Protocol Analysis 201 for High-Speed Fibre Channel Fabrics

Live Webcast
April 11, 2019
10:00 AM PT

Today's Speakers

David Rodgers

Sr. Product Marketing Manager

BoD, FCIA

Teledyne LeCroy

Yamini Shastry

Principal Engineer

Viavi Solutions

Joseph Kimpler

Global Manager, Alliances and Enablement

ATTO Technology

About the FCIA

- The Fibre Channel Industry Association (FCIA) is a mutual benefit, non-profit, international organization of manufacturers, system integrators, developers, vendors, and industry professionals, and end users
  - Promotes the advancement of Fibre Channel technologies and products that conform to the existing and emerging T11 standards.
  - Maintains resources and supports activities to ensure multi-vendor interoperability for hardware, interconnection, and protocol solutions.
  - Promotion and marketing of FC solutions, educational awareness campaigns, hosting public interoperability demonstrations, and fostering technology and standards conformance.

https://fibrechannel.org/

FCIA
FIBRE CHANNEL INDUSTRY ASSOCIATION

FCIA
FIBRE CHANNEL INDUSTRY ASSOCIATION

Agenda

FC SAN Elements

Analyzer placement

Getting most out of the capture

Post-capture analysis

Trace Formats, Trace reading, Graphing

Metric groups - Deeper analysis

Debug Examples

Benefits of purposeful error injection

Layer 2, 3, 4 Jamming examples

Fibre Channel Layers*

FC has functional layers:

- **FC-0**: The interface to the physical media; transceivers, cables, etc.
- **FC-1**: Transmission protocol or data-link layer, encodes and decodes signals
- **FC-2**: Network Layer consists of the low level Fibre Channel protocols; port to port connections.
- **FC-3**: Common services layer, a thin layer that could eventually implement functions like encryption or RAID redundancy algorithms; multiport connections
- **FC-4**: Protocol-mapping layer, in which upper level protocols such as NVMe, SCSI, IP or FICON, are encapsulated into Information Units (IUs) for delivery to FC-2.

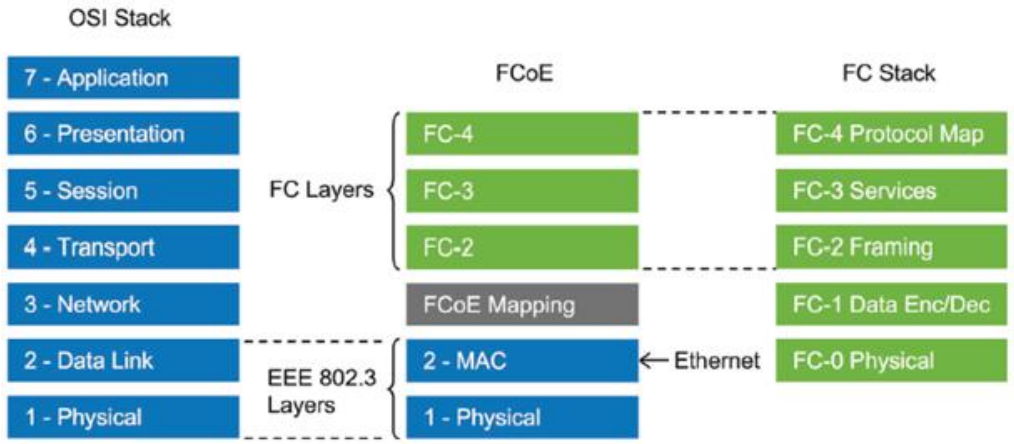

*Reprinted from Fibre Channel Interoperability Webcast, https://fibrechannel.org/webcasts/

Device and Interop Testing

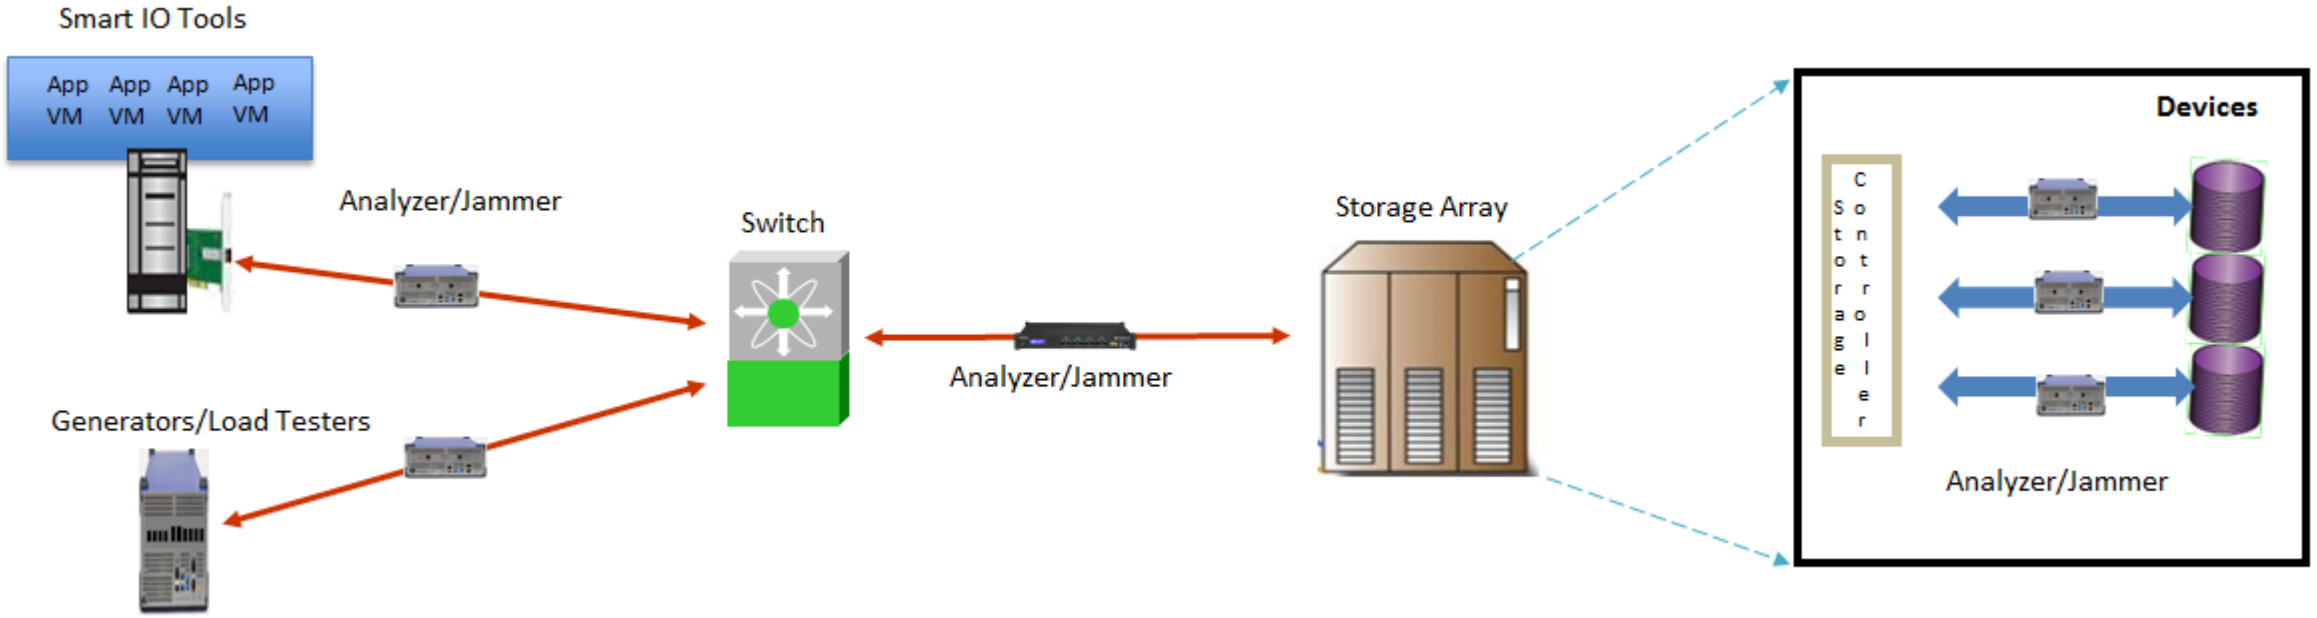

Generators/IO Tools
- Performance Testing
- Compliance Testing
- Functionality Testing
- Data Integrity Testing

