all VSANs for the affected devices?

Look for the * next to all IVR'd devices in both VSANs local active zoneset

Miscellaneous Checks

Is it possible that a natted FCID changed because of a reload causing AIX or HP-UX to have target binding issues?

Ensure HBA is not configured to time out PLOGI to quickly. IVR NAT delays ACC to PLOGI for a few seconds. Most HBAs have a 10 second timeout.

Is the IVR VSAN topology exactly the same in every IVR enabled switch in the fabric? *<show IVR VSAN topology>*

Upgrading & What to Collect for Support

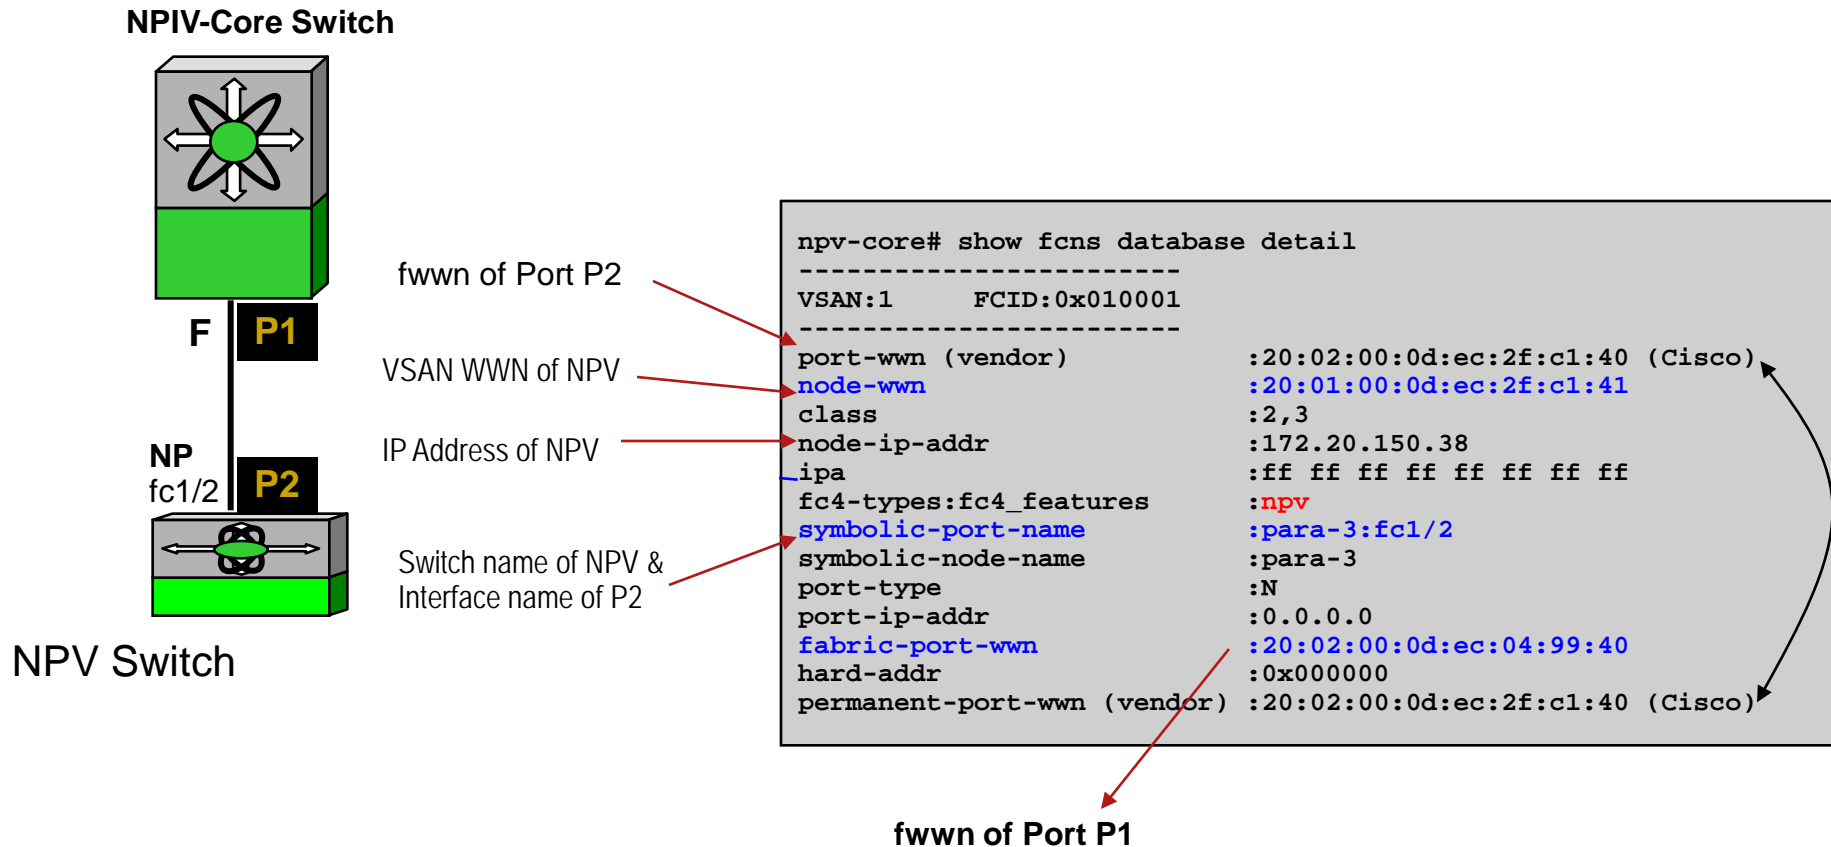
- Upgrades need to be planned and well thought out
- The mix of IVR1 (Non NAT) and IVR2 (NAT) can be tricky and confusing to configure without introducing traffic disruption
- Downgrades might require entire fabric to be downgraded
- show tech-support IVR (on each IVR enabled switch)
- show tech-support details (on the affected switches)
- FM screen snap shots

