

SNIA

STORAGE NETWORKING INDUSTRY ASSOCIATION

EDUCATION

SAS Overview

Manu Gupta

Patni Computer Systems

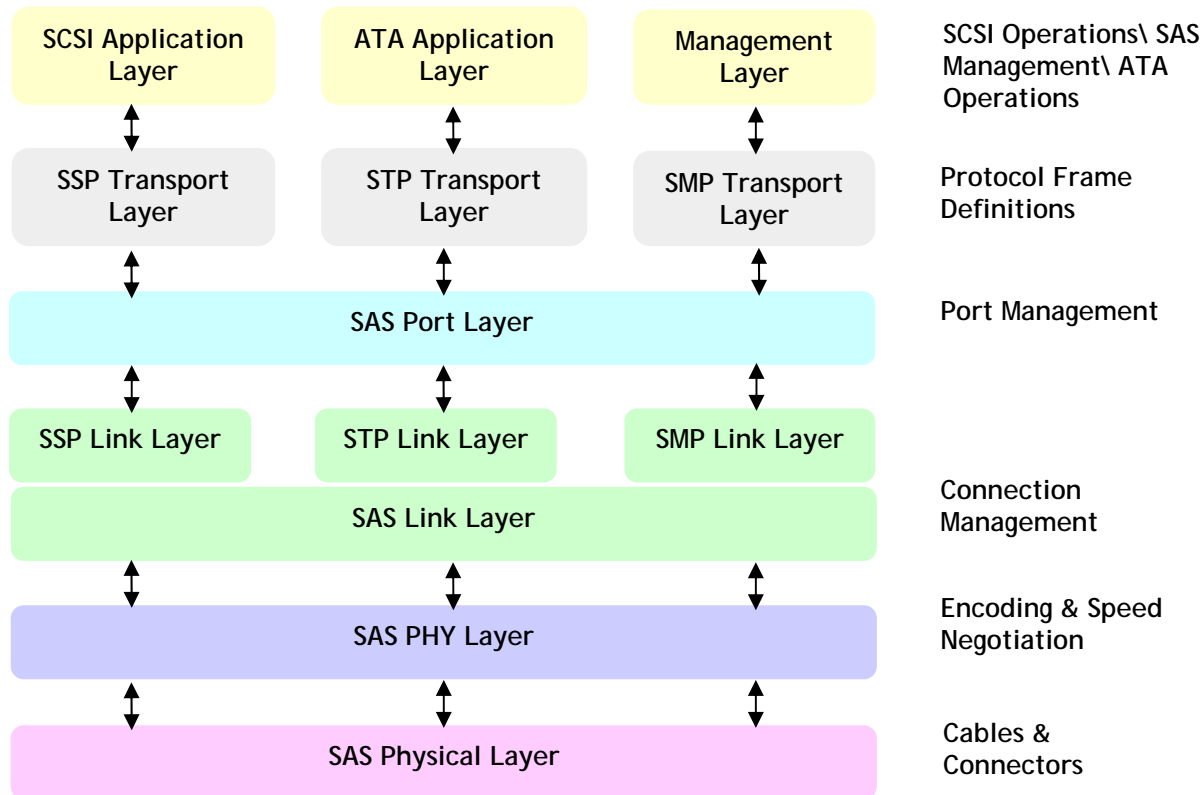
SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
 - Any slide or slides used must be reproduced without modification
 - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.