Analyzer
- Capture and Analysis
- Protocol Violations/Errors
- Interop Testing

Jammer
- Error Injection
- Error Recovery

FC Switched Link Elements

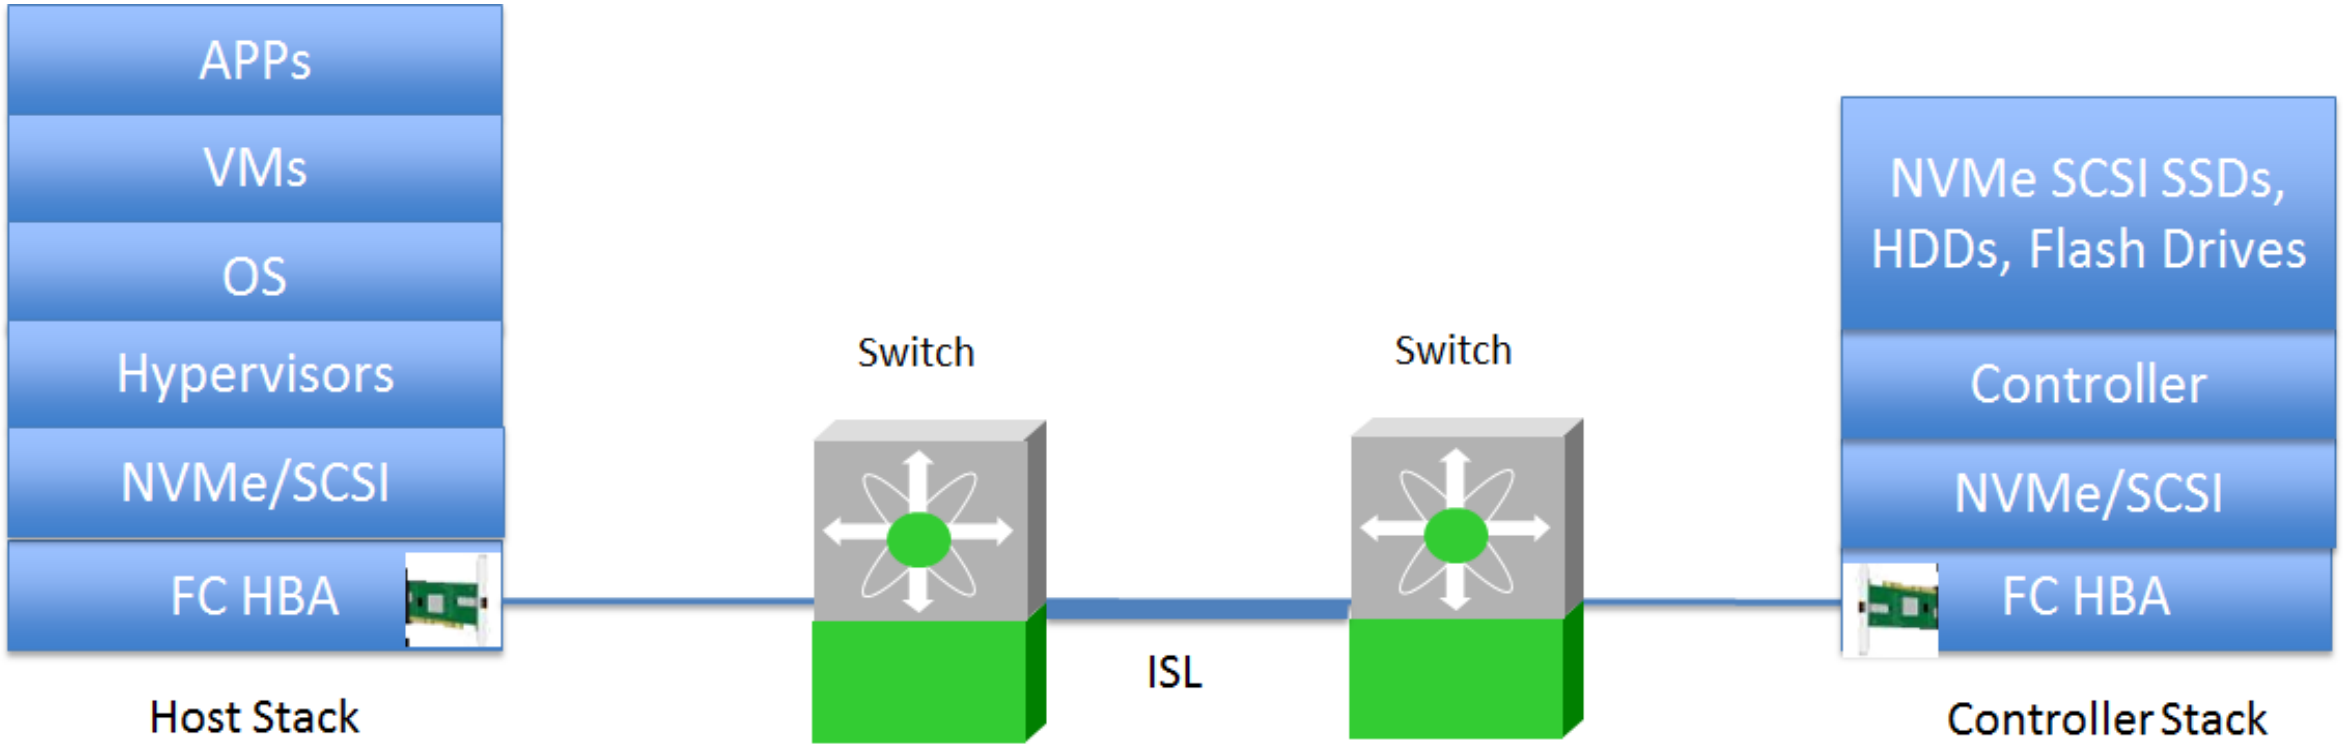

FC Switched Link Elements - Specific to Analyzer

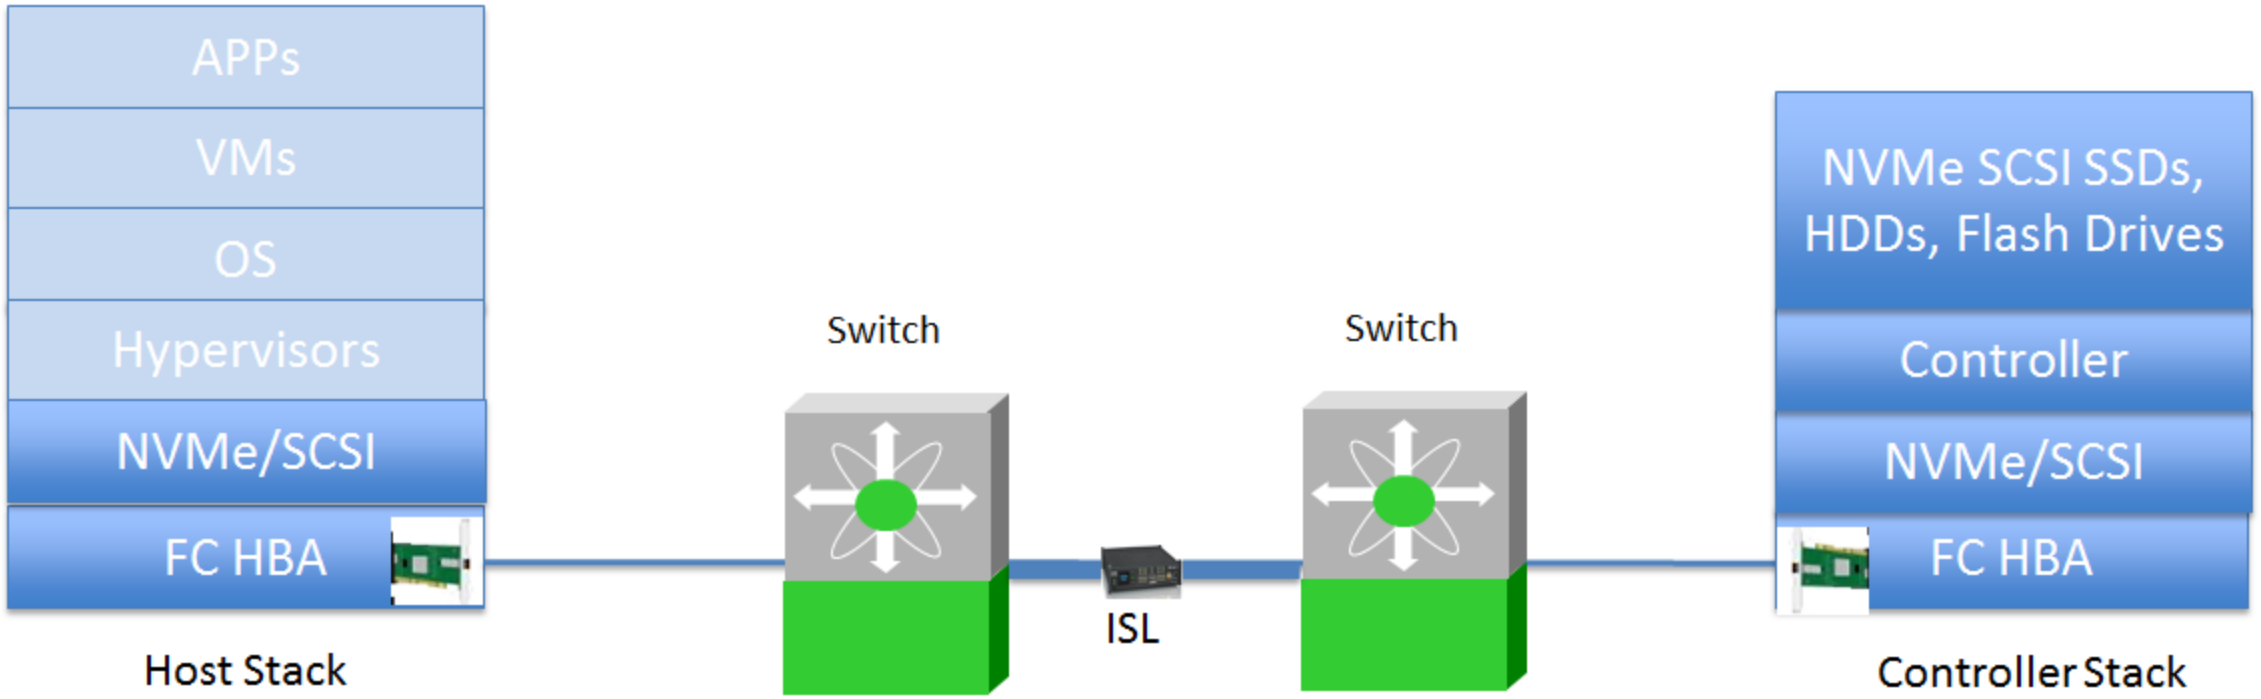

SAN Issues

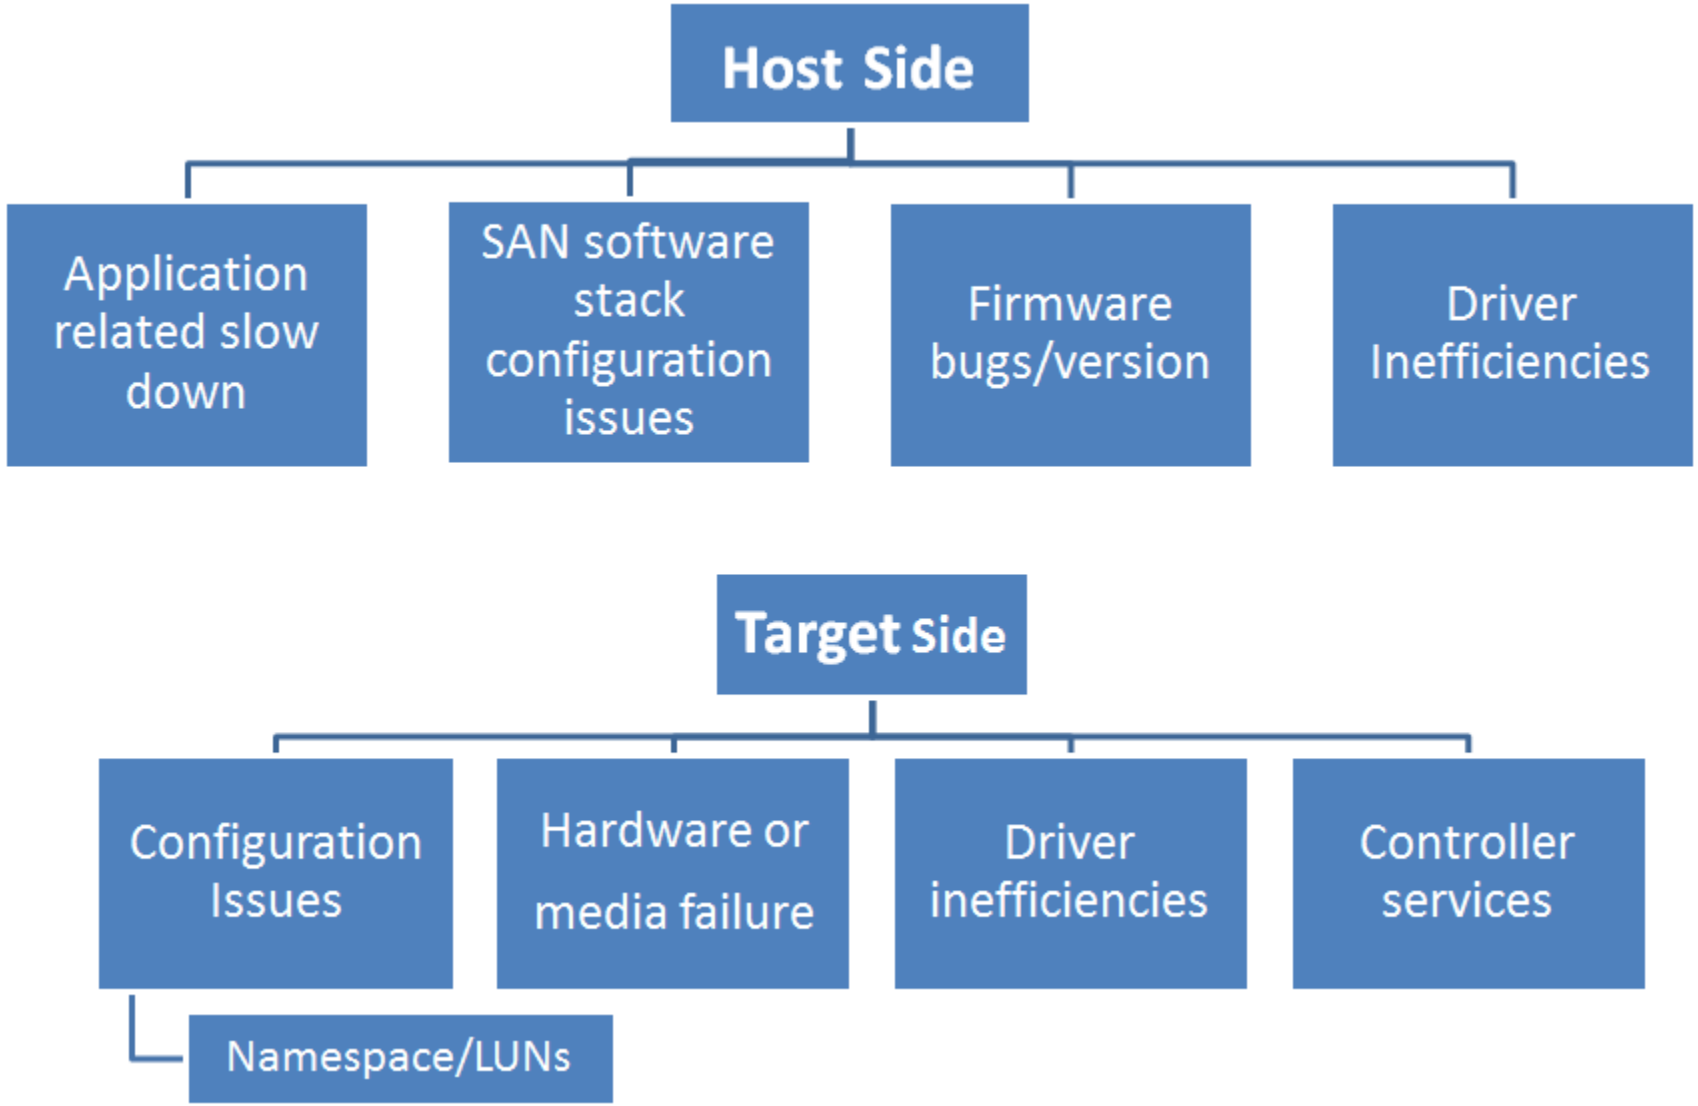
Host Side
Application related slow down
SAN software stack configuration issues
Firmware bugs/version
Driver Inefficiencies
Target Side
Configuration Issues
Hardware or media failure
Driver inefficiencies
Controller services
Namespace/LUNs

SAN Issues

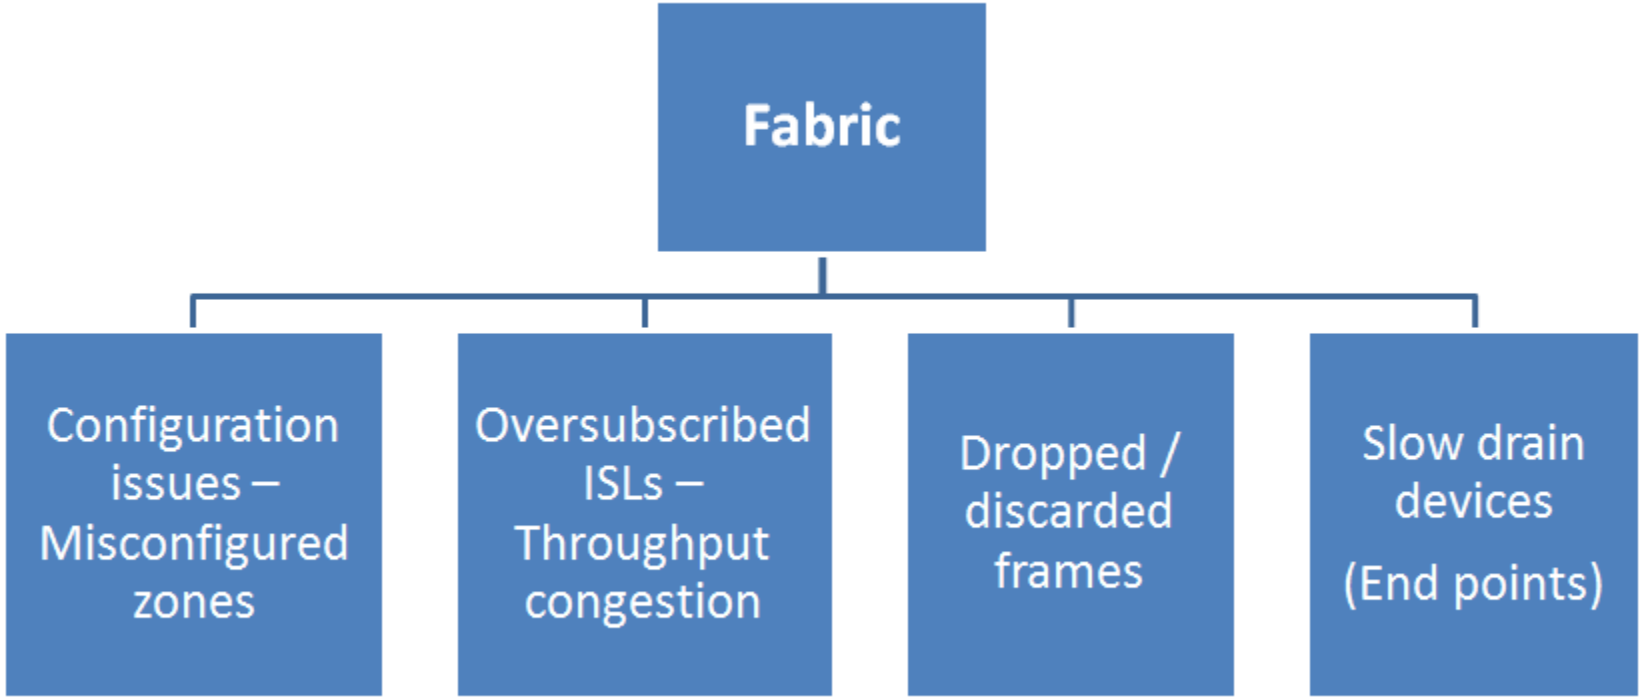
Fabric
Configuration issues – Misconfigured zones
Oversubscribed ISLs – Throughput congestion
Dropped / discarded frames
Slow drain devices
(End points)

Troubleshooting

Problem Detection

- Physical errors
- Link resets/Aborts
- Missing Targets(LUNs) or ports
- Performance problem

Problem Monitoring

- Fabric monitoring tools
- Port logs/Device counters
- Compare: Problematic SAN configs/performance to baseline configs/

Is the problem from the host side, target side, fabric links or can't be determined

Can the issue be recreated?

Can it be captured on analyzer????

Analyzer Placement and Setup

Placement of analyzer

- Host Link
- Target Link
- On the ISL
- Each end of the network

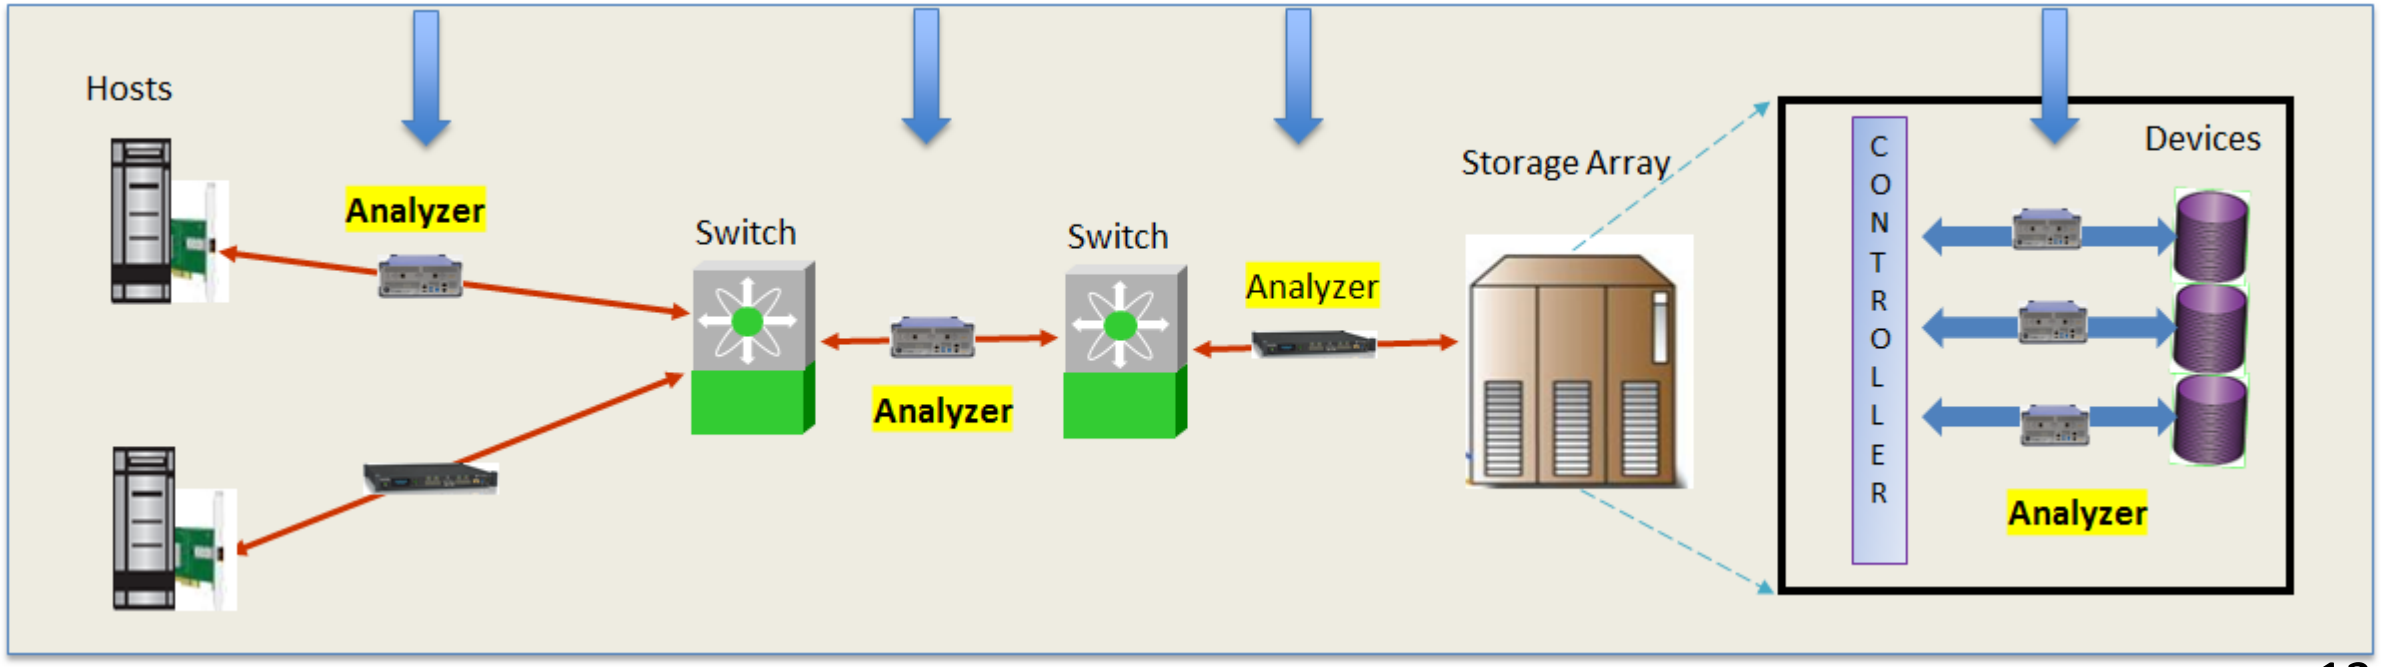

Getting the Most of a Capture

Setting up triggers and filters

- Link issues: Trigger on an error (Capture every word)
  - Stateless errors - simple triggers
  - Stateful errors - complex triggers
- Credit issues: Trigger(frame loss, credit loss), deep buffers, capture all Frames and R_RDYs
- Latency Issues: Visual trigger (throughput dip/threshold triggers), Filter out data frames
- Data Corruption issues: Visual trigger (app failure or bluescreen), capture frames only

Post-capture analysis

- Error Counters
- Frame to R_RDY timings, time on idle link, Out of credit situations w.r.t Frame counts
- Exchange latencies – SCSI/NVMe
- Data comparison, command overlap metrics

Trace Format Options

- **Event/Exchange Based View**
  - Group trace information based on event/exchange
  - Easily visualize and analyze transactions, make measurements
- **Important for the FW/ Driver developer, SW/ application engineer**