NPV Troubleshooting



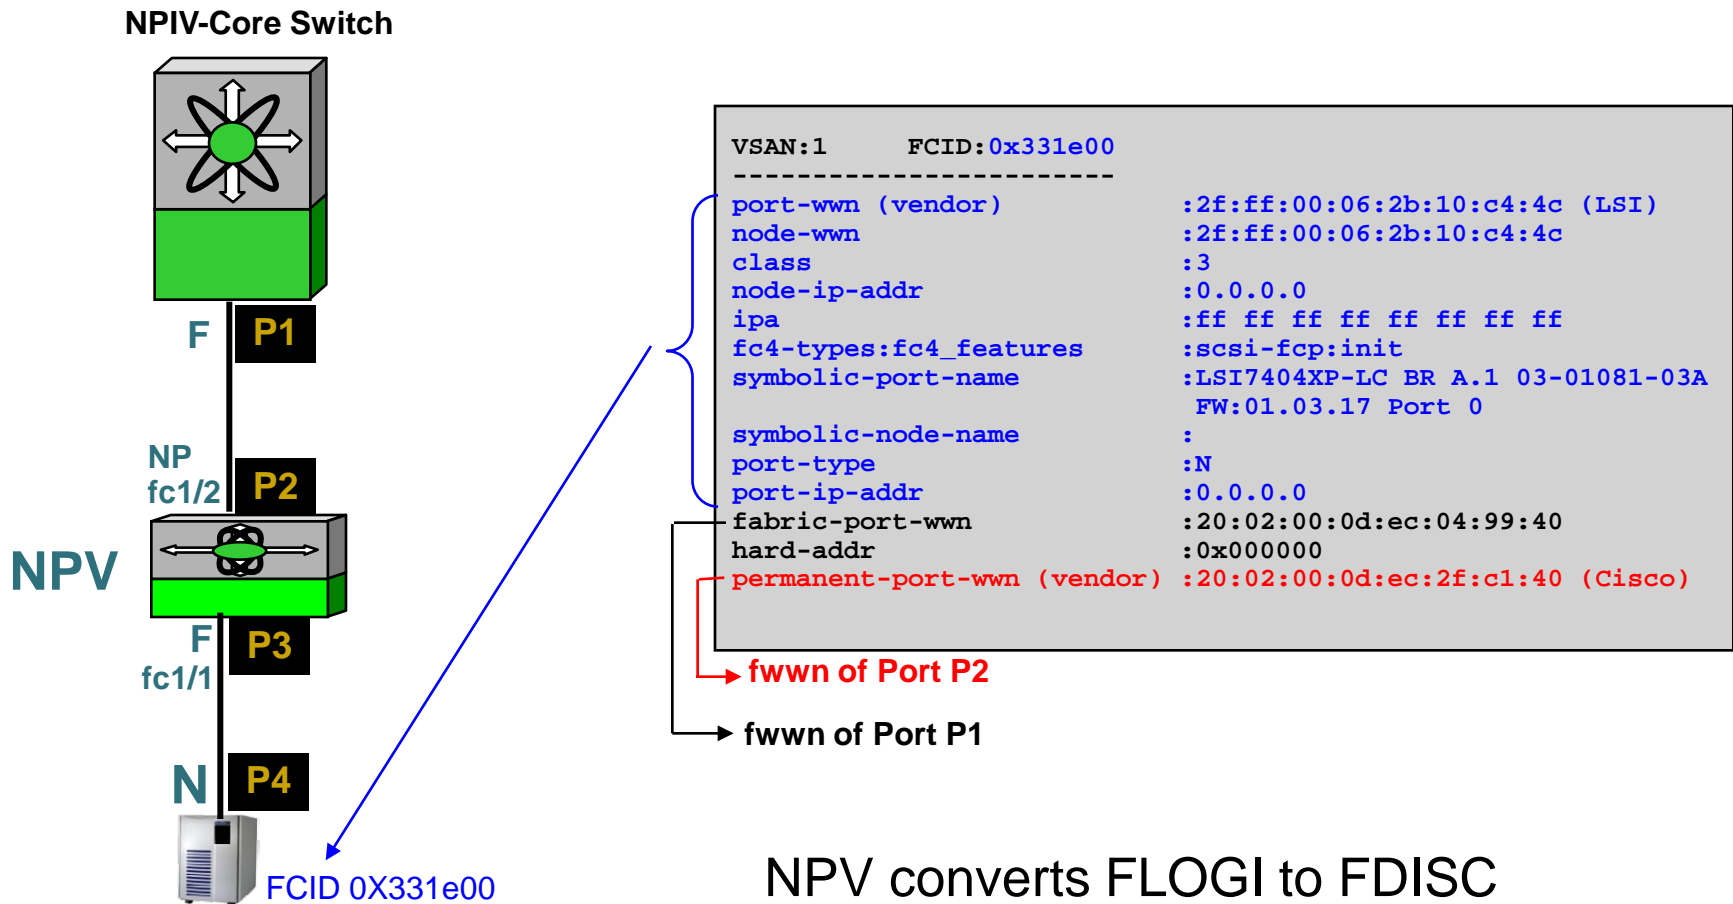
NPV: Internal Logins (FLOGIs)

- When an NP port comes up NPV itself first FLOGIs into the core