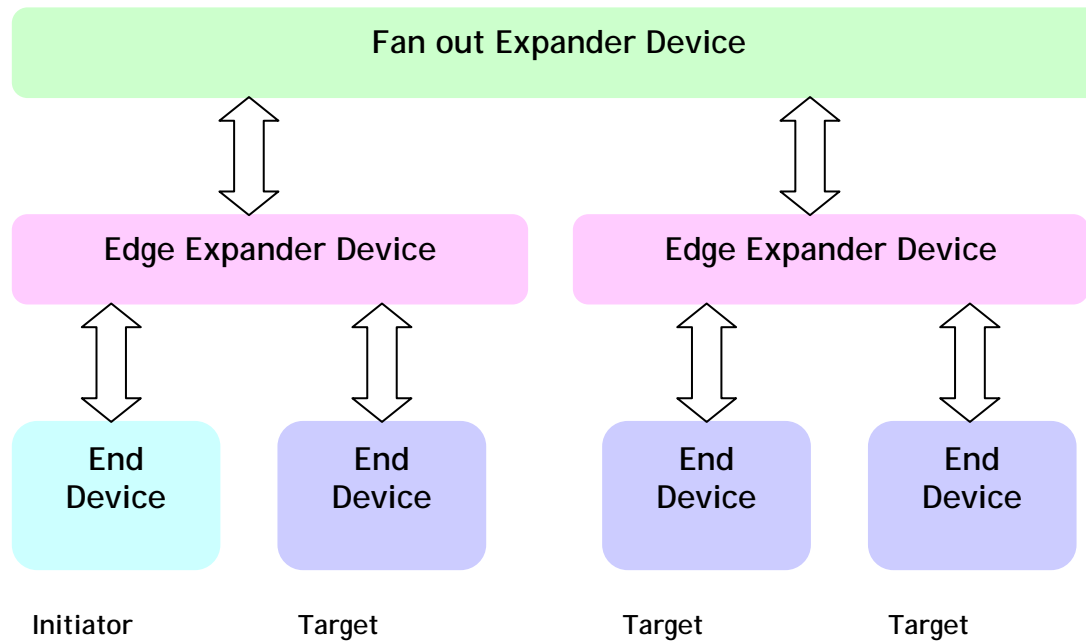
Abstract

This presentation explores the various aspects of SAS Architecture and SAS Transport Protocols. It elaborates the SAS Network Topology with SAS device discovery process. It depicts how scalable storage systems can be build using SAS. This presentation is intended for Technical Architects, OEMs, Storage Protocol Stack Developers and IT Managers

SAS Standard Layers



SAS Network Topology

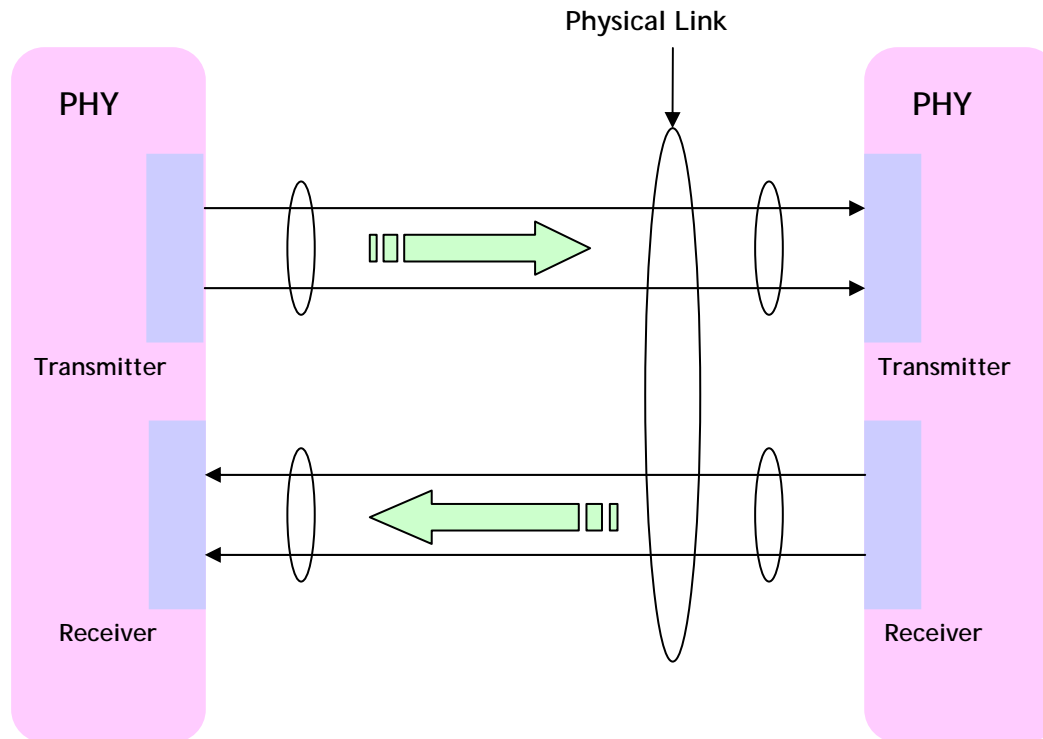


SAS Terminology

- PHY
- Physical Links
- SAS Address
- Port
- SAS Device (Initiator & Target)
- Expander Device

PHY Attributes

- Transmitter & Receiver




PHY Attributes

- Physical Link
 - Set of four wires used as two differential signal pairs
 - Physical Link between two PHYs
- Full Duplex Operation
 - Transmitter and Receiver operates at same physical link rate
- SAS Address inherited from SAS Port
- PHY Identifier unique within device
- SAS uses 8b/10b encoding

SAS Address


- Each SAS Port has worldwide unique 64-bit SAS Address.
- 24-bit Company ID assigned by IEEE
- 36-bit Vendor Specific ID assigned by organization associated with Company ID