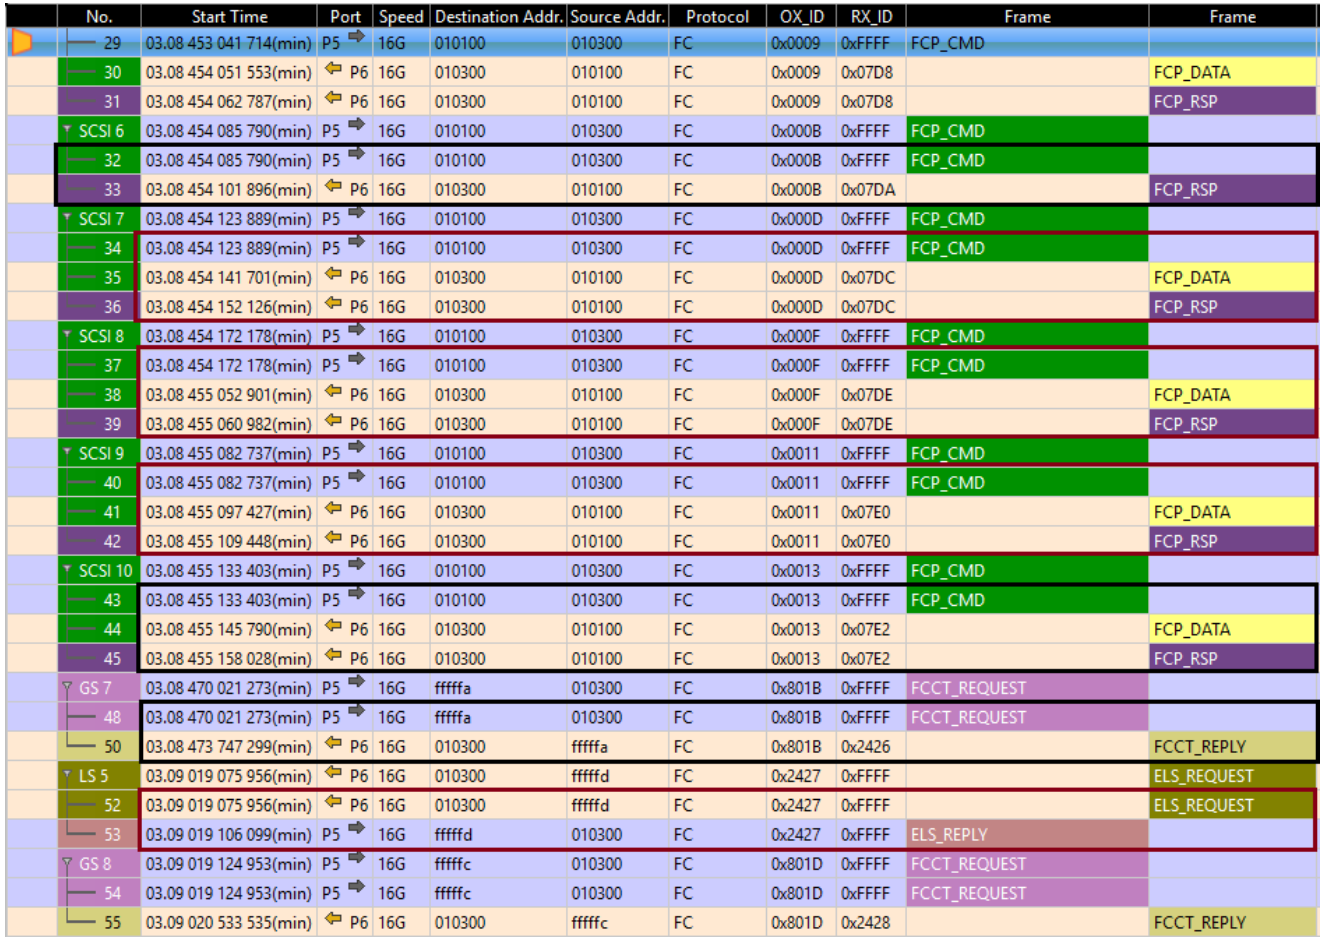

No.	Start Time	Port	Speed	Destination Addr.	Source Addr.	Protocol	OX_ID	RX_ID	Frame	Frame
29	03.08 453 041 714(min)	P5 →	16G	010100	010300	FC	0x0009	0xFFFF	FCP_CMD	
30	03.08 454 051 553(min)	← P6	16G	010300	010100	FC	0x0009	0x07D8		FCP_DATA
31	03.08 454 062 787(min)	← P6	16G	010300	010100	FC	0x0009	0x07D8		FCP_RSP
SCSI 6	03.08 454 085 790(min)	P5 →	16G	010100	010300	FC	0x000B	0xFFFF	FCP_CMD	
32	03.08 454 085 790(min)	P5 →	16G	010100	010300	FC	0x000B	0xFFFF	FCP_CMD	
33	03.08 454 101 896(min)	← P6	16G	010300	010100	FC	0x000B	0x07DA		FCP_RSP
SCSI 7	03.08 454 123 889(min)	P5 →	16G	010100	010300	FC	0x000D	0xFFFF	FCP_CMD	
34	03.08 454 123 889(min)	P5 →	16G	010100	010300	FC	0x000D	0xFFFF	FCP_CMD	
35	03.08 454 141 701(min)	← P6	16G	010300	010100	FC	0x000D	0x07DC		FCP_DATA
36	03.08 454 152 126(min)	← P6	16G	010300	010100	FC	0x000D	0x07DC		FCP_RSP
SCSI 8	03.08 454 172 178(min)	P5 →	16G	010100	010300	FC	0x000F	0xFFFF	FCP_CMD	
37	03.08 454 172 178(min)	P5 →	16G	010100	010300	FC	0x000F	0xFFFF	FCP_CMD	
38	03.08 455 052 901(min)	← P6	16G	010300	010100	FC	0x000F	0x07DE		FCP_DATA
39	03.08 455 060 982(min)	← P6	16G	010300	010100	FC	0x000F	0x07DE		FCP_RSP
SCSI 9	03.08 455 082 737(min)	P5 →	16G	010100	010300	FC	0x0011	0xFFFF	FCP_CMD	
40	03.08 455 082 737(min)	P5 →	16G	010100	010300	FC	0x0011	0xFFFF	FCP_CMD	
41	03.08 455 097 427(min)	← P6	16G	010300	010100	FC	0x0011	0x07E0		FCP_DATA
42	03.08 455 109 448(min)	← P6	16G	010300	010100	FC	0x0011	0x07E0		FCP_RSP
SCSI 10	03.08 455 133 403(min)	P5 →	16G	010100	010300	FC	0x0013	0xFFFF	FCP_CMD	
43	03.08 455 133 403(min)	P5 →	16G	010100	010300	FC	0x0013	0xFFFF	FCP_CMD	
44	03.08 455 145 790(min)	← P6	16G	010300	010100	FC	0x0013	0x07E2		FCP_DATA
45	03.08 455 158 028(min)	← P6	16G	010300	010100	FC	0x0013	0x07E2		FCP_RSP
GS 7	03.08 470 021 273(min)	P5 →	16G	fffffa	010300	FC	0x801B	0xFFFF	FCCT_REQUEST	
48	03.08 470 021 273(min)	P5 →	16G	fffffa	010300	FC	0x801B	0xFFFF	FCCT_REQUEST	
50	03.08 473 747 299(min)	← P6	16G	010300	fffffa	FC	0x801B	0x2426		FCCT_REPLY
LS 5	03.09 019 075 956(min)	← P6	16G	010300	fffffd	FC	0x2427	0xFFFF		ELS_REQUEST
52	03.09 019 075 956(min)	← P6	16G	010300	fffffd	FC	0x2427	0xFFFF		ELS_REQUEST
53	03.09 019 106 099(min)	P5 →	16G	fffffd	010300	FC	0x2427	0xFFFF	ELS_REPLY	
GS 8	03.09 019 124 953(min)	P5 →	16G	fffffc	010300	FC	0x801D	0xFFFF	FCCT_REQUEST	
54	03.09 019 124 953(min)	P5 →	16G	fffffc	010300	FC	0x801D	0xFFFF	FCCT_REQUEST	
55	03.09 020 533 535(min)	← P6	16G	010300	fffffc	FC	0x801D	0x2428		FCCT_REPLY

Trace Format Options

- **Chronological Based View**
  - Group trace based on events over time
  - Represents the transactions over the wire in order presented
  - Correlate information port by port
- **Important to HW/PHY layer engineering effort**

No.	Start Time	Port	Speed	Destination Addr.	Source Addr.	Protocol	OX_ID	RX_ID	Frame	Frame
29	03.08 453 041 714(min)	P5 →	16G	010100	010300	FC	0x0009	0xFFFF	FCP_CMD	
30	03.08 454 051 553(min)	← P6	16G	010300	010100	FC	0x0009	0x07D8		FCP_DATA
31	03.08 454 062 787(min)	← P6	16G	010300	010100	FC	0x0009	0x07D8		FCP_RSP
32	03.08 454 085 790(min)	P5 →	16G	010100	010300	FC	0x000B	0xFFFF	FCP_CMD	
33	03.08 454 101 896(min)	← P6	16G	010300	010100	FC	0x000B	0x07DA		FCP_RSP
34	03.08 454 123 889(min)	P5 →	16G	010100	010300	FC	0x000D	0xFFFF	FCP_CMD	
35	03.08 454 141 701(min)	← P6	16G	010300	010100	FC	0x000D	0x07DC		FCP_DATA
36	03.08 454 152 126(min)	← P6	16G	010300	010100	FC	0x000D	0x07DC		FCP_RSP
37	03.08 454 172 178(min)	P5 →	16G	010100	010300	FC	0x000F	0xFFFF	FCP_CMD	
38	03.08 455 052 901(min)	← P6	16G	010300	010100	FC	0x000F	0x07DE		FCP_DATA
39	03.08 455 060 982(min)	← P6	16G	010300	010100	FC	0x000F	0x07DE		FCP_RSP
40	03.08 455 082 737(min)	P5 →	16G	010100	010300	FC	0x0011	0xFFFF	FCP_CMD	
41	03.08 455 097 427(min)	← P6	16G	010300	010100	FC	0x0011	0x07E0		FCP_DATA
42	03.08 455 109 448(min)	← P6	16G	010300	010100	FC	0x0011	0x07E0		FCP_RSP
43	03.08 455 133 403(min)	P5 →	16G	010100	010300	FC	0x0013	0xFFFF	FCP_CMD	
44	03.08 455 145 790(min)	← P6	16G	010300	010100	FC	0x0013	0x07E2		FCP_DATA
45	03.08 455 158 028(min)	← P6	16G	010300	010100	FC	0x0013	0x07E2		FCP_RSP
46	03.08 459 175 660(min)	← P6	16G	010300	fffc01	FC	0x8019	0x2425		ELS_REPLY
47	03.08 469 983 077(min)	← P6	16G	010300	fffffa	FC	0x8011	0x2423		FCCT_REPLY
48	03.08 470 021 273(min)	P5 →	16G	fffffa	010300	FC	0x801B	0xFFFF	FCCT_REQUEST	
49	03.08 470 778 927(min)	← P6	16G	010300	fffffa	FC	0x8013	0x2424		FCCT_REPLY
50	03.08 473 747 299(min)	← P6	16G	010300	fffffa	FC	0x801B	0x2426		FCCT_REPLY
52	03.09 019 075 956(min)	← P6	16G	010300	fffffd	FC	0x2427	0xFFFF		ELS_REQUEST
53	03.09 019 106 099(min)	P5 →	16G	fffffd	010300	FC	0x2427	0xFFFF	ELS_REPLY	
54	03.09 019 124 953(min)	P5 →	16G	fffffc	010300	FC	0x801D	0xFFFF	FCCT_REQUEST	
55	03.09 020 533 535(min)	← P6	16G	010300	fffffc	FC	0x801D	0x2428		FCCT_REPLY
56	03.09 020 560 851(min)	P5 →	16G	010100	010300	FC	0x801F	0xFFFF	ELS_REQUEST	
57	03.09 020 593 983(min)	← P6	16G	010300	010100	FC	0x801F	0xFFFF		ELS_REPLY

Determine Metrics

- **Common Exchange Characteristics Compare expected with Actual**
  - FCCT_REQUEST to FCCT_REPLY
    - Directory Services request to the Name Server with a GS_ACCEPT followed by the ESL REQUEST for PLOGI
  - Time Stamp delta
    - Request at: 03.08.379.946.147min
    - Reply at: 03.08.399.325.845min
    - ET: **19.379.698(ms)**
  - Time Stamp delta
    - Request at: 03.08.419.321.626min
    - Reply at: 03.08.420.053.156min
    - ET: **731.530(us)**

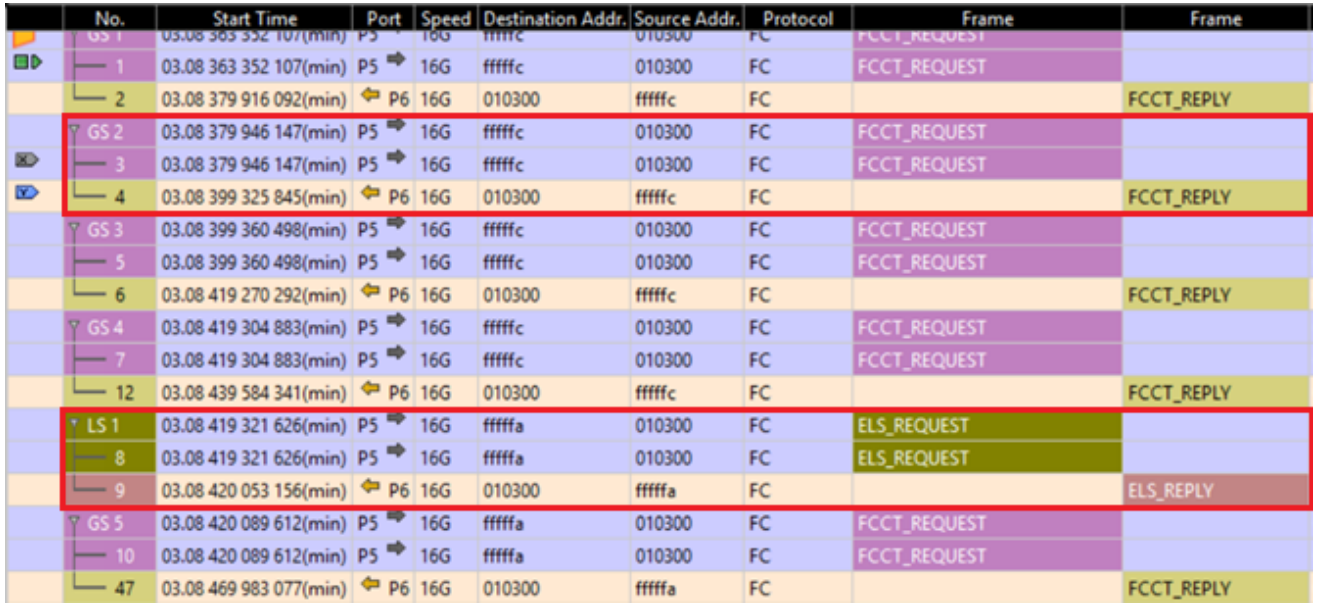

No.	Start Time	Port	Speed	Destination Addr.	Source Addr.	Protocol	Frame	Frame
1	03.08 363 352 107(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
2	03.08 379 916 092(min)	P6	16G	010300	fffffc	FC		FCCT_REPLY
GS 2	03.08 379 946 147(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
3	03.08 379 946 147(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
4	03.08 399 325 845(min)	P6	16G	010300	fffffc	FC		FCCT_REPLY
GS 3	03.08 399 360 498(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
5	03.08 399 360 498(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
6	03.08 419 270 292(min)	P6	16G	010300	fffffc	FC		FCCT_REPLY
GS 4	03.08 419 304 883(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
7	03.08 419 304 883(min)	P5	16G	fffffc	010300	FC	FCCT_REQUEST	
12	03.08 439 584 341(min)	P6	16G	010300	fffffc	FC		FCCT_REPLY
LS 1	03.08 419 321 626(min)	P5	16G	fffffa	010300	FC	ELS_REQUEST	
8	03.08 419 321 626(min)	P5	16G	fffffa	010300	FC	ELS_REQUEST	
9	03.08 420 053 156(min)	P6	16G	010300	fffffa	FC		ELS_REPLY
GS 5	03.08 420 089 612(min)	P5	16G	fffffa	010300	FC	FCCT_REQUEST	
10	03.08 420 089 612(min)	P5	16G	fffffa	010300	FC	FCCT_REQUEST	
47	03.08 469 983 077(min)	P6	16G	010300	fffffa	FC		FCCT_REPLY

Determine Metrics

Exchange Characteristics Ex.1

- FCP_CMD to FCP_RSP
- Time Stamp delta
- Data Length = 36Bytes
  - Request at: 03.08.439.852.471min
  - Reply at: 03.08.440.069.484min
  - ET: **217.013(us)**

Exchange Characteristics Ex. 2

- FCP_CMD to FCP_RSP
- Time Stamp delta
- Data Length = 64Bytes
  - Request at: 03.08.440.086.418min
  - Reply at: 03.08.441.079.592min
  - ET: **993.174(us)**

SCSI 1	03.08 439 852 471(min)	P5	16G	010100	010300	FC	FCP_CMD		0x12:Inquiry ; 0x00:None ; FCP LUN=0x0000 ; FCP_DL=0x00000024
18	03.08 439 852 471(min)	P5	16G	010100	010300	FC	FCP_CMD		0x12:Inquiry ; 0x00:None ; FCP LUN=0x0000 ; FCP_DL=0x00000024
19	03.08 440 052 760(min)	P6	16G	010300	010100	FC		FCP_DATA	Data Length=36 Byte(s)
20	03.08 440 069 484(min)	P6	16G	010300	010100	FC		FCP_RSP	0x00:Good
SCSI 2	03.08 440 086 418(min)	P5	16G	010100	010300	FC	FCP_CMD		0x12:Inquiry ; 0x00:None ; FCP LUN=0x0000 ; FCP_DL=0x000000FF
21	03.08 440 086 418(min)	P5	16G	010100	010300	FC	FCP_CMD		0x12:Inquiry ; 0x00:None ; FCP LUN=0x0000 ; FCP_DL=0x000000FF
22	03.08 441 067 955(min)	P6	16G	010300	010100	FC		FCP_DATA	Data Length=64 Byte(s)
23	03.08 441 079 592(min)	P6	16G	010300	010100	FC		FCP_RSP	0x00:Good

FC Trace Events, SCSI Performance

- **Example of Summary performance for SCSI Events**
  - Extracted directly from the captured transactions
  - Operations represented in S_ID/D_ID pairing
    - **Is Min-Max response time within expected values?**

SCSI.FC-SCSI

Port No	FC Source ID	FC Destination ID	Command	Status	Count	%
P5	010300	010100	Inquiry	Good	4	0.09
P5	010300	010100	Mode Sense (6)	Check Condition	1	0.02
P5	010300	010100	Test Unit Ready	Good	1	0.02
P5	010300	010100	Report LUNS	Good	1	0.02
P5	010300	010100	Report Target Port Groups	Good	1	0.02
P5	010300	010100	Read(10)	Good	4218	98.32
P5	010300	010100	Read Capacity(16)	Good	1	0.02
P5	010300	010100	Read(10)	Incomplete	62	1.45
P5	010300	010100	Mode Sense (6)	Good	1	0.02
					Total: 4290	

SCSI.FC-SCSI.SCSI Timing

Port No	FC Source ID	FC Destination ID	Min Response Time	Max Response Time	Average Response Time	Total	%
P5	010300	010100	015.697(us)	001.021 110(ms)	394.846(us)	4290	100.00
						Total: 4290	

SCSI.FC-SCSI.SCSI Timing.SCSI Read Operation

Port No	FC Source ID	FC Destination ID	Min Response Time	Max Response Time	Average Response Time	Min Throughput (MB/s)	Max Throughput (MB/s)	Average Throughput (MB/s)	Min Latency Time	Max Latency Time	Average Latency Time	Total Byte Transferred	Total Duration	Count	%
P5	010300	010100	269.970(us)	546.845(us)	394.787(us)	0.8929	1.8086	1.2393	136.628(us)	366.020(us)	206.014(us)	2175488	001.724 052(ms)	4280	100.00
														Total: 4280	

FC Metrics, ELS and FCP

FC.ELS/AL

Port No	FC Source ID	FC Destination ID	Type	Count	%
P5	010300	fffffd	ELS_REPLY	1	5.56
P6	fffc01	010300	ELS_REQUEST	3	16.67
P6	fffffa	010300	ELS_REPLY	1	5.56
P6	010100	010300	ELS_REPLY	3	16.67
P6	fffc01	010300	ELS_REPLY	1	5.56
P5	010300	fffffa	ELS_REQUEST	1	5.56
P5	010300	fffc01	ELS_REQUEST	1	5.56
P5	010300	fffc01	ELS_REPLY	3	16.67
P5	010300	010100	ELS_REQUEST	3	16.67
P6	fffffd	010300	ELS_REQUEST	1	5.56
				Total: 18	