NPV: Logins from End Devices (FDISCs)

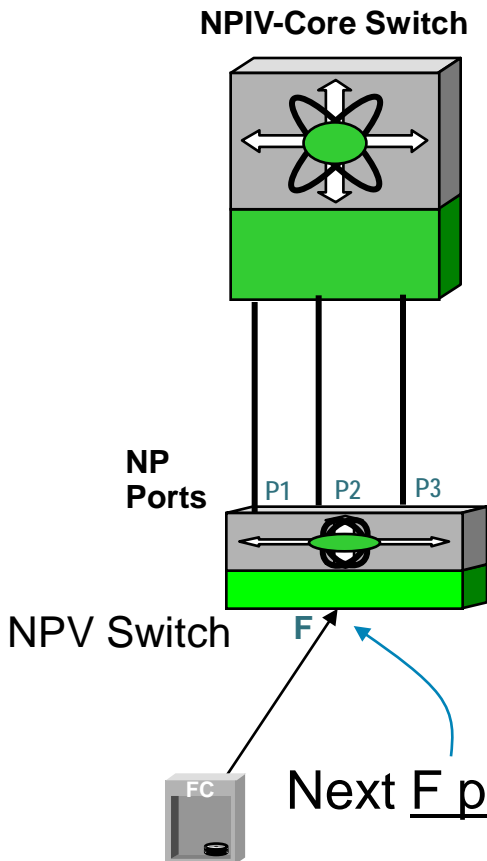
- An end port logs into the npv-core as follows:



NPV: Distribution of End Device Logins

An example of current mapping of ports

<u>NP Port</u>	<u>No. of mapped F ports</u>
P1	5 (fc1/1, 1/5 ...)
P2	2 (fc1/7, fc1/21)
P3	9 (fc1/2, fc1/8, ...)




Next F port on NPV would be assigned to NP Port P2

(NP port with *minimum* number of mapped F ports)


NPV: What Happens When an NP Port Goes Down?

- All F ports mapped to that NP port are re-initialized (shut → no shut)
- These N ports would attempt re-login
- If another NP port is available
 - N ports would be logged in via the available NP port(s)
 - Logins would be distributed as per the previous slide
- If no NP port is available
 - F ports would remain in down state waiting for an NP port
- When the failed NP port comes back up
 - The logins are NOT re-distributed (to avoid disruption)



NPV: What Happens When an F Port Goes Down?

- NPV sends a LOGO to NPV-core on behalf of the N port
 - If there were multiple FCIDs assigned to that N port (NPIV-enabled host) then a LOGO is sent to the NPV-core for each FCID
- The F port is not allowed to come back up until the LOGO(s) are completed
- Conflict in Port-security
 - DPVM or FC-SP configuration on NPV-code switch can logout an FCID or prevent it from coming up



NPV Switch Logon –

Span of F-Port on Core MDS to an SD Port with PAA attached

npv.cap - Wireshark

File Edit View Go Capture Analyze Statistics Help

Filter: Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Info
1	0.000000	00.00.00	ff.ff.fe	FC ELS	FLOGI
2	0.038857	ff.ff.fe	0b.00.00	FC ELS	ACC (FLOGI)
3	0.042278	0b.00.00	ff.ff.fc	FC ELS	PLOGI
4	0.042860	ff.fc.0b	0b.00.00	FC ELS	PLOGI
5	0.049379	ff.ff.fc	0b.00.00	FC ELS	ACC (PLOGI)
6	0.050030	0b.00.00	ff.ff.fc	dns	RFT_ID
7	0.058718	ff.ff.fc	0b.00.00	dns	ACC (RFT_ID)
8	0.059603	0b.00.00	ff.ff.fc	dns	RIP_NN
9	0.061105	ff.ff.fc	0b.00.00	dns	ACC (RIP_NN)
10	0.061991	0b.00.00	ff.ff.fc	dns	RSNN_NN
11	0.063150	ff.ff.fc	0b.00.00	dns	ACC (RSNN_NN)
12	0.064072	0b.00.00	ff.ff.fc	dns	RSPN_ID
13	0.065963	ff.ff.fc	0b.00.00	dns	ACC (RSPN_ID)
14	0.067063	0b.00.00	ff.ff.fa	FC ELS	PLOGI
15	0.069327	ff.ff.fa	0b.00.00	FC ELS	ACC (PLOGI)
16	0.070036	0b.00.00	ff.ff.fa	FC-FCS	GMAL
17	0.071211	ff.ff.fa	0b.00.00	FC-FCS	MSG_ACC (GMAL)
18	0.796523	00.00.00	ff.ff.fe	FC ELS	FDISC
19	0.811353	ff.ff.fe	0b.00.02	FC ELS	ACC (ACC)
20	1.043834	ff.fc.0b	0b.00.00	FC ELS	LOGO
21	1.044466	ff.fc.0b	0b.00.02	FC ELS	PLOGI
22	1.045369	0b.00.02	ff.fc.0b	FC ELS	ACC (PLOGI)
23	1.045516	ff.fc.0b	0b.00.02	FC ELS	PRLI
24	1.046275	0b.00.02	ff.fc.0b	FC ELS	LS_RJT (PRLI)
25	1.047076	ff.fc.0b	0b.00.02	FC ELS	LOGO
26	1.047452	ff.fc.0b	0b.00.00	FC ELS	PLOGI
27	1.047847	0b.00.02	ff.fc.0b	FC ELS	ACC (LOGO)
28	2.043623	ff.fc.0b	0b.00.00	FC ELS	LOGO
29	2.103110	ff.fc.0b	0b.00.02	FC ELS	PLOGI
30	2.104030	0b.00.02	ff.fc.0b	FC ELS	ACC (PLOGI)
31	2.104221	ff.fc.0b	0b.00.02	FC ELS	PRLI
32	2.105002	0b.00.02	ff.fc.0b	FC ELS	LS_RJT (PRLI)
33	2.105787	ff.fc.0b	0b.00.02	FC ELS	LOGO
34	2.106559	0b.00.02	ff.fc.0b	FC ELS	ACC (LOGO)