SAS Port & SAS Device

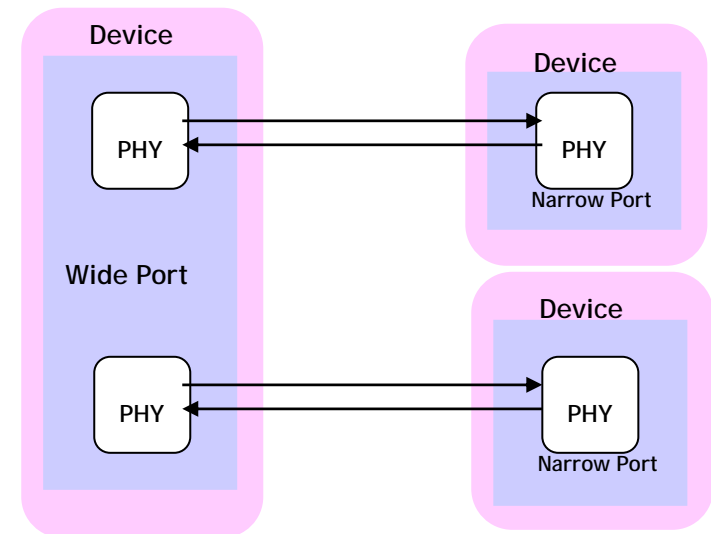
- Port contains one or more PHYs
 - Group of PHYs having same SAS Address
 - Narrow Port
 - If there is only one PHY in the Port
 - Wide Port
 - If there are more than one PHY in the Port
- SAS Device contains one or more ports
- Connections are made between PHYs but addressed to Ports

SAS Port & SAS Device



Identification Sequence

- Each PHY transmits IDENTIFY address frame
 - SAS Address of Port
 - Device Type
 - PHY Identifier
 - Role/Capabilities
- Each PHY receives same set of information from peer PHY in IDENTIFY address frame