FC.ELS/AL.ELS Command

Port No	FC Source ID	FC Destination ID	Command	Response	Count	%
P5	010300	fffc01	RPSC	Accept	1	11.11
P6	fffffd	010300	RSCN	Accept	1	11.11
P6	fffc01	010300	PLOGI	Accept	1	11.11
P5	010300	010100	PRLI	Accept	1	11.11
P6	fffc01	010300	LOGO	Accept	1	11.11
P5	010300	010100	PLOGI	Accept	1	11.11
P6	fffc01	010300	PRLI	Accept	1	11.11
P5	010300	fffffa	PLOGI	Accept	1	11.11
P5	010300	010100	ADISC	Accept	1	11.11
					Total: 9	

FC.FCP Frames

Port No	FC Source ID	FC Destination ID	Frame Type	Count	%
P6	010100	010300	FCP_DATA	4257	33.32
P5	010300	010100	FCP_CMD	4290	33.58
P6	010100	010300	FCP_RSP	4228	33.10
				Total: 12775	

- Compare and contrast events across multiple links, by OX_ID and RX_ID pairs
- Tabular view of commands, ie PLOGI, indicating “normal” link functionality
- FCP request/response ratio indicating normative operation

Performance Metrics, Events/Utilization

FC.GS

Port No	FC Source ID	FC Destination ID	Type	Count	%
P6	fffffc	010300	FCCT_REPLY	5	22.73
P6	fffffa	010300	FCCT_REPLY	6	27.27
P5	010300	fffffa	FCCT_REQUEST	6	27.27
P5	010300	fffffc	FCCT_REQUEST	5	22.73
				Total: 22	

FC.GS.Command

Port No	FC Source ID	FC Destination ID	Type	Sub Type	Command	Response	Count	%
P5	010300	fffffc	Directory Service	Name Server	GID_PN	Accept	1	9.09
P5	010300	fffffa	Management Service	Fabric Device Management Interface	RPA	Accept	1	9.09
P5	010300	fffffc	Directory Service	Name Server	RSPN_ID	Accept	1	9.09
P5	010300	fffffc	Directory Service	Name Server	GID_FF	Accept	1	9.09
P5	010300	fffffc	Directory Service	Name Server	RFF_ID	Accept	1	9.09
P5	010300	fffffa	Management Service	Fabric Configuration Server	GFN	Accept	1	9.09
P5	010300	fffffa	Management Service	Fabric Configuration Server	GMAL	Reject	3	27.27
P5	010300	fffffc	Directory Service	Name Server	RFT_ID	Accept	1	9.09
P5	010300	fffffa	Management Service	Fabric Device Management Interface	RHBA	Accept	1	9.09
							Total: 11	

FC

Port No	FC Source ID	FC Destination ID	Type	Count	%
P5	010300	fffffa	ELS Frame	1	0.01
P6	fffc01	010300	ELS Frame	4	0.03
P5	010300	fffffa	GS Frame	6	0.05
P6	fffffd	010300	ELS Frame	1	0.01
P6	fffffa	010300	GS Frame	6	0.05
P6	fffffc	010300	GS Frame	5	0.04
P5	010300	010100	ELS Frame	3	0.02
P5	010300	fffc01	ELS Frame	4	0.03
P5	010300	fffffd	ELS Frame	1	0.01
P6	fffffa	010300	ELS Frame	1	0.01
P5	010300	fffffc	GS Frame	5	0.04
P6	010100	010300	ELS Frame	3	0.02
P6	010100	010300	FCP Frame	8485	66.21
P5	010300	010100	FCP Frame	4290	33.48
				Total: 12815	

Metric Groups (for deeper analysis)

- Error Counts
- Bandwidth vs Throughput
- Protocol overheads, Time on Link, IFG
  - *Host/target/network/idle time*
- Queue Depth vs Pending IOs
  - *Issued IOs*
  - *Completed IOs*
- MB/Sec, Command Latencies **
  - *Command to Command Time*
  - *First response time*
  - *Data Time*
  - *Data to Status Time*

**Credit analysis and command latencies are discussed in detail in
Protocol Analysis for High-Speed Fibre Channel Fabrics
https://www.brighttalk.com/webcast/14967/333863

Debug Example 1- Analyzer Placement

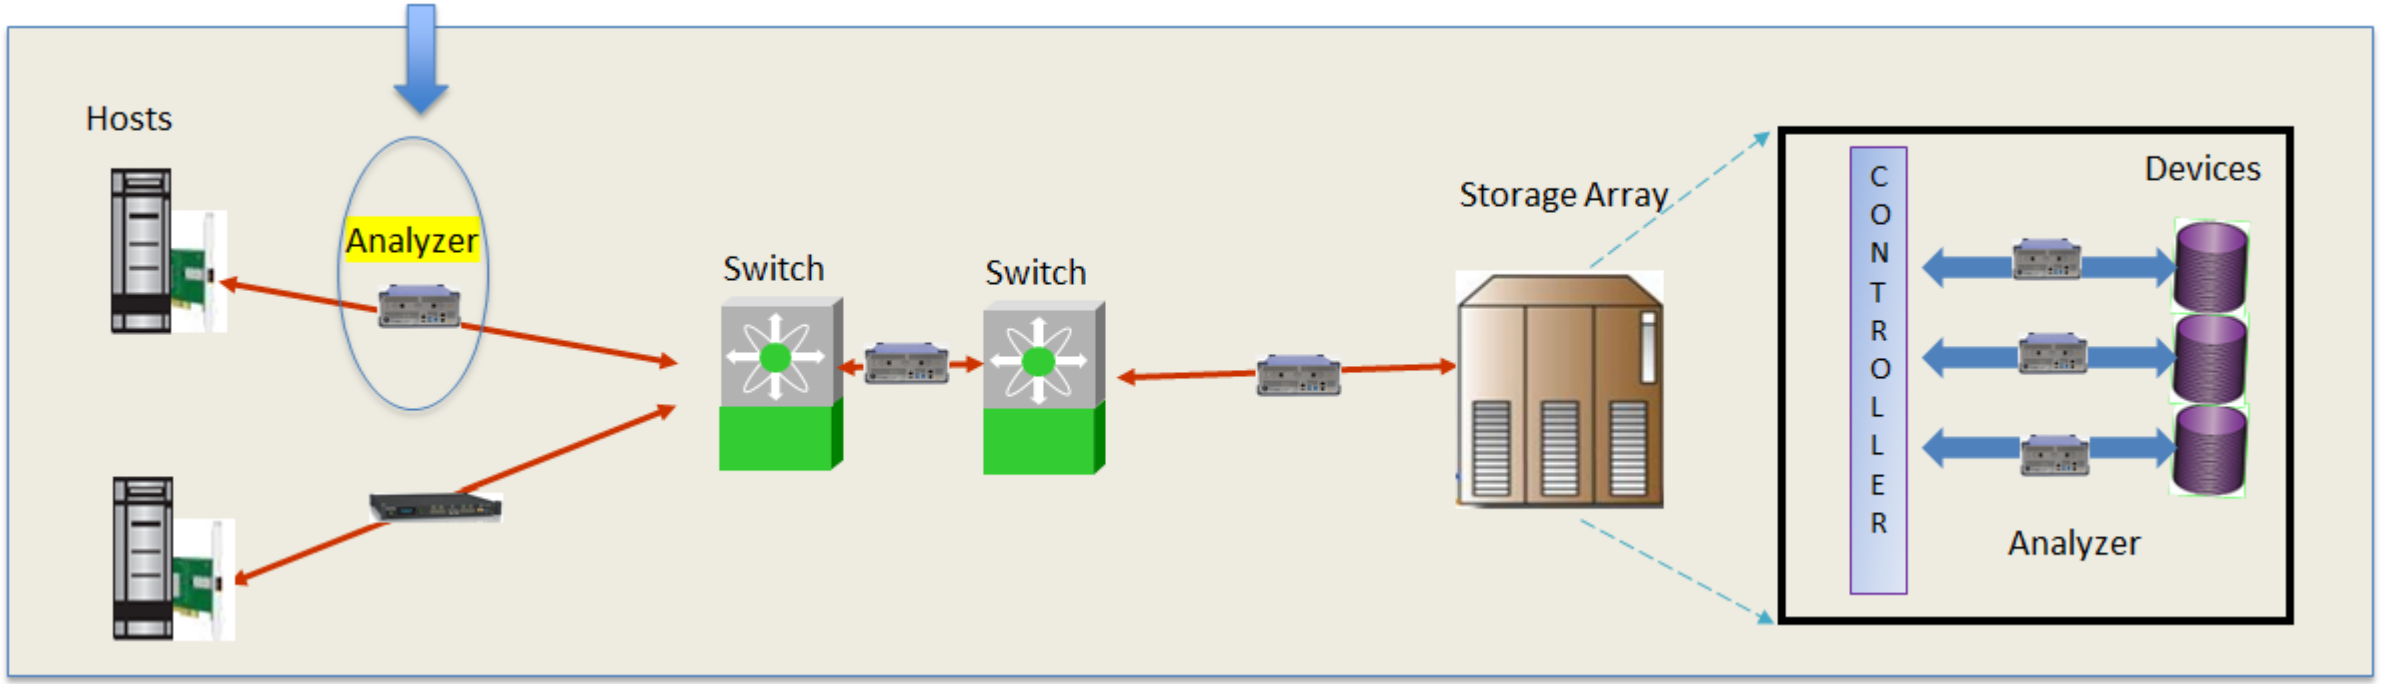

Debug Example 1 – Host Delay

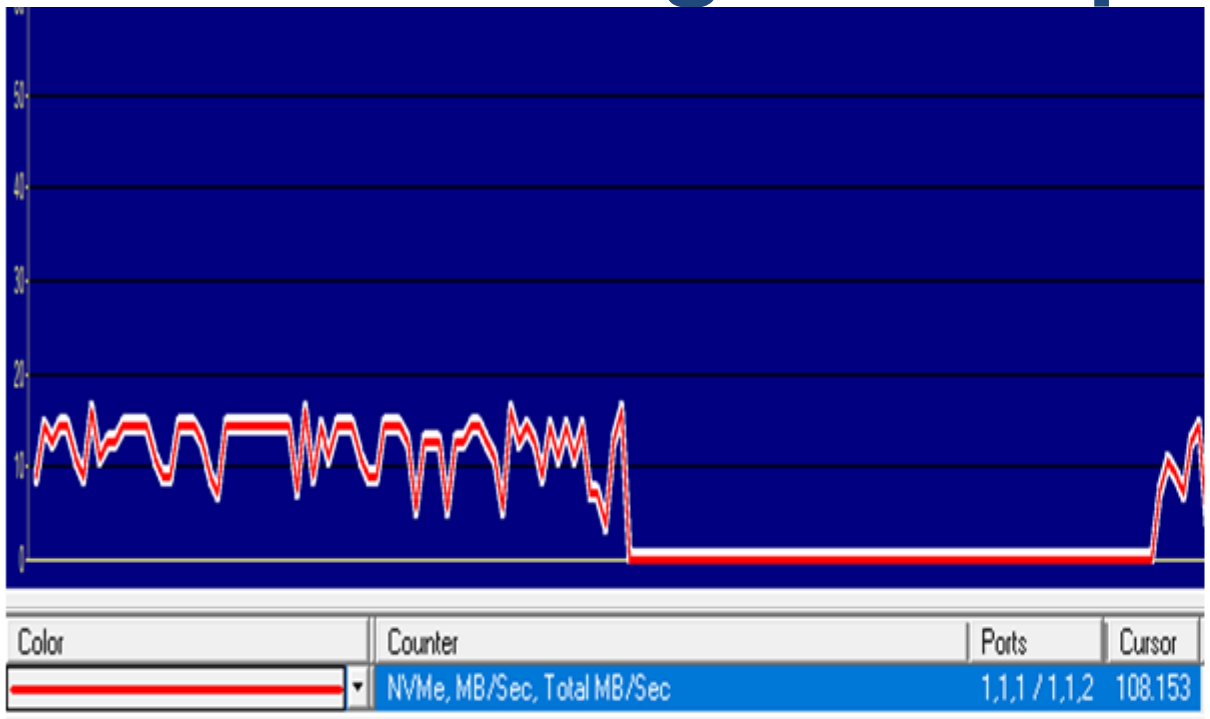

MB/Sec,

Pending IOs

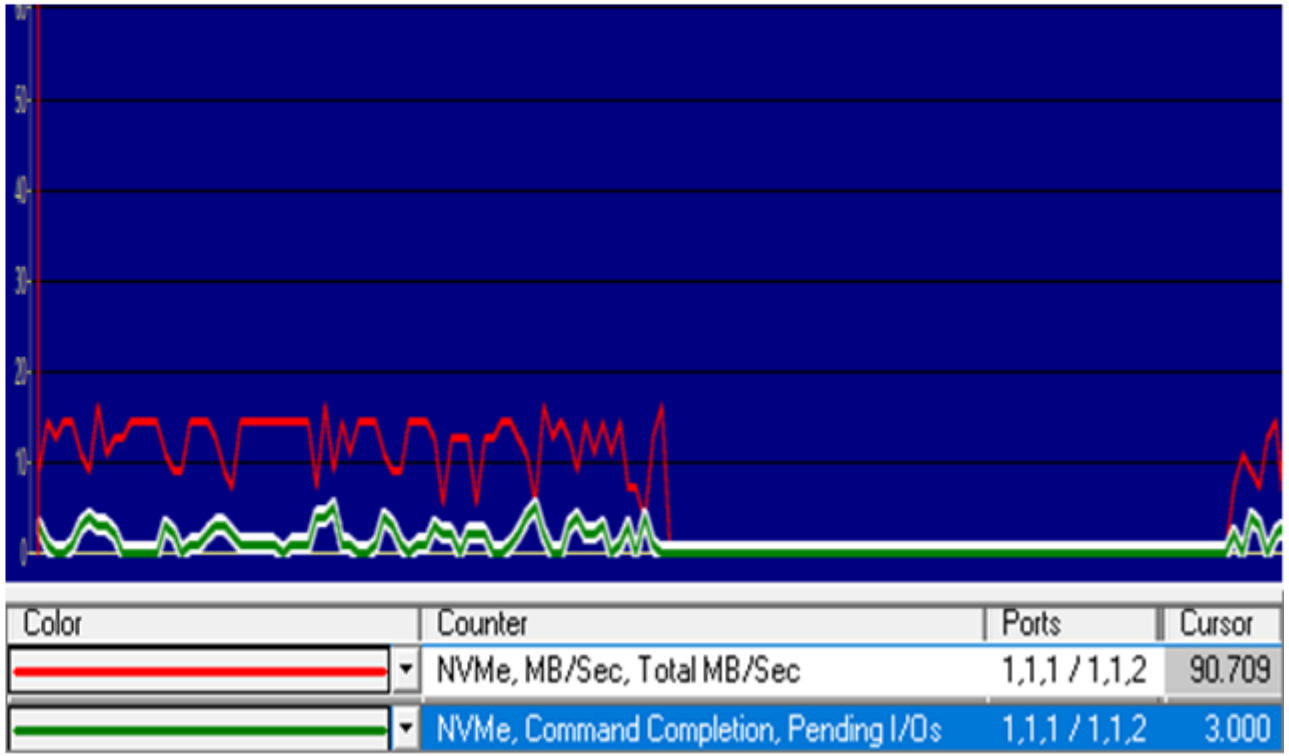

- Pending IOs graph dips with MB/Sec graph
- Indicates a host side delay.

Debug Example 1 – Host Delay contd.

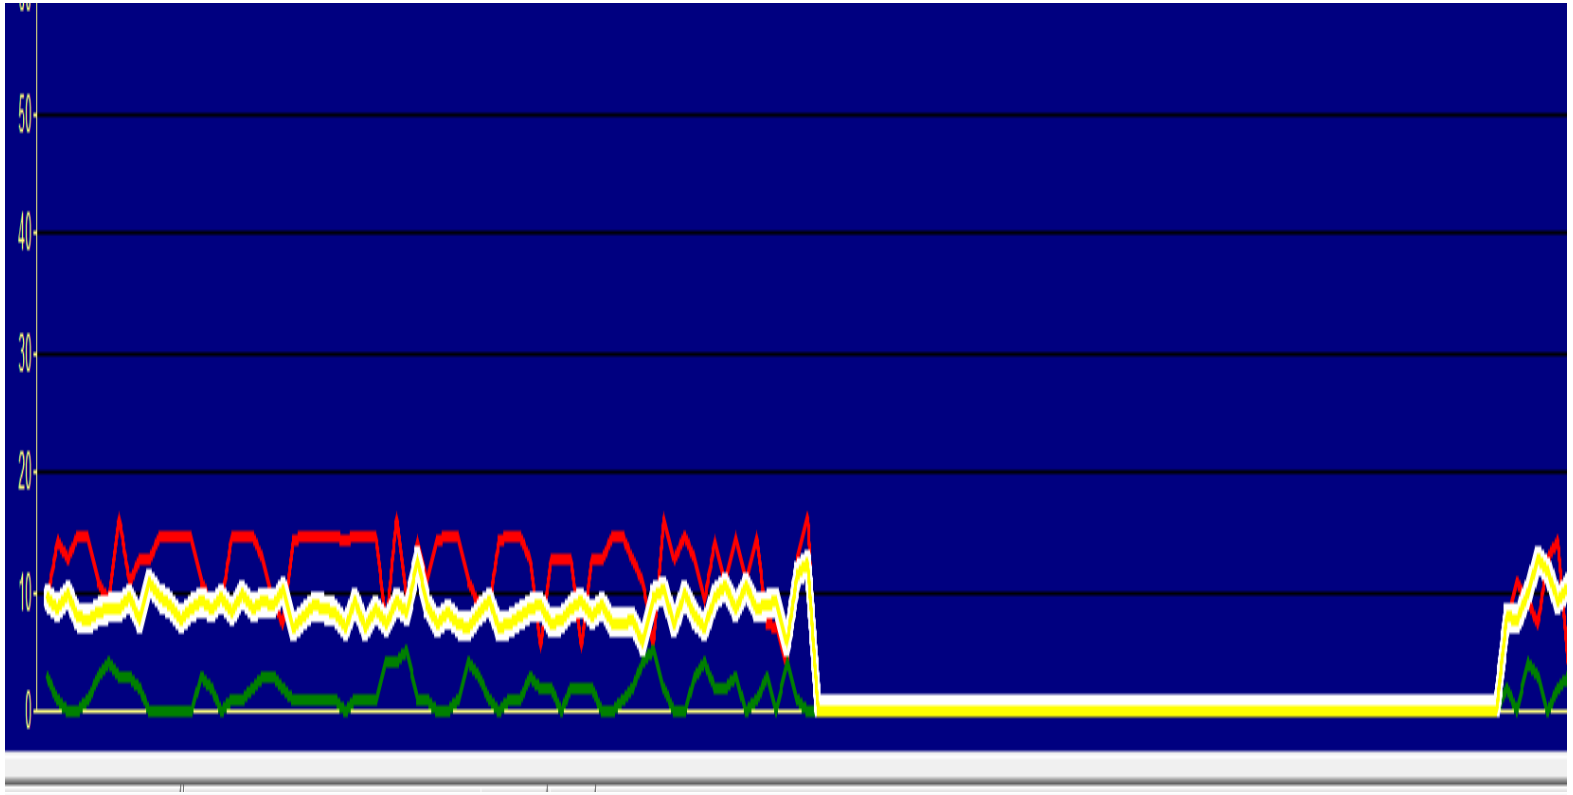

Color	Counter	Ports	Cursor
	NVMe, Command Completion, Pending I/Os	1,1,1 / 1,1,2	3.000
	NVMe, Exchanges (Any), Cmd to Completion Time (Max - ms)	1,1,1 / 1,1,2	0.125
	NVMe, MB/Sec, Total MB/Sec	1,1,1 / 1,1,2	71.953

Plot FC/NVMe exchange completion time to confirm.