FLOGI From NPV MDS switch to Core MDS

PLOGI of NPV Switch to Core MDS

Other Plogi Parameter exchanges such as a HBA would do

FDISC- converted FLOGI or connecting host or storage device on the NPV switch

PLOGI accept to the real HBA on the MDS NPV switch, FCID of attached device is 0b0002

Remaining exchanges are from device 0b0002 to the code MDS switch

NPV Related Show Commands on NPV Switch

- The following show commands can be used on the NPV switch to display info. on the NPV devices
- There is no familiar fcns or flogi database to view, for these services are not running in a NPV enabled switch

```
show npv flogi-table
show npv status
show tech-support npv
```

```
show npv internal event-history { errors | events | ext-if-fsm [ <interface > ]
| flogi-fsm [ interface < interface > | pwwn <wwn> ] }
show npv internal event-history msgs
show npv internal event-history svr-if-fsm [ interface < interface > ]
```

```
show npv internal info
show npv internal info external-interface { all | < interface > }
show npv internal info global
show npv internal info interface { all | < interface > }
show npv internal info server-interface { all | < interface > }
show npv internal mem-stats [ detail ]
show npv internal pending-queue interface { all | < interface > }
```

```
debug npv { all | erros | events | ... }
show debug npv
```

NPV Related Show Commands on NPV-Core Switch

- The following show commands can be used in the NPV-core switch to display information on the NPV devices. Since these outputs are based on the name server information, this command can be run from any non-NPV MDS switch running 3.2(1) or later release

```
show fcns database npv [ detail [ vsan <vsan range> ] ]
show fcns database npv [ node_wwn <wwn> ] [ vsan <vsan range> ]
```

- Example Outputs

```
npv# show fcns database npv
```

```
VSAN 1:
```

```
-----
NPV NODE-NAME      NPV IP_ADDR      NPV IF  CORE SWITCH WWN      CORE IF
-----
20:00:00:0d:ec:3d:62:80  10.1.96.24      fc1/20  20:00:00:0d:ec:2d:af:40  fc4/4
20:00:00:0d:ec:3d:62:80  10.1.96.24      fc1/19  20:00:00:0d:ec:2d:af:40  fc4/3
20:00:00:0d:ec:3d:62:80  10.1.96.24      fc1/17  20:00:00:0d:ec:2d:af:40  fc4/1
```

```
...
```


NPV Related Commands on NPV-Core Switch

```
npv# show fcns database npv detail
```

```
-----  
VSAN:1    NPV Node Name: 20:00:00:0d:ec:3d:62:80  
-----  
NPV Fabric Port-WWN      :20:14:00:0d:ec:3d:62:80  
class                    :2,3  
NPV IP Address           :10.1.96.24  
ipa                      :ff ff ff ff ff ff ff ff  
fc4-types:fc4_features   :npv  
NPV Switch Name:Interface :sw24-gd96:fc1/20  
port-type                :N P  
Core Switch fabric-port-wwn :20:c4:00:0d:ec:2d:af:40  
permanent-port-wwn (vendor) :20:14:00:0d:ec:3d:62:80 (Cisco)
```