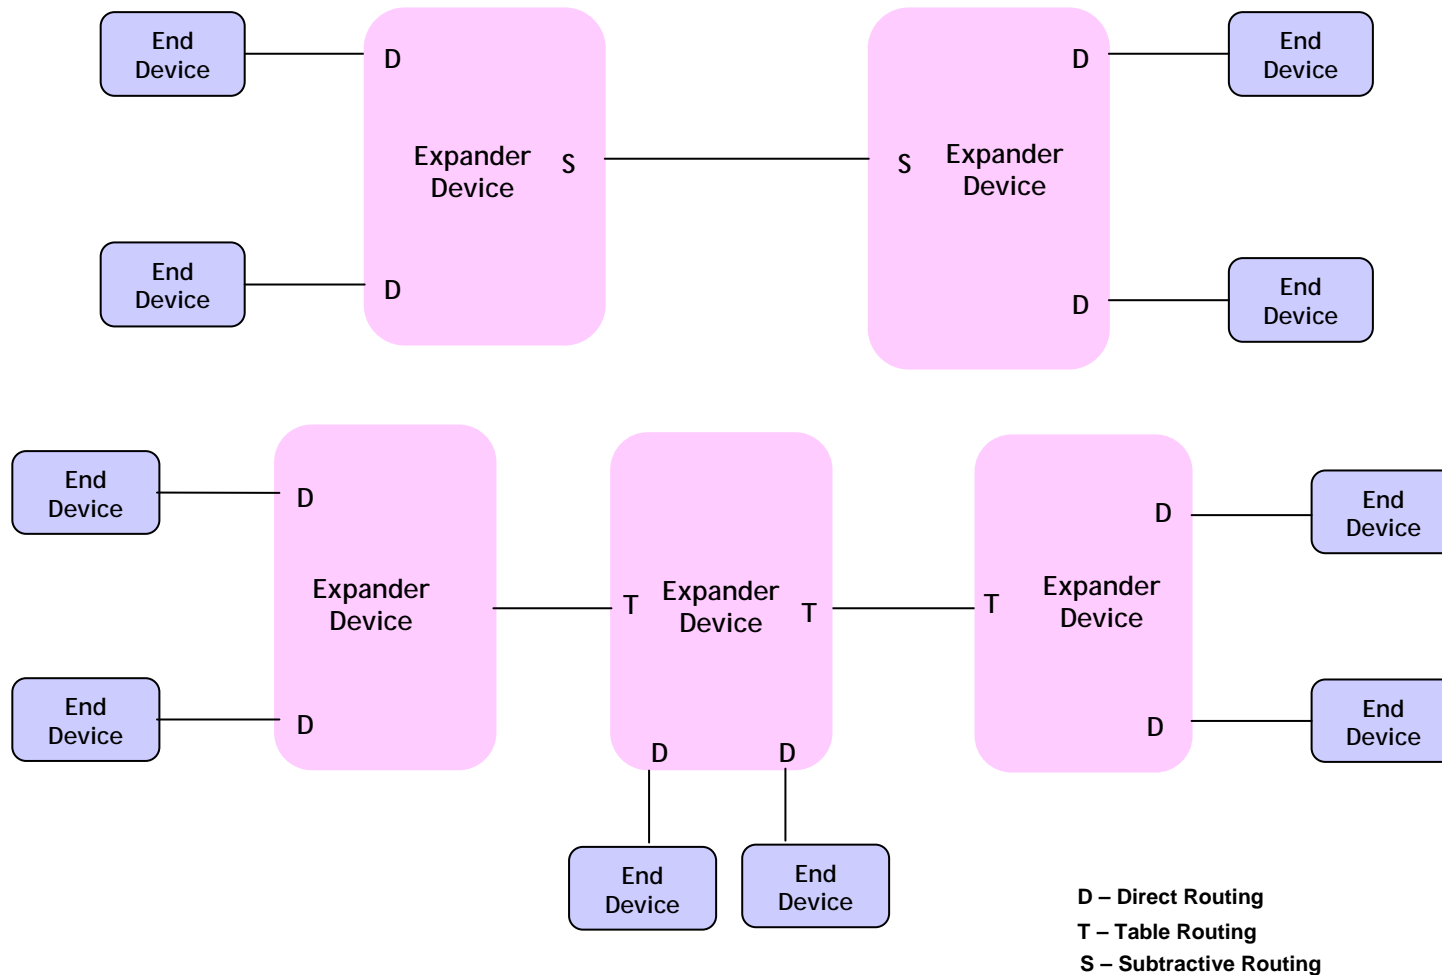
Expander Device

- Facilitates communication between SAS Devices
- Can have one or more PHYs
- Each PHY has same SAS address
- Expander ports are created based on attached SAS devices
- Each PHY in an expander had “Routing Attributes”
- Expanders may contain devices for enclosure management

Routing Attributes

- Direct – routes connection requests to attached end devices
- Table – routes connection requests to attached expander devices using Route Table
- Subtractive – routes unresolved connection requests to attached expander device
 - An expander device can have at most one defined port for subtractive routing

Routing Attributes



Route Table

- An association between SAS Destination Address and expander PHY to forward connection request

Expander Route Index	PHY Identifier				
	0	1	2	N
0					
1					
....					
M					

SAS Address & Disable Bit

PHY-based Route Table

Routed SAS Address	PHY Bit Map				
	0	1	2	N
SAS Address A					
SAS Address B					
SAS Address C					
...					