Various metric groups can bring us to the same conclusion
Ex: Completed IOs vs Issued IOs

Debug Example 1 – Host Delay contd.

Event View

Icon	us.ns_ps	Delta Time	Port	Side A	Side B	Summary
32 FR	5771761659.050_0	0.307_6	FC Port(1,1,1)	FC4Status		Good Status
32 FR	5771761660.000_0	0.950_0	FC Port(1,1,1)	FC4Status		Good Status
32 FR	5771761663.324_4	3.324_4	FC Port(1,1,1)	FC4XRdy		DATA_RO =
32 FR	5771761671.696_4	8.372_0	FC Port(1,1,2)		FC4SData	FC4SData; S
32 FR	5771761672.315_2	0.618_8	FC Port(1,1,2)		FC4SData	FC4SData; S
32 FR	5771761676.112_0	3.796_8	FC Port(1,1,2)		FC4SData	FC4SData; S
32 FR	5771761676.731_5	0.619_5	FC Port(1,1,2)		FC4SData	FC4SData; S
32 FR	5771761704.028_0	27.296_5	FC Port(1,1,1)	FC4Status		Good Status;
32 FR	5771761719.368_0	15.340_0	FC Port(1,1,1)	FC4Status		Good Status
32 FR	5771776591.357_5	14871.989_5	FC Port(1,1,2)		FC4Cmd	Read; NSID
32 FR	5771776594.077_5	2.720_0	FC Port(1,1,2)		FC4Cmd	Write; NSID
32 FR	5771776638.464_0	44.386_5	FC Port(1,1,1)	FC4SData		FC4SData; S
32 FR	5771776639.086_0	0.622_0	FC Port(1,1,1)	FC4SData		FC4SData; S
32 FR	5771776639.706_0	0.620_0	FC Port(1,1,1)	FC4XRdy		DATA_RO =
32 FR	5771776650.466_4	10.760_4	FC Port(1,1,1)	FC4Status		Good Status;
32 FR	5771776653.462_0	2.995_6	FC Port(1,1,2)		FC4SData	FC4SData; S
32 FR	5771776654.082_0	0.620_0	FC		FC4SData	FC4SData; S
32 FR	5771776658.973_2	4.891_2	FC Port(1,1,2)		FC4Cmd	Write; NSID

FC Port(1,1,1)

Delta time is 14871.9895 us. Host stalls before issuing commands

Debug Example 2- Analyzer Placement

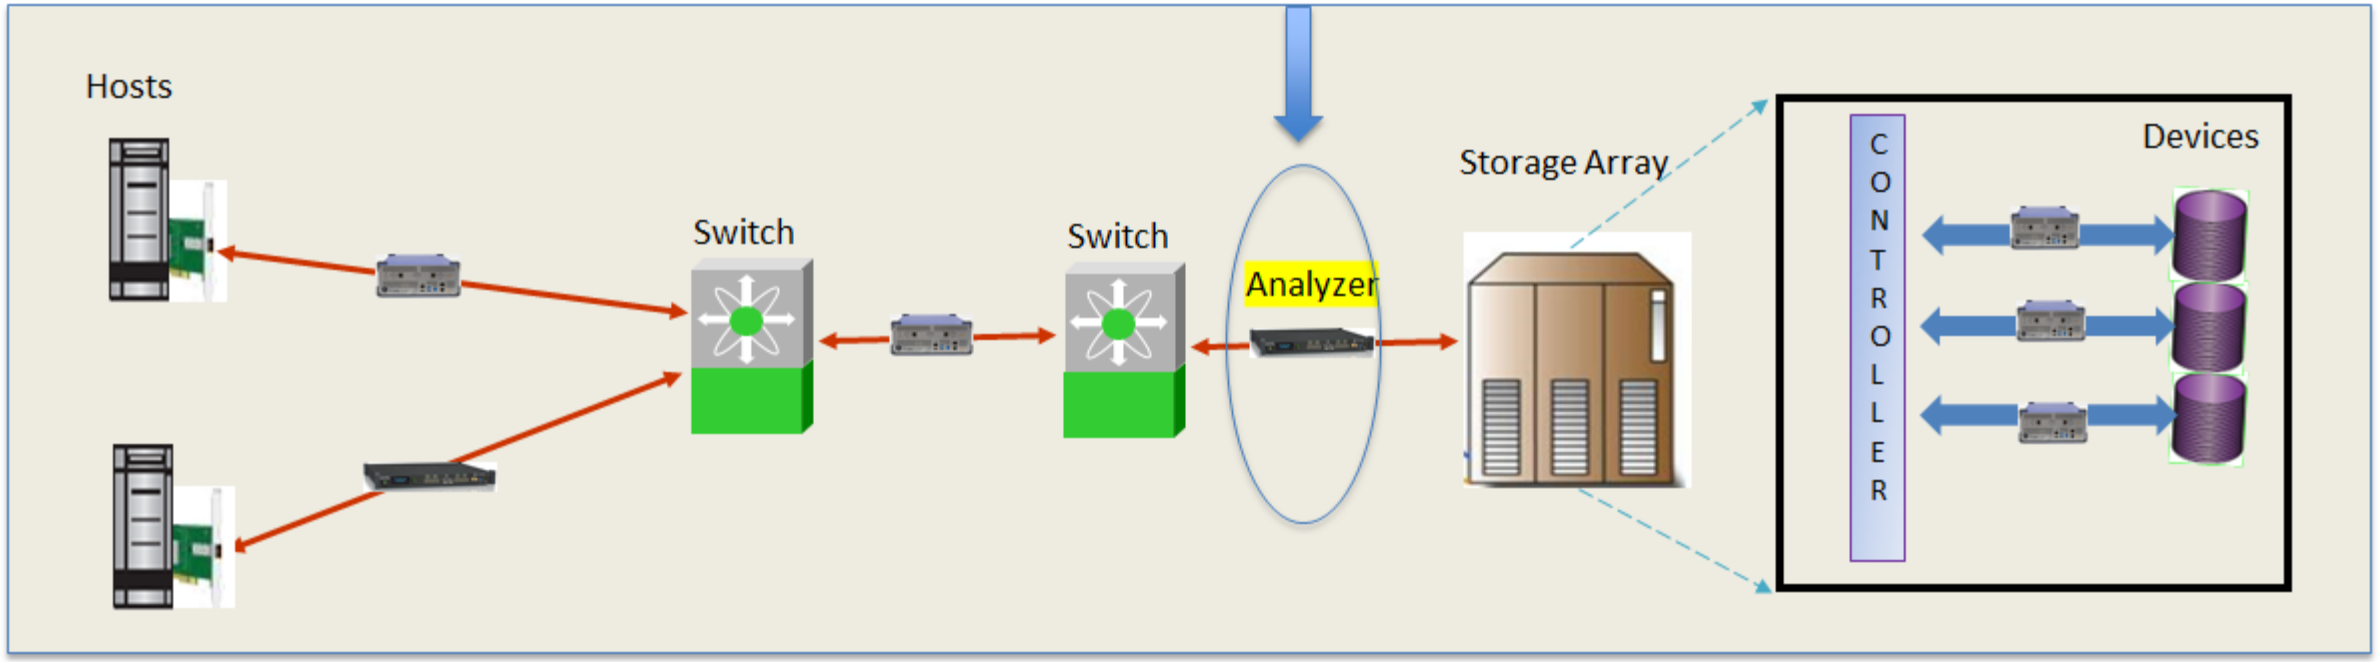

Debug Example 2- Controller Bug

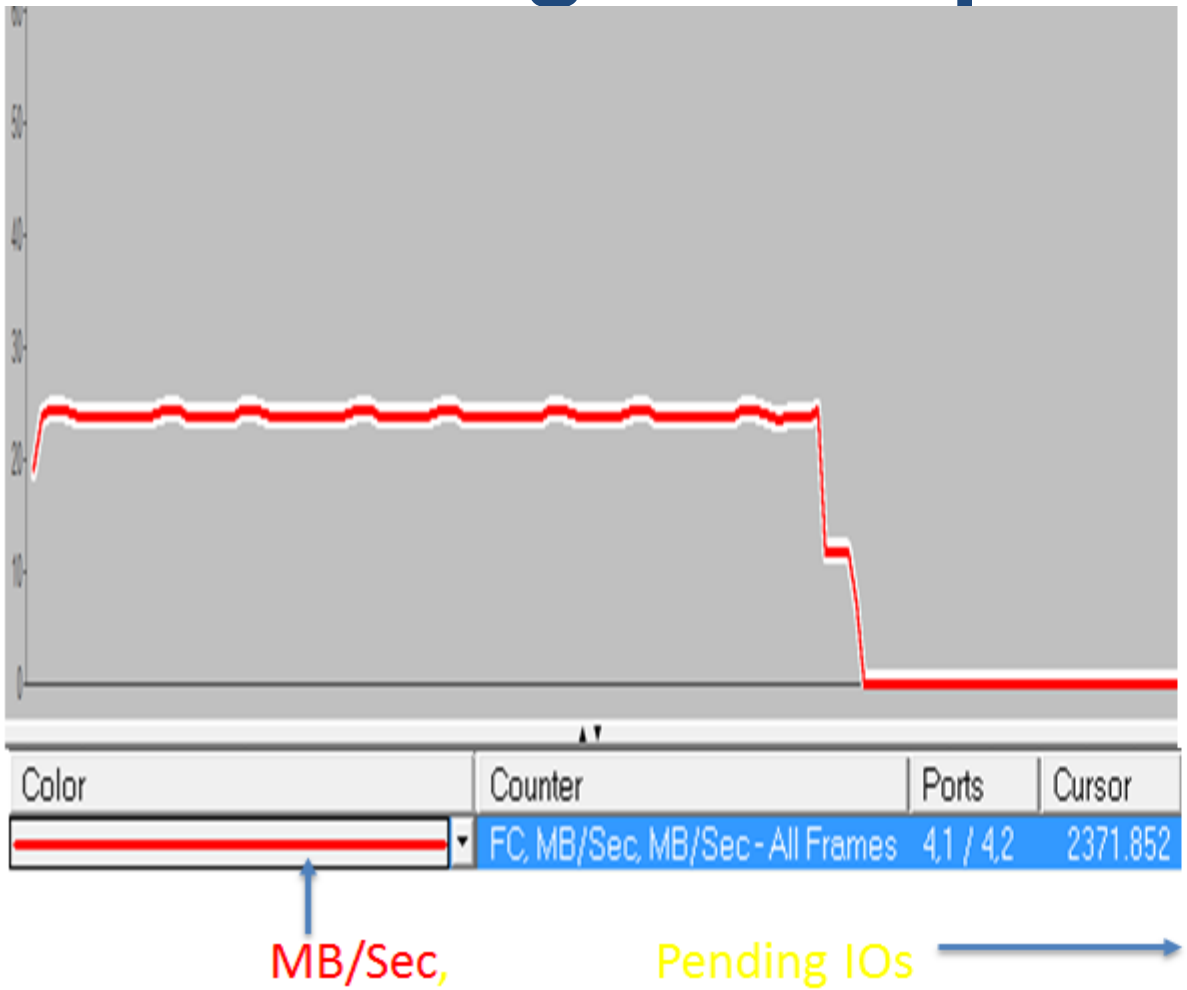

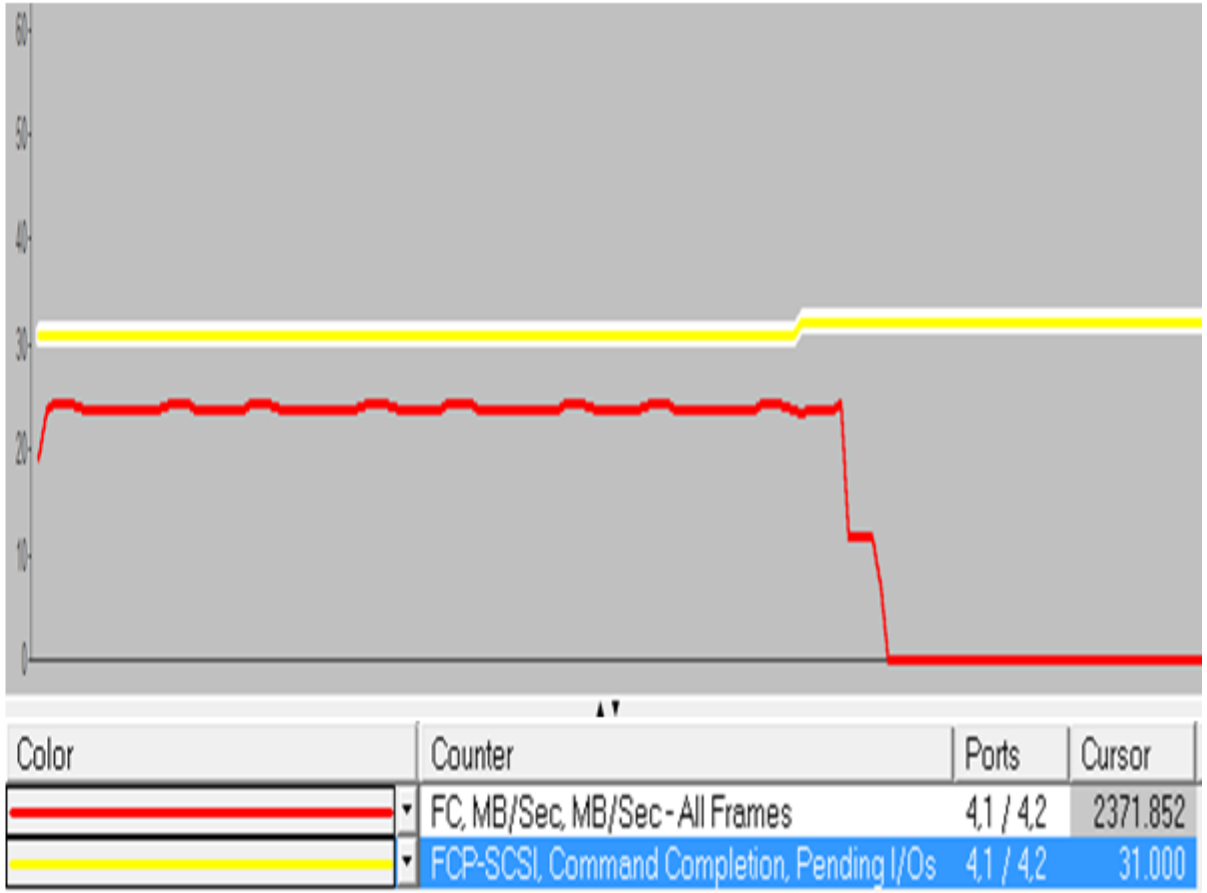

- Pending IOs increase when MB/Sec dips
- May indicate a controller side delay.

Debug Example 2- Controller Bug contd.

Exchanges View

Bo...	Side A	Side B	Age(us)	Pendi...	LBA	Tag	Source	Destination
⊞	Write(10)	Good	10642.188_5	1	3800	5C2	020100	010100
⊞	Read(10)	Pending	13791.615_2	1	1000	239	020000	010000
⊞	Read(10)	Pending	13766.591_8	2	1408	207	020000	010000
⊞	Read(10)	Pending	13382.434_6	3	800	25E	020000	010000
⊞	Read(10)	Pending	13344.095_7	4	C08	208	020000	010000
⊞	Write(10)	Pending	13305.773_4	5	4000	24C	020000	010000
⊞	Write(10)	Pending	13267.437_5	6	4408	22B	020000	010000
⊞	Read(10)	Pending	83.322_4	7	1800	1F0	020000	010000
⊞	Read(10)	Pending	10723.172_9	3	1000	630	020100	010100
⊞	Read(10)	Pending	10703.137_7	4	1408	597	020100	010100
⊞	Read(10)	Pending	10347.069_3	5	800	57E	020100	010100
⊞	Read(10)	Pending	10327.037_1	6	C08	5E6	020100	010100
⊞	Read(10)	Pending	9620.183_6	7	4000	5E3	020100	010100
⊞	Read(10)	Pending	9600.152_3	8	4408	58A	020100	010100
⊞	Write(10)	Pending	7609.912_6	9	2000	551	020100	010100
⊞	Write(10)	Pending	7521.600_6	10	2408	593	020100	010100
⊞	Write(10)	Pending	6164.569_8	11	2800	572	020100	010100
⊞	Write(10)	Pending	6107.956_5	12	2C08	595	020100	010100
⊞	Write(10)	Pending	4852.116_2	13	3000	5C7	020100	010100
⊞	Write(10)	Pending	4813.795_4	14	3408	5E4	020100	010100