....

```
npv# show fcns database npv node_wwn 20:00:00:0d:ec:3d:42:40
```

```
VSAN 1:
```

```
-----  
FCID      TYPE  PWWN                (VENDOR)      FC4-TYPE:FEATURE  
-----  
0x330f00  N    2f:ff:00:06:2b:10:c7:b2 (LSI)        scsi-fcp:init  
0x331000  N    2f:ff:00:06:2b:10:c7:b3 (LSI)        scsi-fcp:init
```

```
Total number of npv-attached entries = 2
```

Wrap-Up

- Cornerstone to SAN network troubleshooting is to understand Standards operation
- Each SAN/OS feature has methods of troubleshooting, Basic most seen issues were covered in this session
- Familiarity with available tools in SAN/OS will narrow time to resolution if either O/S bug or operational issues
- Have other insight on HBA operations & firmware, Array tools and Applications
- Interop with Brocade/McData requires another skill set
- Understanding architecture of Cisco Fabric Manager deployments and best practices also a plus

Core Dumps

- Show cores

```
Guernsey# sh cores
```

Module-num	Process-name	PID	Core-create-time
-----	-----	-----	-----
1	cimxmlserver	20029	Jul 18 08:39

- Configure switch for core dumps

```
Switch# sh system cores
```

```
Cores are transferred to tftp://10.91.42.133/
```

Show process log will display cores dumped to server

```
Switch# sh processes log
```

Process	PID	Normal-exit	Stack	Core	Log-create-time
-----	-----	-----	-----	-----	-----
SystemHealth	27828	N	Y	N	Tue Dec 7 19:08:09 2004
SystemHealth	27880	N	Y	N	Tue Dec 7 19:08:20 2004
SystemHealth	27934	N	Y	N	Tue Dec 7 19:08:30 2004
sme	2030	N	Y	N	Sun Sep 23 18:47:15 2007
sme	2306	N	Y	N	Sun Sep 23 18:47:17 2007
syslogd	2271	N	N	N	Thu Sep 7 13:29:12 2006
syslogd	2442	N	N	N	Thu Sep 7 13:30:12 2006
syslogd	2510	N	N	N	Thu Sep 7 13:31:12 2006

Q and A



Interested in Data Center?

- Discover the Data Center of the Future

Cisco booth: #617

See a simulated data center and discover the benefits including investing to save, energy efficiency and innovation.

- Data Center Booth

Come by and see what's happening in the world of Data Center – demos; social media activities; bloggers; author signings

Demos include:

- Unified Computing Systems

- Cisco on Cisco Data Center Interactive Tour

- Unified Service Delivery for Service Providers

- Advanced Services


Interested in Data Center?

- Data Center Super Session
 - Data Center Virtualization Architectures, Road to Cloud Computing (UCS)
 - Wednesday, July 1, 2:30 – 3:30 pm, Hall D
 - Speakers: John McCool and Ed Bugnion
- Panel: 10 Gig LOM
 - Wednesday 08:00 AM Moscone S303
- Panel: Next Generation Data Center
 - Wednesday 04:00 PM Moscone S303
- Panel: Mobility in the DC Data
 - Thursday 08:00 AM Moscone S303

Please Visit the Cisco Booth in the World of Solutions


See the technology in action

- Data Center and Virtualization
 - DC1 – Cisco Unified Computing System
 - DC2 – Data Center Switching: Cisco Nexus and Catalyst
 - DC3 – Unified Fabric Solutions
 - DC4 – Data Center Switching: Cisco Nexus and Catalyst
 - DC5 – Data Center 3.0: Accelerate Your Business, Optimize Your Future
 - DC6 – Storage Area Networking: MDS
 - DC7 – Application Networking Systems: WAAS and ACE



Recommended Reading

- Storage Networking Protocol Fundamentals, ISBN: 1-58705-160-5
- Cisco Storage Networking Architectures Poster, ISBN: 1-58720-102-X
- The Business Case for Storage Networks, ISBN: 1-58720-118-6
- Storage Networking Fundamentals: An Introduction to Storage Devices, Subsystems, Applications, Management, and File Systems, ISBN: 1-58705-162-1



Complete Your Online Session Evaluation

- Give us your feedback and you could win fabulous prizes. Winners announced daily.
- Receive 20 Passport points for each session evaluation you complete.
- Complete your session evaluation online now (open a browser through our wireless network to access our portal) or visit one of the Internet stations throughout the Convention Center.