Incomplete exchanges shown in the trace

Debug Example 3 – Analyzer Placement

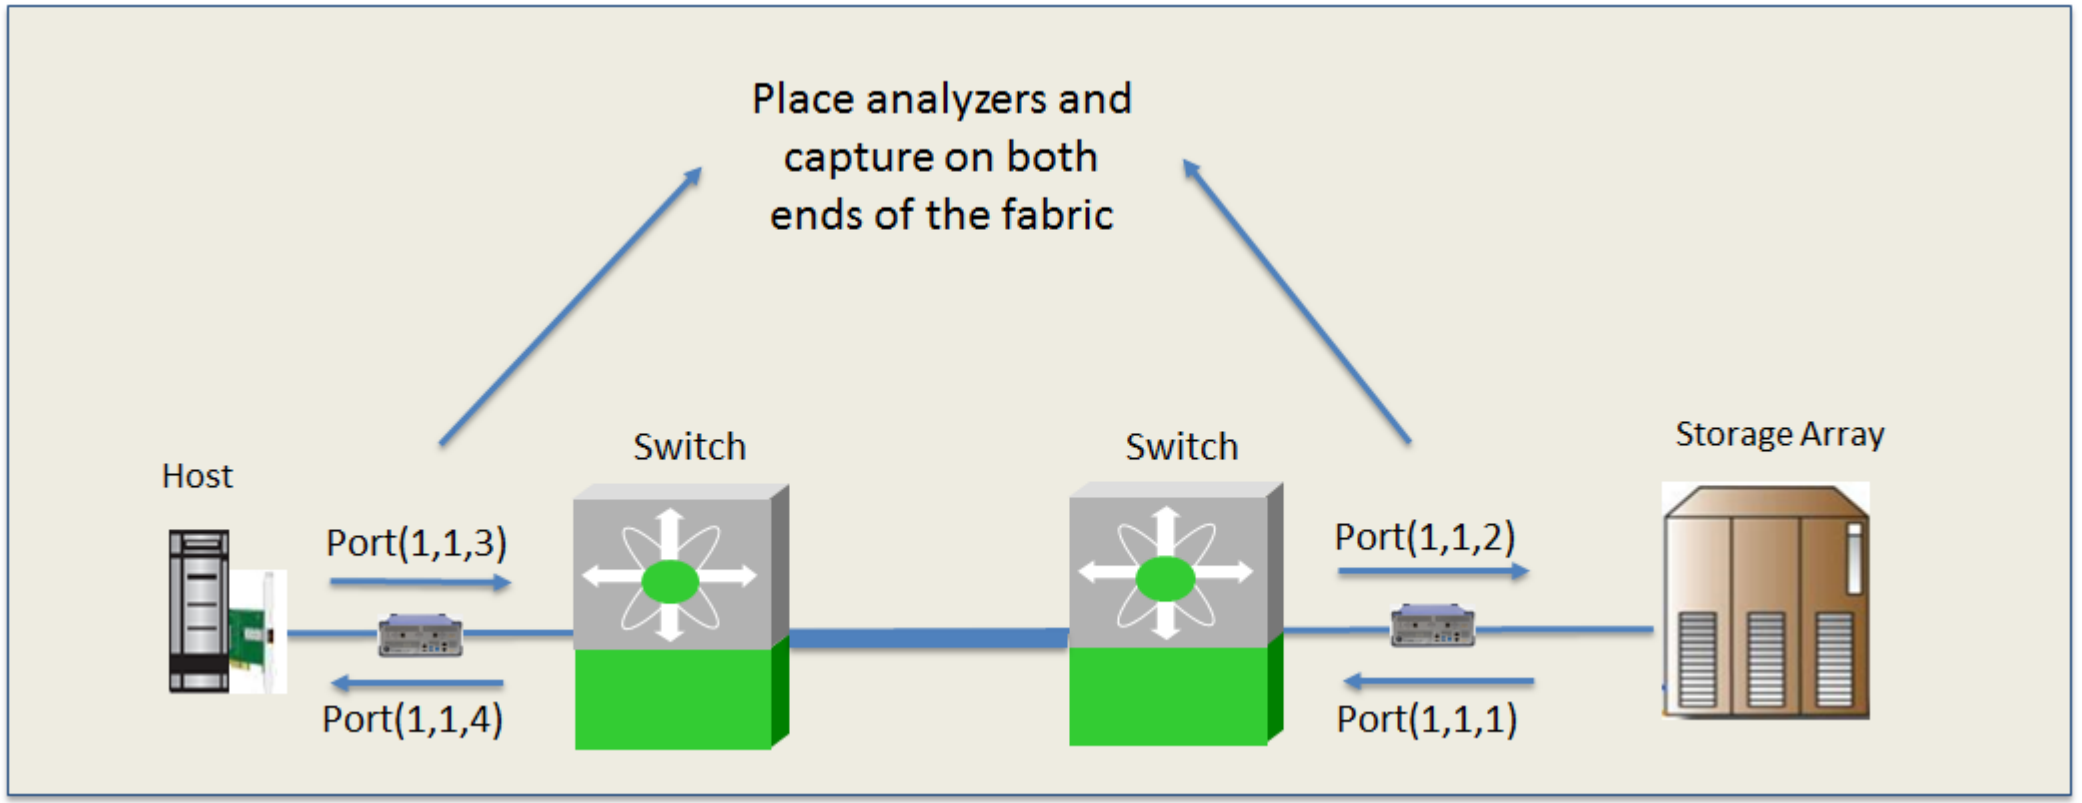

Debug Example 3 - Dropped Frames

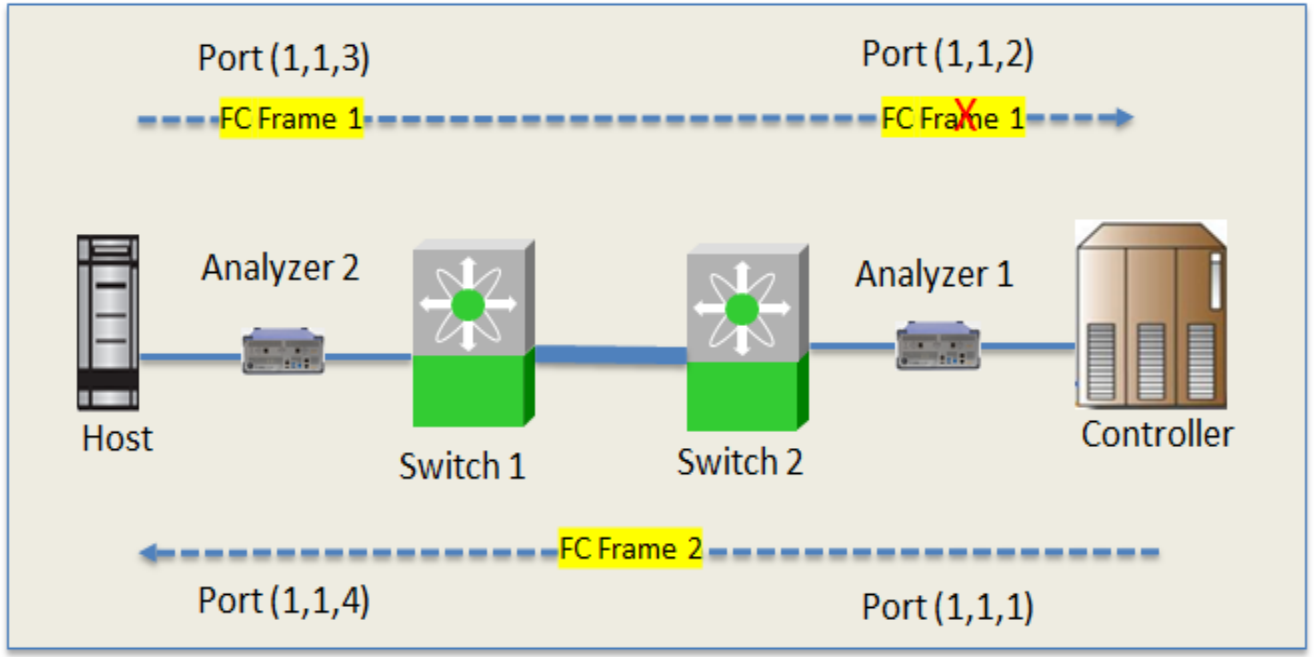

	Description	Timestamp	Source	Destination	Port Name	Type	Total Count	Value1	Value2
1	Out of Order Frame by Src/Dst	122001789.613_5	FC Port(1,1,3)	FC Port(1,1,2)	1,1,1 / 1,1,2	CrossPort	3	Time Detected 122001800.506_3	Delta Time 10.8927 (us)

Debug Example 3 - Frame Tracing

32 FR		122001787.915_2	Analyzer 1		FC4SData	00005000	010D00	C80001	B711A419
32 BM	Frame 1	122001787.943_8	Analyzer 2	FC4SData		00008000	010D00	C80001	EECD7557
32 BM	Frame 2	122001788.563_5	Analyzer 2	FC4SData		00008800	010D00	C80001	D7682A1E
32 BM	Frame 3	122001789.613_5	Analyzer 2	FC4SData		00009000	010D00	C80001	B17A710F
32 FR		122001789.887_2	Analyzer 1		FC4SData	00005800	010D00	C80001	B994FF95
32 BM	Frame 4	122001790.235_5	Analyzer 2	FC4SData		00009800	010D00	C80001	91E39D02
32 FR		122001790.505_5	Analyzer 1		FC4SData	00006000	010D00	C80001	56C8A9F1
32 BM	Frame 5	122001791.317_5	Analyzer 2	FC4SData		0000A000	010D00	C80001	50A37CE7
32 FR		122001791.883_2	Analyzer 1		FC4SData	00006800	010D00	C80001	730927B0
32 BM	Frame 6	122001791.939_5	Analyzer 2	FC4SData		0000A800	010D00	C80001	5B7E4527
32 FR		122001792.501_5	Analyzer 1		FC4SData	00007000	010D00	C80001	097FADA9
32 FR		122001792.860_4	Analyzer 2	FC4SData		0000B000	010D00	C80001	0F1478BF
32 FR		122001793.480_4	Analyzer 2	FC4SData		0000B800	010D00	C80001	1DF5F23B
32 FR		122001794.025_5	Analyzer 1		FC4SData	00007800	010D00	C80001	358290AC
32 FR		122001794.535_5	Analyzer 2	FC4SData		0000C000	010D00	C80001	D31616EC
32 FR		122001795.157_5	Analyzer 2	FC4SData		0000C800	010D00	C80001	CF45F46C
32 FR		122001796.249_5	Analyzer 2	FC4SData		0000D000	010D00	C80001	8CA112B4
32 BM	Frame out 1	122001796.329_2	Analyzer 1		FC4SData	00008000	010D00	C80001	EECD7557
32 FR		122001796.869_5	Analyzer 2	FC4SData		0000D800	010D00	C80001	89CE4370
32 BM	Frame out 2	122001796.947_5	Analyzer 1		FC4SData	00008800	010D00	C80001	D7682A1E
32 FR		122001797.923_5	Analyzer 2	FC4SData		0000E000	010D00	C80001	6D781F5C
32 FR		122001798.545_5	Analyzer 2	FC4SData		0000E800	010D00	C80001	43539B55
32 FR		122001799.598_4	Analyzer 2	FC4SData		0000F000	010D00	C80001	32CF1B04
32 FR		122001800.220_4	Analyzer 2	FC4SData		0000F800	010D00	C80001	8C121770
32 BM	Frame out 6	122001800.506_3	Analyzer 1		FC4SData	0000A800	010D00	C80001	5B7E4527
32 FR		122001802.501_5	Analyzer 1		FC4SData	0000B000	010D00	C80001	0F1478BF

Frame 1 seen on Analyzer 2 and then on Analyzer 1.

Debug Example 3 - Frame Tracing

32 FR		122001787.915_2	Analyzer 1		FC4SData	00005000	010D00	C80001	B711A419
32 BM	Frame 1	122001787.943_8	Analyzer 2	FC4SData		00008000	010D00	C80001	EECD7557
32 BM	Frame 2	122001788.563_5	Analyzer 2	FC4SData		00008800	010D00	C80001	D7682A1E
32 BM	Frame 3	122001789.613_5	Analyzer 2	FC4SData		00009000	010D00	C80001	B17A710F
32 FR		122001789.887_2	Analyzer 1		FC4SData	00005800	010D00	C80001	B994FF95
32 BM	Frame 4	122001790.235_5	Analyzer 2	FC4SData		00009800	010D00	C80001	91E39D02
32 FR		122001790.505_5	Analyzer 1		FC4SData	00006000	010D00	C80001	56C8A9F1
32 BM	Frame 5	122001791.317_5	Analyzer 2	FC4SData		0000A000	010D00	C80001	50A37CE7
32 FR		122001791.883_2	Analyzer 1		FC4SData	00006800	010D00	C80001	730927B0
32 BM	Frame 6	122001791.939_5	Analyzer 2	FC4SData		0000A800	010D00	C80001	5B7E4527
32 FR		122001792.501_5	Analyzer 1		FC4SData	00007000	010D00	C80001	097FADA9
32 FR		122001792.860_4	Analyzer 2	FC4SData		0000B000	010D00	C80001	0F1478BF
32 FR		122001793.480_4	Analyzer 2	FC4SData		0000B800	010D00	C80001	1DF5F23B
32 FR		122001794.025_5	Analyzer 1		FC4SData	00007800	010D00	C80001	358290AC
32 FR		122001794.535_5	Analyzer 2	FC4SData		0000C000	010D00	C80001	D31616EC
32 FR		122001795.157_5	Analyzer 2	FC4SData		0000C800	010D00	C80001	CF45F46C
32 FR		122001796.249_5	Analyzer 2	FC4SData		0000D000	010D00	C80001	8CA112B4
32 BM	Frame out 1	122001796.329_2	Analyzer 1		FC4SData	00008000	010D00	C80001	EECD7557
32 FR		122001796.869_5	Analyzer 2	FC4SData		0000D800	010D00	C80001	89CE4370
32 BM	Frame out 2	122001796.947_5	Analyzer 1		FC4SData	00008800	010D00	C80001	D7682A1E
32 FR		122001797.923_5	Analyzer 2	FC4SData		0000E000	010D00	C80001	6D781F5C
32 FR		122001798.545_5	Analyzer 2	FC4SData		0000E800	010D00	C80001	43539B55
32 FR		122001799.598_4	Analyzer 2	FC4SData		0000F000	010D00	C80001	32CF1B04
32 FR		122001800.220_4	Analyzer 2	FC4SData		0000F800	010D00	C80001	8C121770
32 BM	Frame out 6	122001800.506_3	Analyzer 1		FC4SData	0000A800	010D00	C80001	5B7E4527
32 FR		122001802.501_5	Analyzer 1		FC4SData	0000B000	010D00	C80001	0F1478BF

Frames 3, 4, 5 are dropped enroute and cannot be seen on Analyzer 1.

Frame 1, 2 and 6 seen on Analyzer 2 and then on Analyzer 1.

Benefits of Purposeful Error Injection and Packet Modification in FC0

- **Create specific "errors" to test fabrics and their capacity to recover from failure scenarios, examples**;
  - Emulate Cable Pull testing (force disconnect/reconnect)
  - Alter Speed-Negotiation advertisement, add Transmitter Training Errors
  - Alter Primitives
    - Ex. change SoF1 to EOF
  - Introduce latency
- **Stress Test Fabric-Link Participants**
- **Replicate Customer environments/issues**
- **Optimize Error recovery algorithms**

Benefits of Purposeful Error Injection in Transmitter Training Routines

- **Test fabrics participants and their capacity to recover from failure scenarios, examples**;
  - Alter Speed-Negotiation advertisement, add Transmitter Training Errors
  - Introduce latency in the linkup exchange
- **Modify Transmitter Training Control and Status Field exchanges to examine link interoperation**
  - **TT sequence has a greater impact with the increase of fabric speeds**
    - **Modify Control Field values**
      - **Change EXTENDED MARKER value**
      - **Turn "Off" FECReq**
      - **Turn "On" Parallel Lane Support**
    - **Modify Status Field values**
      - **Change FECCap value**
      - **Spoof Coefficient status**
  - **Modify values to unexpected or "bad" values**
- **Establish worse case limitations**
- **Emulate link partner characteristics**

Field	Value
Training Sequence	0x8A022A80
Control Field	0x8A02
Extended Marker	0x2
Preset	0x0
Initialize	0x0
FECReq	0x1
Parallel Lane Support	0x0
C1Upd	0x0
C0Upd	0x0
C-1Upd	0x2
Status Field	0x2A80
TC	0x0
SN	0x0
FECCap	0x1
TF	0x0
C1Stat	0x0
C0Stat	0x0
C-1Stat	0x0

Benefits of Purposeful Error Injection and FEC Effects

- **Testing FEC Limitations**
  - FEC is mandatory in PAM4 fabrics
  - How do the link participants respond to Correctable Errors
  - How do the link participants respond to Uncorrectable Errors
- **Insertion of FEC decay conditions errors**
- **Characterize the capabilities of the fabric under examination!**

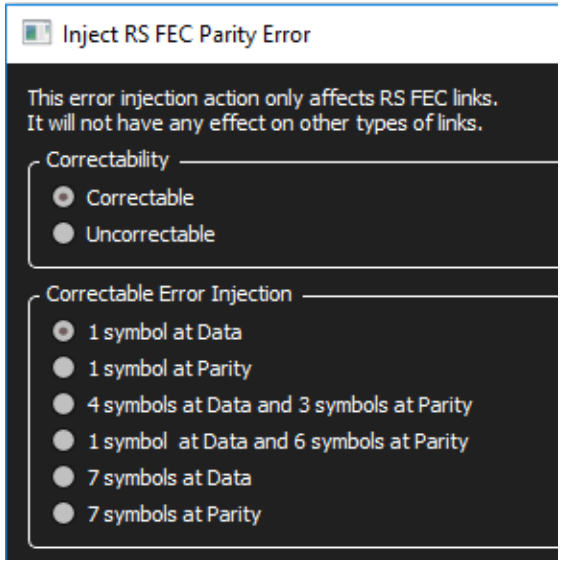

FC-1, Link Layer

- **Supports the Specifics for Link Control including;**
  - Ordered Sets
    - Frame Delimiters
      - SOF, EOF
    - Primitive Signaling
      - Idles, R_RDY, ARB
    - Primitive Sequences
      - LR, LRR, NOS, OLS
  - Encoding/Decoding
  - Transmitter Training
- **Characterize the transactions**
  - Timing, Control Messaging, Verification

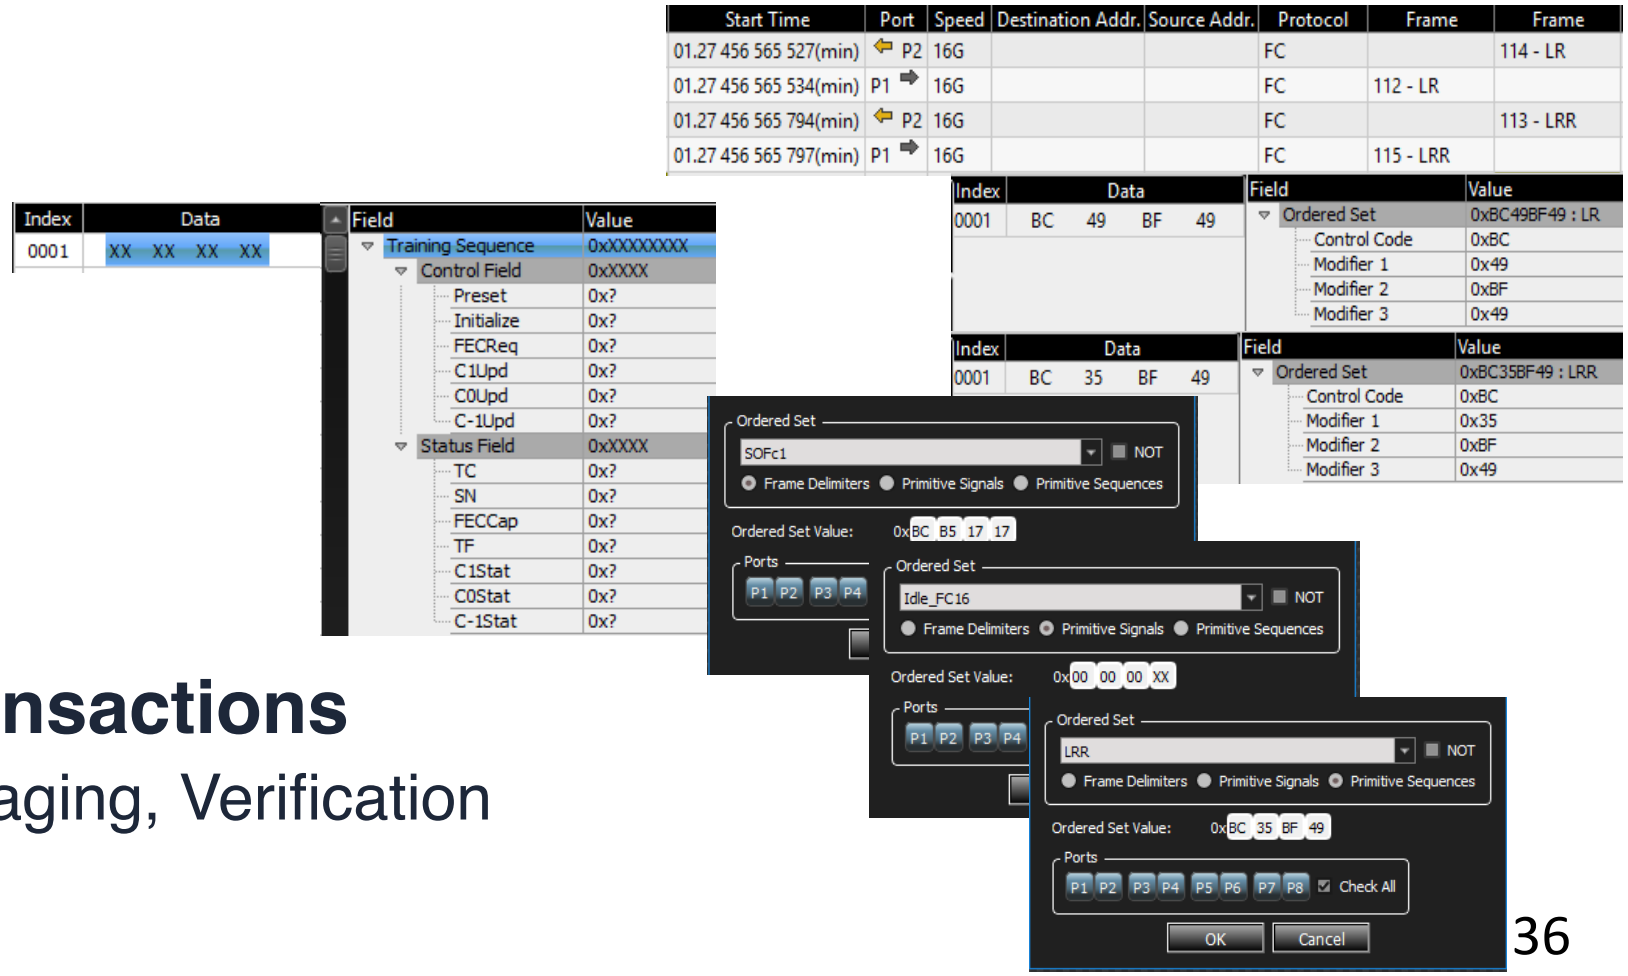

Benefits of Purposeful Error Injection and Packet Modification in FC1

- **Create specific TX/RX conditions to test link recovery**
  - Alter Primitives
    - Change SOF to EOF
    - Force ABTS
  - Force Link Resets
    - Force link failure to test link failover
    - Change R_RDY to IDLE
    - Change LRR to IDLE
- Test/Optimize Error recovery algorithms

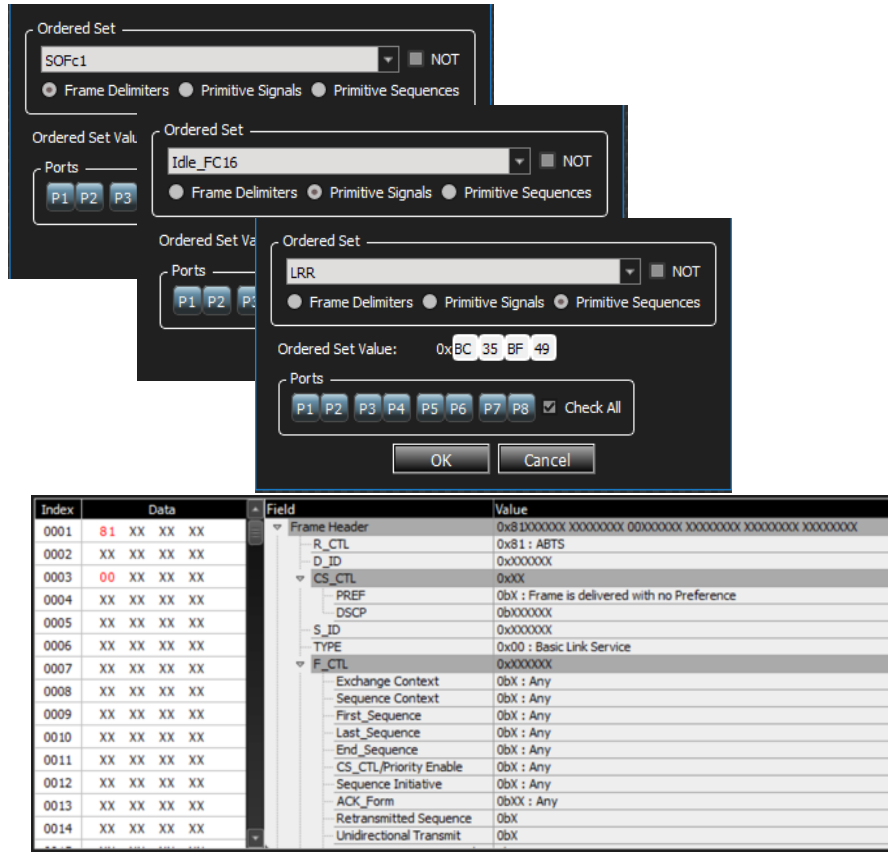

Errors in FC2 Framing Protocol

- **Create specific Framing/Exchange errors to test framing**
  - Alter Class of Service parameters
  - Change OX_ID or RX_ID information
  - Modify SEQ_ID count
  - Change Frame/Data/Link control values
  - Introduce ABTS
  - Obfuscate Buffer to Buffer Credits
- **May be required of FW and Driver Engineering teams**

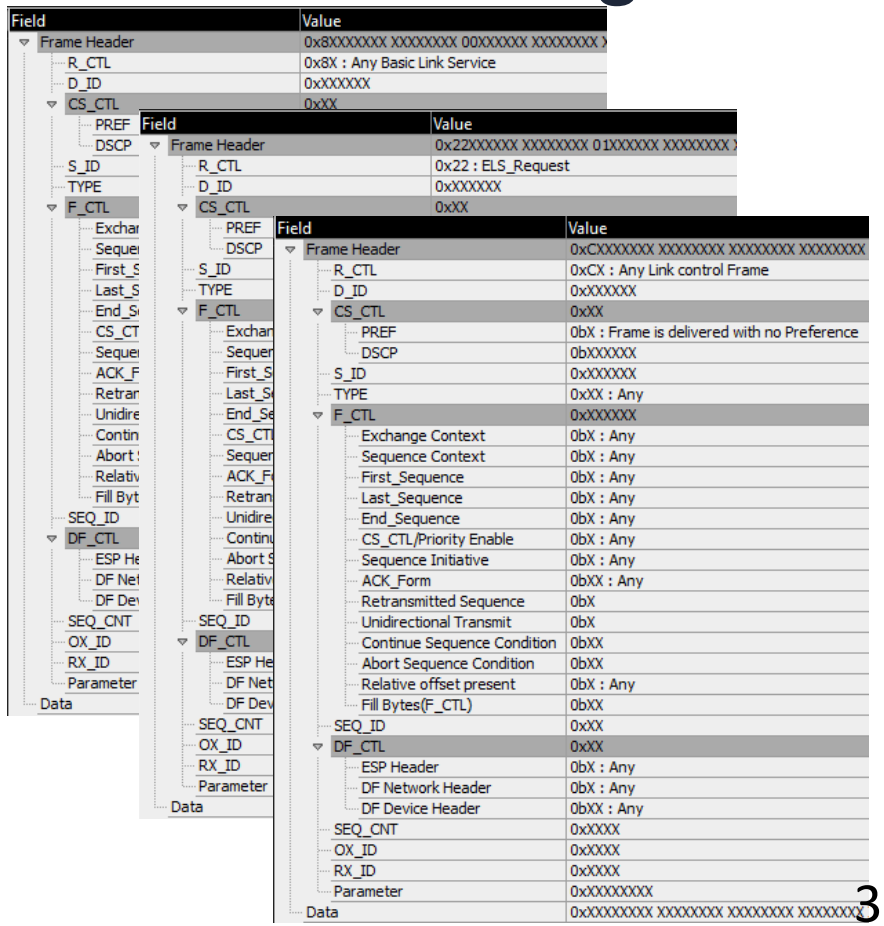

Example #1, BBC Jamming

- Emulating Credit Starvation by modifying BBC response, turn R_RDY into IDLE or OLS – Continuous, One Time, Variable
  - Test the effects of maximum Link budget and latency characteristics

09.267 741 340(s)	P5 →	16G	010100	010300	FC	0x0009	0xFFFF	FCP_CMD	
09.268 014 886(s)	← P6	16G	010300	010100	FC	0x0009	0x0C4C		FCP_XFER_RDY
09.268 017 307(s)	P5 →	16G	010100	010300	FC	0x0009	0x0C4C	FCP_DATA	
09.268 253 763(s)	← P6	16G	010300	010100	FC	0x0009	0x0C4C		FCP_RSP
09.267 741 985(s)	← P6	16G			FC				R_RDY

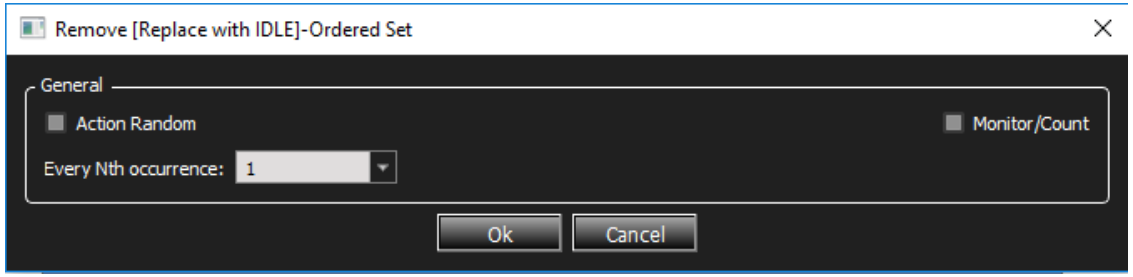

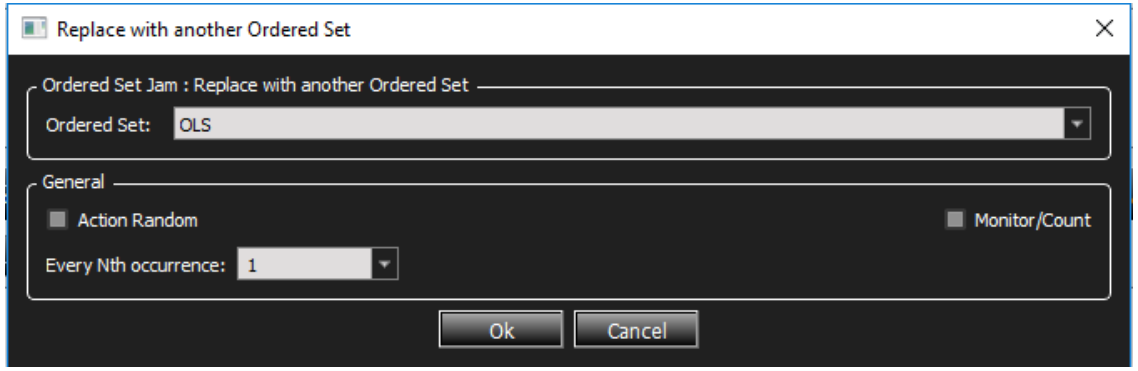

Extended Link Services

- **Common ELS Commands and Response**
  - PLOGI
  - FLOGI
  - PRLI
  - PRLO
  - LS_ACC
  - LS_RJT

Field	Value
ELS Command	0x03 : PLOGI
Common Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Buffer-to-Buffer Credit	0xXXXX
Common Features	0xXXXX
BB_SC_N	0xX
Buffer-to-Buffer Receive Data...	0xXXX
Nx_Port Total Concurrent Seq...	0xXX
Relative Offset By Info Categ...	0xXXXX
E_D_TOV Value	0xXXXXXXXX
Port_Name	0xXXXXXXXX XXXXXXXX
Node_ Or Fabric_Name	0xXXXXXXXX XXXXXXXX
Class 1 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Class 2 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Class 3 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Service Options	0xXXXX
Class Validity	0bX : Any
Priority/Preemption	0bX : Any
Preference	0bX : Any
DiffServ QoS	0bX : Any
Initiator Control	0bXXXXXXXX XXXXXXXX
Initial Responder Process_...	0x? : Any
Clock Synchronization ELS ...	0bX : Any
Recipient Control	0xXXXX
E_D_TOV Resolution	0bXX : Any
Categories Per Sequence	0bXX : Any
Clock Synchronization ELS ...	0bX : Any
Receive Data_Field Size	0xXXX
Concurrent Sequences	0xXX
Open Sequences Per Exchange	0xXX
Vendor Version Level	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX

Field	Value
ELS Command	0x20 : PRLI
Page Len.	0xXX
Payload Len.	0xXXXX
Service Parameter Page	0xXXXXXXXX XXXXXXXX
Type Code Or Common Servic...	0xXX : Any
Type Code Extension	0xXX
Originator Process_Associator ...	0bX
Responder Process_Associato...	0bX
Establish Image Pair	0bX
Originator Process_Associator	0xXXXXXXXX
Responder Process_Associator	0xXXXXXXXX
Service Parameters	0xXXXXXXXX XXXXXXXX

Field	0x04 : FLOGI
ELS Command	
Common Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Buffer-to-Buffer Credit	0xXXXX
Common Features	0xXXXX
Multiple N_Port_ID Support	0bX : Any
Virtual Fabrics Bit	0bX : Any
Valid Vendor Version Level	0bX : Any
N_Port/F_Port	0bX : Any
BB_Credit Management	0bX : Any
Name Server Session Begin	0bX
Energy Efficient LPI Mode ...	0x?
Priority Tagging Supported	0x?
Query Data Buffer Conditi...	0bX : Any
Security Bit	0bX : Any
Clock Synchronization Prim...	0bX : Any
R_T_TOV Value	0bX : Any
Dynamic Half Duplex Supp...	0bX : Any
Valid Vendor Version Level	0bX : Any
Payload Bit	0bX : Any
BB_SC_N	0xX
Buffer-to-Buffer Receive Data...	0xXXX
Port_Name	0xXXXXXXXX XXXXXXXX
Node_ Or Fabric_Name	0xXXXXXXXX XXXXXXXX
Class 1 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Class 2 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Class 3 Service Parameters	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX
Service Options	0xXXXX
Class Validity	0bX : Any
Sequential Delivery	0bX : Any
Priority/Preemption	0bX : Any
Preference	0bX : Any
DiffServ QoS	0bX : Any
Initiator Control	0bXXXXXXXX XXXXXXXX
Clock Synchronization ELS ...	0bX : Any
Recipient Control	0xXXXX
Clock Synchronization ELS ...	0bX : Any
Vendor Version Level	0xXXXXXXXX XXXXXXXX XXXXXXXX XXXXXXXX

Field	Value
ELS Command	0x21 : PRLO
Payload Len.	0xXXXX
Logout Parameter Page	0xXXXXXXXX XXXXXXXX
Type Code Or Common Logout...	0xXX
Type Code Extension	0xXX
Originator Process_Associator ...	0bX
Responder Process_Associato...	0bX
Originator Process_Associator	0xXXXXXXXX
Responder Process_Associator	0xXXXXXXXX
Logout Service Parameters	0xXXXXXXXX

Example #2, ELS Jamming

- Port Login on **S_ID 010300**
  Modify ELS Request to reduce initially reported Buffer to Buffer Credits – Continuous jam