Don't forget to activate your Cisco Live Virtual account for access to all session material, communities, and on-demand and live activities throughout the year. Activate your account at the Cisco booth in the World of Solutions or visit www.ciscolive.com.



Appendix - Extras



Necessary Network Settings

- Switch Fabric clocks synchronized
 - Use NTP along with CFS to simplify job
 - Use fabric manager to set all clocks and time zones
- Setup system logs, syslogd server
 - Use CFS to simplify
- Save off switch configurations or even complete show tech-support details regularly
 - Use CLI scheduler in SAN/OS to simplify process

Errors on Application Server Logs, User Having Performance Problems

- Use logging capabilities on switch to piece together network issues
- Put together:
 - What interfaces are on the fabric (up, down, flapping)
 - What was health of each interface (virtual and physical interfaces)
 - Examine interface health (FC protocol errors, physical layer errors)
 - Are effects being seen on the complete network fabric or within the VSAN (RSCNs, Zones, ISLs, Errors on common controllers)

Gathering Internal Counter Information for Unknown Issues and Plaguing Connectivity Problems

- Determine what ports on which switches you need to examine, this is done by narrowing to physical switch, VSAN, zone
- Look at path through complete network, initiator and target-side along with ISLs
- One stop shopping on each switch for all the data:
Attach to Line Card Module that has the interface you need information on

```
module-1# terminal length 0
```

Set scroll back buffer or log to file on telnet tool

```
module-1# show hardware internal debug-info interface fc1/2
```

Based on Line Card type proper information will be output

Error Disabled Ports

```
fc1/13 is down (Error disabled - bit error rate too high)
Port description is MP3_JBOD
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:0d:00:0d:ec:00:ea:40
Admin port mode is auto, trunk mode is auto
Port vsan is 5
Receive data field Size is 2112
Beacon is turned off
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
6700141425 frames input, 9795578489624 bytes
 0 discards, 0 errors
 0 CRC, 0 unknown class
 0 too long, 0 too short
8930554528 frames output, 15656447669120 bytes
 0 discards, 0 errors
0 input OLS, 0 LRR, 0 NOS, 2 loop inits
0 output OLS, 0 LRR, 0 NOS, 2 loop inits
0 receive B2B credit remaining
0 transmit B2B credit remaining
```

- In FC-MAC driver the ITW (Invalid Transmission Words) errors counter for the port is used to calculate bit-error-rate thresholding as described above. The ITW counter value is counted only when the port is in up state i.e. ACTIVE/MONITORING state is reached on the link. The ITW error can occur due to the following reasons –
 - 1. Faulty or bad cable
 - 2. Faulty or bad GBIC/SFP
 - 3. GBIC/SFP specified to operate in 1G speed but it is being used for 2G speed
 - 4. Short haul cable used for long haul or vice versa.
 - 5. Momentary sync loss most probably due to one of the above reasons.
 - 6. Loose cable connection at one or both ends.
 - 7. Improper GBIC/SFP connection at one or both ends

Core Dumps

- Show cores

```
Guernsey# sh cores
```

Module-num	Process-name	PID	Core-create-time
1	cimxmlserver	20029	Jul 18 08:39

- Configure switch for core dumps

```
Switch# sh system cores
```

```
Cores are transferred to tftp://10.91.42.133/
```

Show process log will display cores dumped to server

```
Switch# sh processes log
```

Process	PID	Normal-exit	Stack	Core	Log-create-time
SystemHealth	27828	N	Y	N	Tue Dec 7 19:08:09 2004
SystemHealth	27880	N	Y	N	Tue Dec 7 19:08:20 2004
SystemHealth	27934	N	Y	N	Tue Dec 7 19:08:30 2004
sme	2030	N	Y	N	Sun Sep 23 18:47:15 2007
sme	2306	N	Y	N	Sun Sep 23 18:47:17 2007
syslogd	2271	N	N	N	Thu Sep 7 13:29:12 2006
syslogd	2442	N	N	N	Thu Sep 7 13:30:12 2006
syslogd	2510	N	N	N	Thu Sep 7 13:31:12 2006