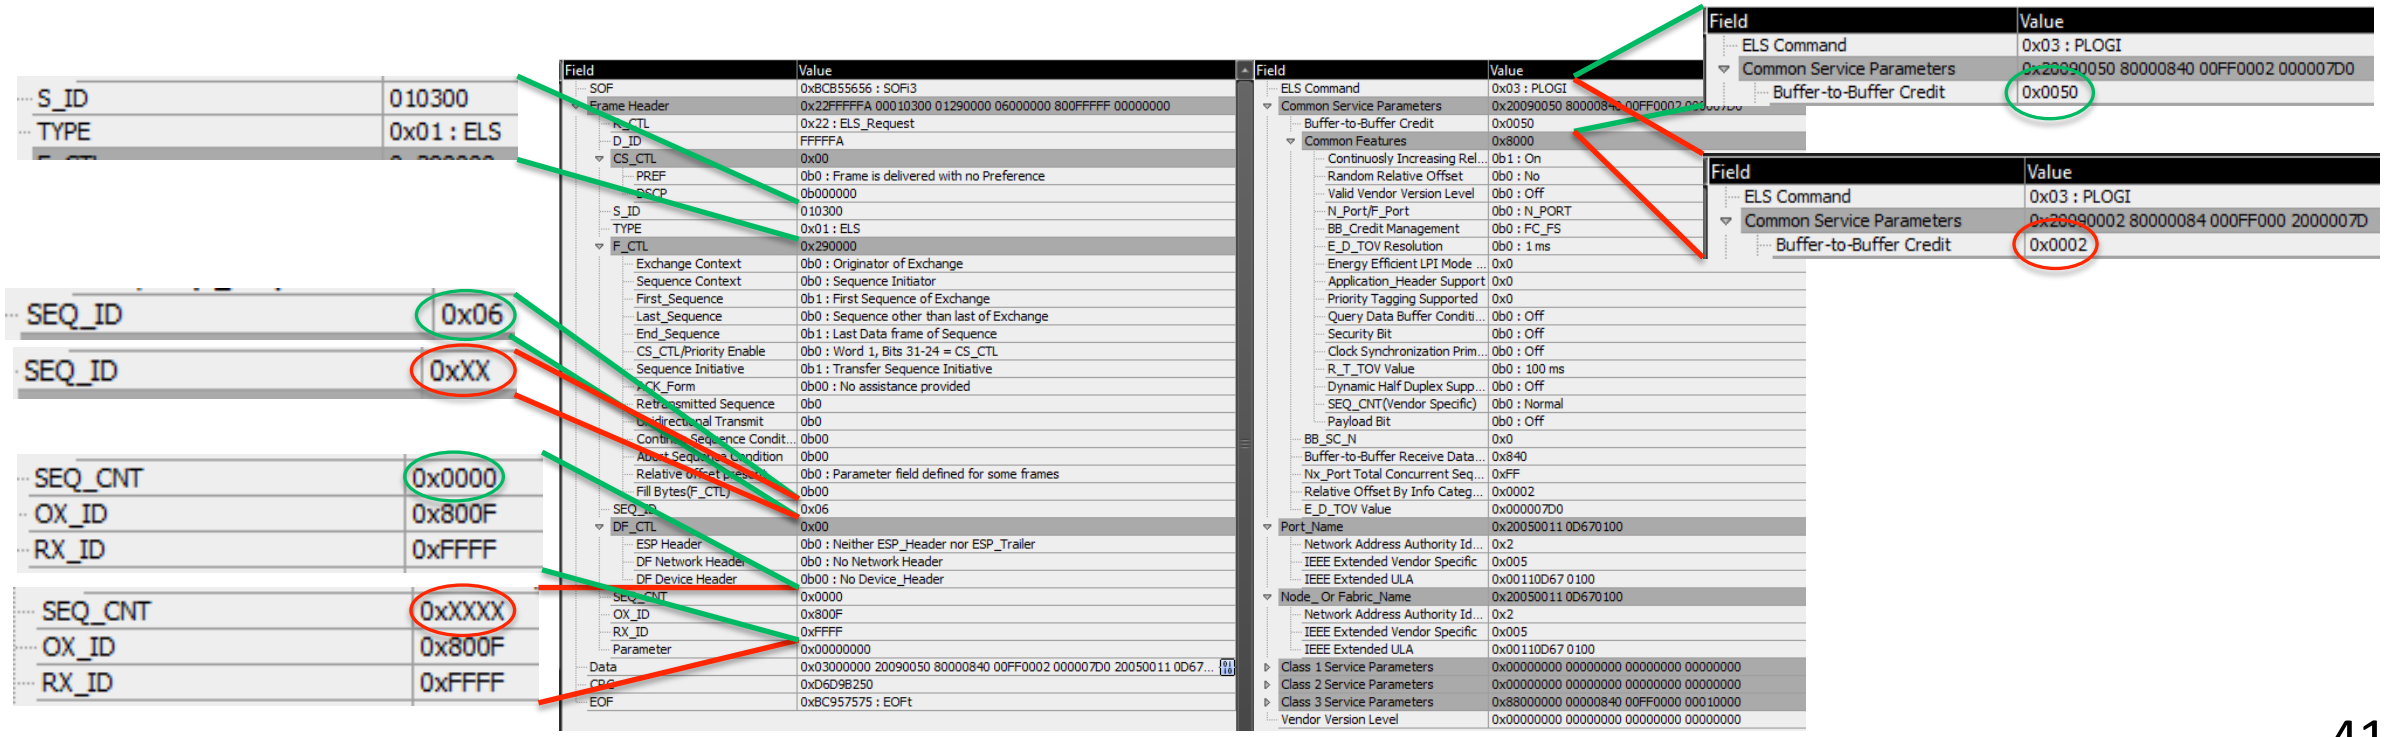

Example #3, ELS Jamming

- Port Login on **S_ID FFFFFA**
  Modify ELS Reply to Reject PLOGI – One time or Continuous jam

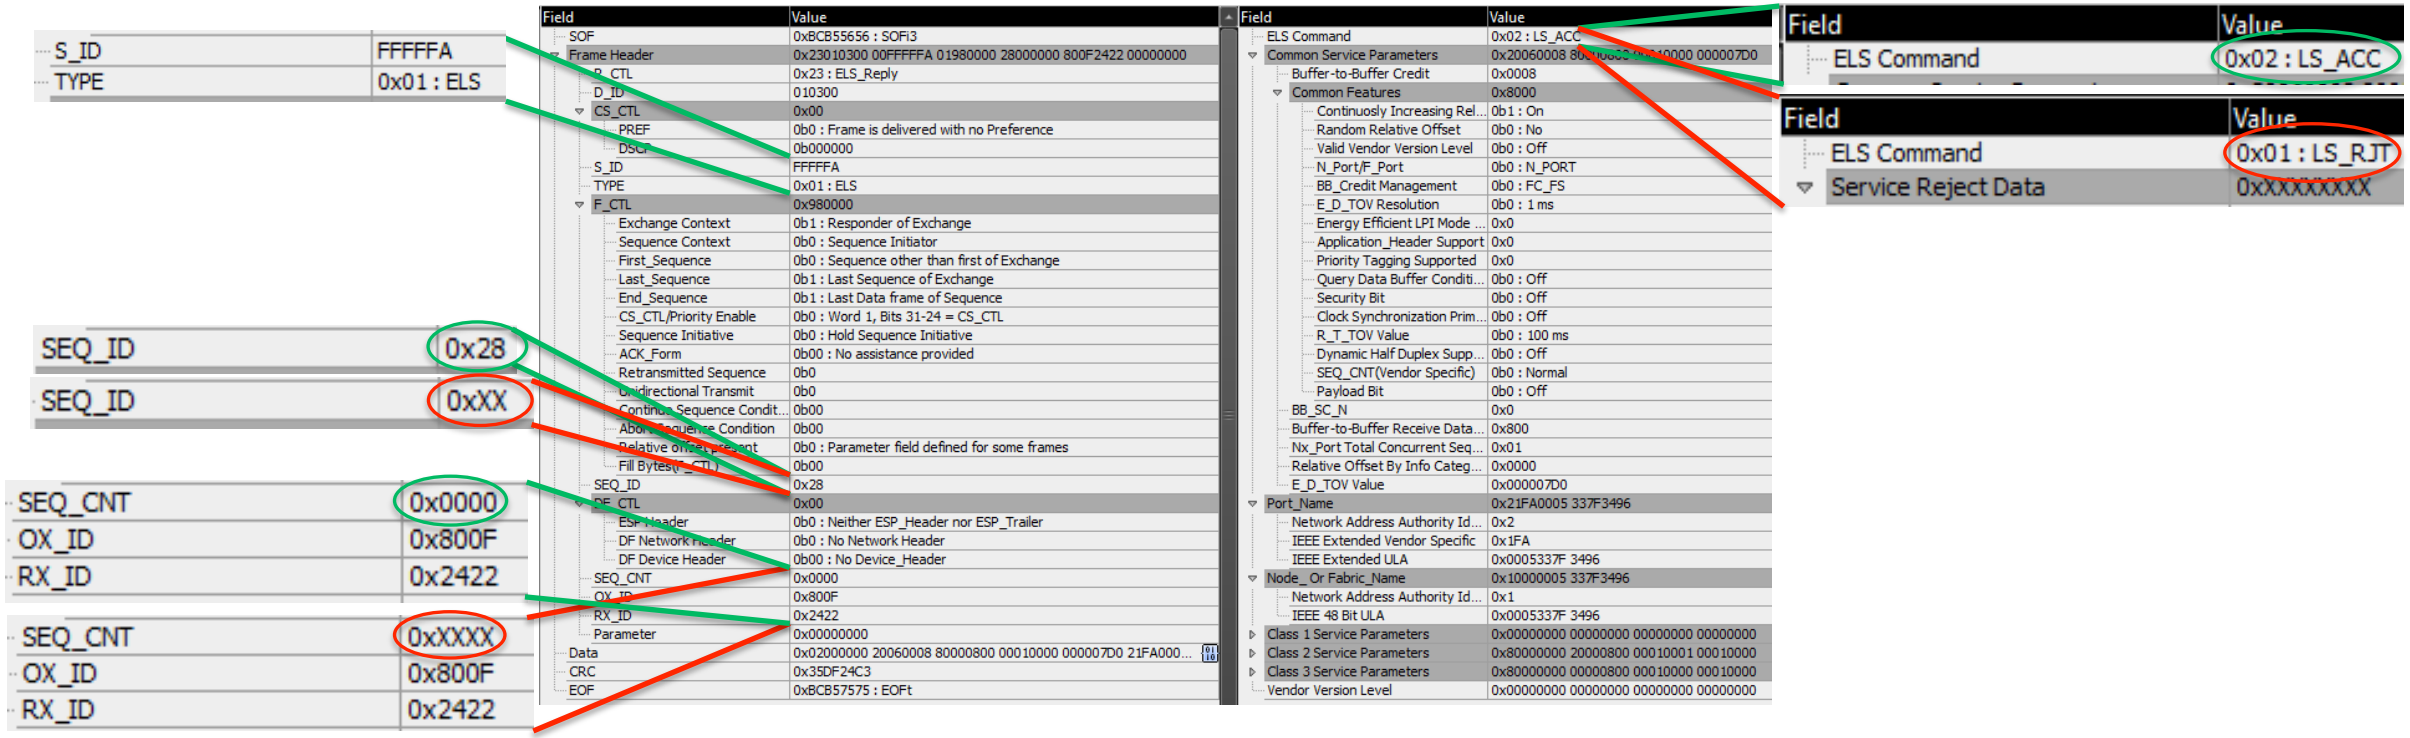

FC-NVMe Jamming Example #4

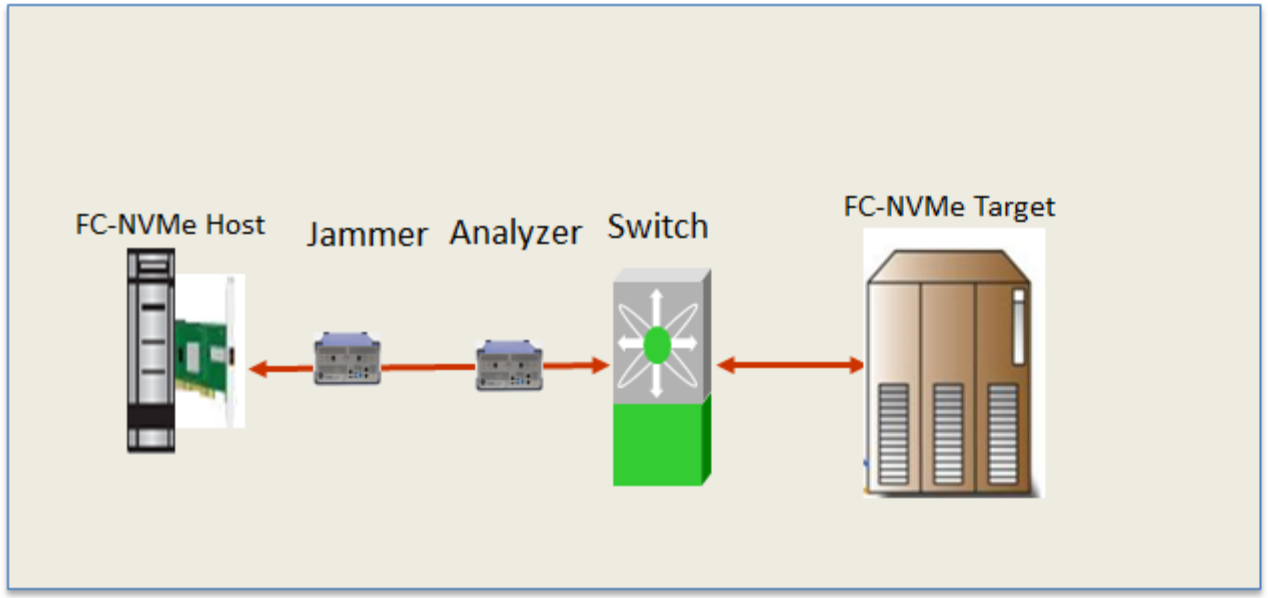

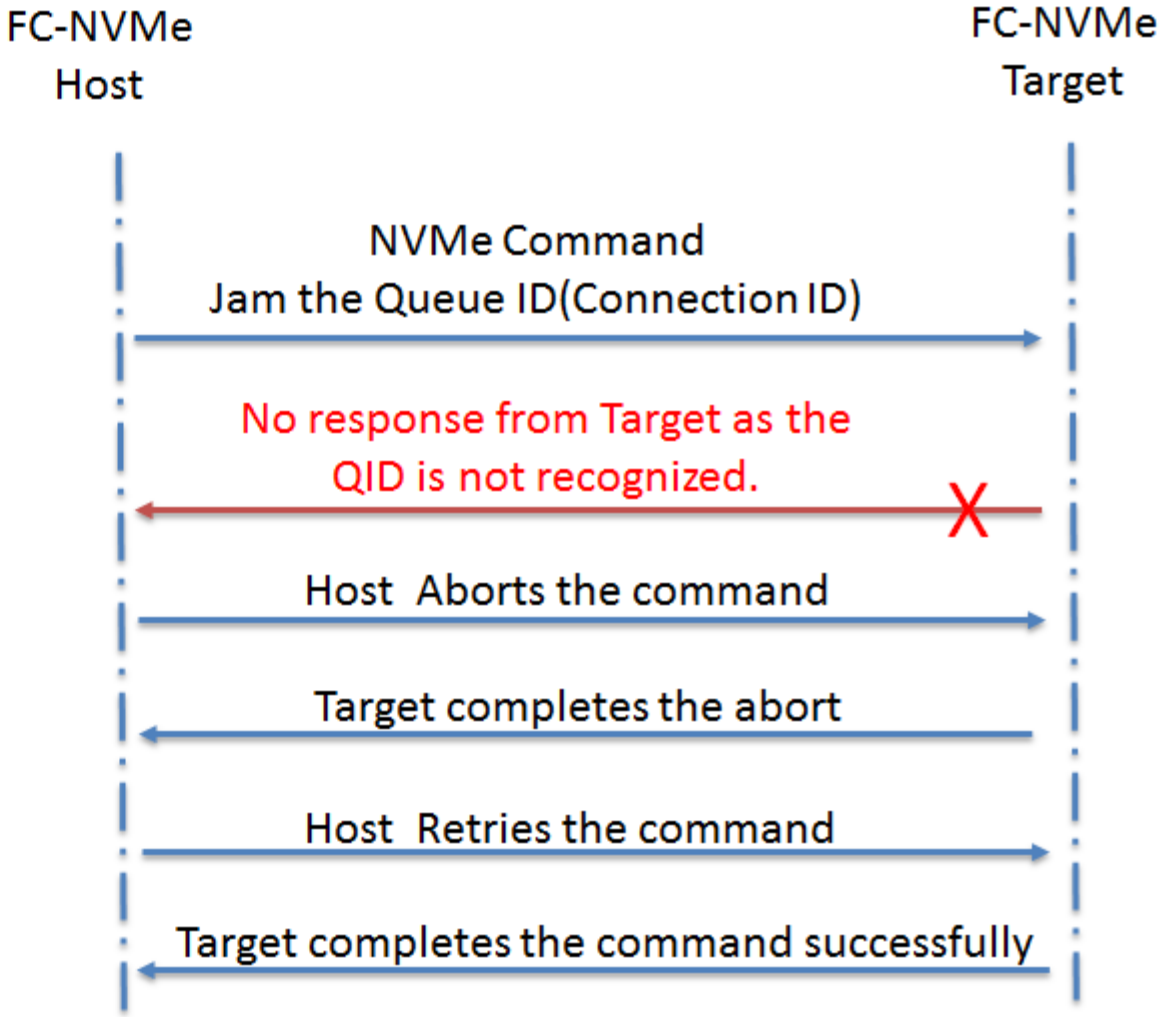

FC-NVMe Jamming Example #4 Contd.

Event View

Ic...	Bookmark	us.ns	Delta Time	Port	Side A	Side B	Summary	NVMe Connection Id...	Tag
Beg									
32 Bk	1) JAM	10492649.732		FC Port(1,1,2)		FC4Cmd	Write; NSID = 0x00000001; LBA	0000000000000FF2	0A55
32 Bk	2) Abort	31847642.178	21354992.446	FC Port(1,1,2)		FC4Cmd	Abort; SQID To Abort = 0x0002;	0000000000002400	0A56
32 Bk	3) Success	31893680.784	46038.607	FC Port(1,1,1)	FC4Status		Good Status;		0A56
32 Bk	4) ABTS	33851817.182	1958136.398	FC Port(1,1,2)		ABTS	ABTS; Basic Link Service; Abort		0A55
32 Bk	5) Retry	33851935.346	118.164	FC Port(1,1,2)		FC4Cmd	Write; NSID = 0x00000001; LBA	0000000000000F03	0A57
32 FR		33852110.772	175.426	FC Port(1,1,1)	FC4XRdy		DATA_RO = 0x00000000; BURS		0A57
32 FR		33852121.702	10.930	FC Port(1,1,2)		FC4SData	FC4SData; Len = 0x0800;		0A57
32 FR		33852122.332	0.630	FC Port(1,1,2)		FC4SData	FC4SData; Len = 0x0800;		0A57
32 Bk	6) Success	33853139.860	1017.528	FC Port(1,1,1)	FC4Status		Good Status;		0A57
32 Bk	7) BA_RJT	33893975.976	40836.116	FC Port(1,1,1)	BA_RJT		BA_RJT; Basic Link Service; Una		0A55

1) NVMe Connection ID Jammed from 0xF03 -> 0xFF2.
   - Host sent a command to a valid queue ID 0x03
   - Target receives a command to unknown queue 0xF2
   - Target Ignores the command and doesn't send a response.
2) Host times out after about ~21 seconds and issues an NVMe Abort to Abort the jammed command
3) Target completes the abort
4) Host sends a ABTS(FC layer)
5) Host retries the command with the correct Connection ID 0xFF2
6) Target Completes the command successfully.

Our Next FCIA Webcast:

Fibre Channel Zoning Basics

June 27, 2019

10:00 am PT/1:00 pm ET

Register at:

https://www.brighttalk.com/webcast/14967/355259

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at http://fibrechannel.org/ with answers to all the questions we received today
- Follow us on Twitter @FCIAnews for updates on future FCIA webcasts
- Visit our library of FCIA on-demand webcasts at http://fibrechannel.org/webcasts/ to learn about:
  - Fibre Channel Fundamentals
  - FC-NVMe
  - Long Distance Fibre Channel
  - Fibre Channel Speedmap
  - FCIP (Extension): Data Protection and Business Continuity
  - Fibre Channel Performance
  - FICON
  - Fibre Channel Cabling
  - 64GFC

Thank You