Forward Bit

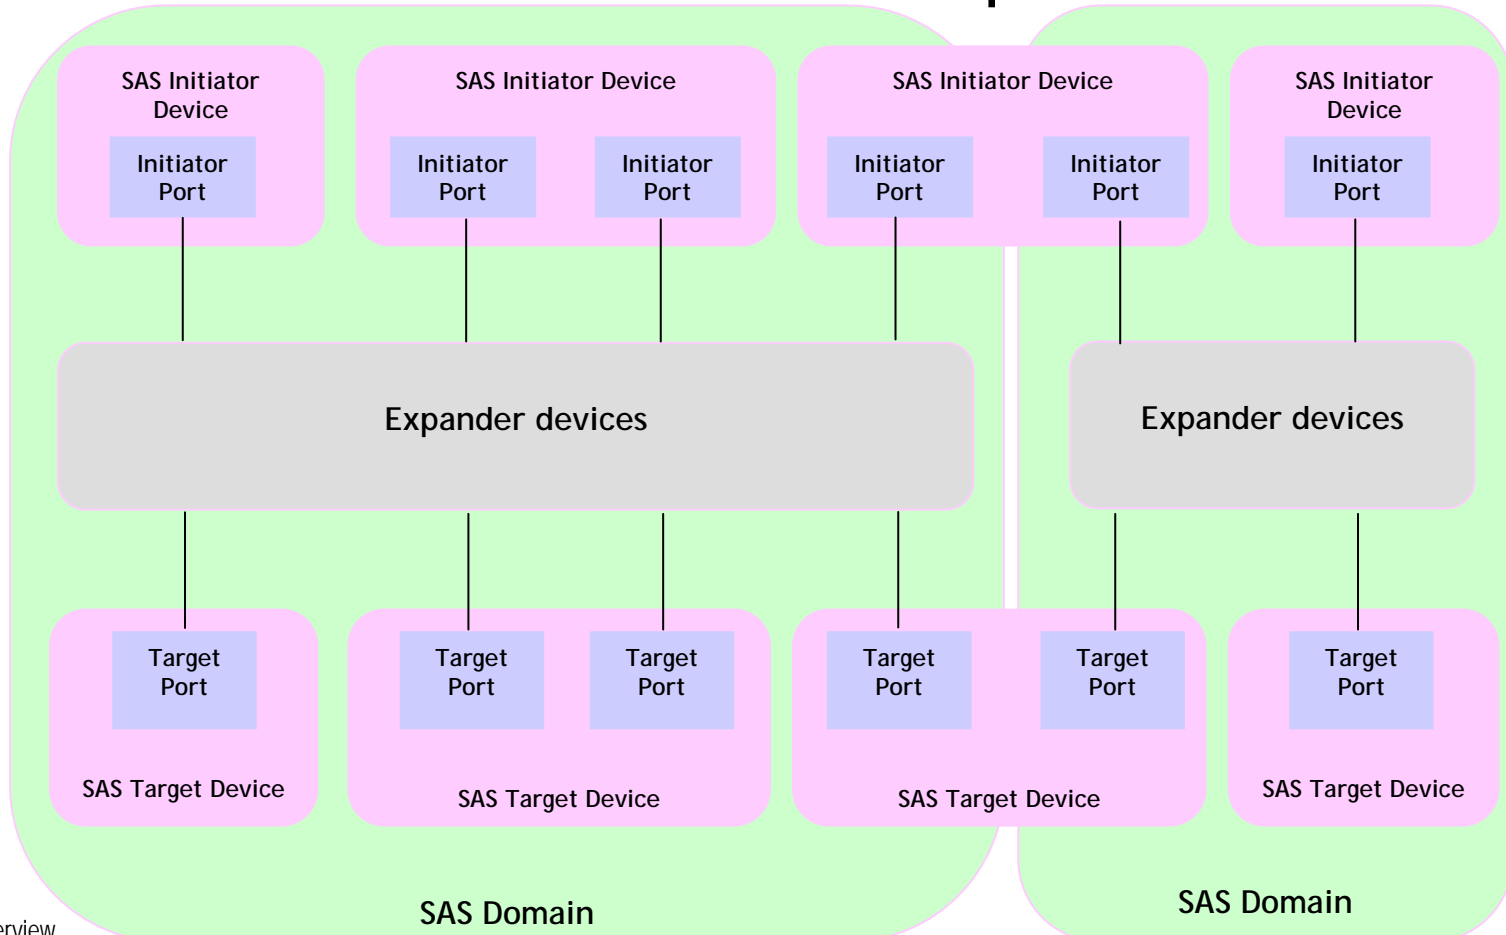
Expander-based Route Table

Expander Devices

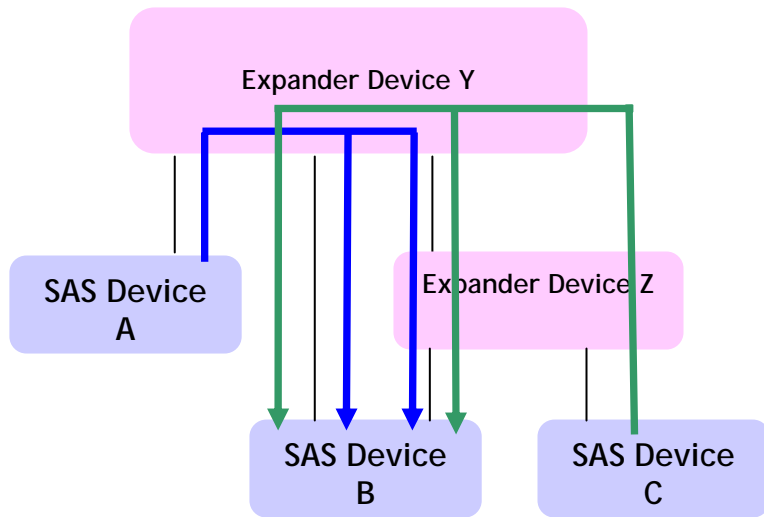
- Fan-out Expander Devices
 - Expanders without subtractive routing PHYs
 - Uses table routing when connected to other expander devices
- Edge Expander Devices
 - All subtractive PHYs must be connected to PHYs with same SAS address
- Edge Expander Device Set
 - Group of one or more edge expander devices
 - Table Routing is used to connect to Expander devices within Expander device set
 - Subtractive Routing is used to connect to other expander device set or Fan-out Expander Device

SAS Domain

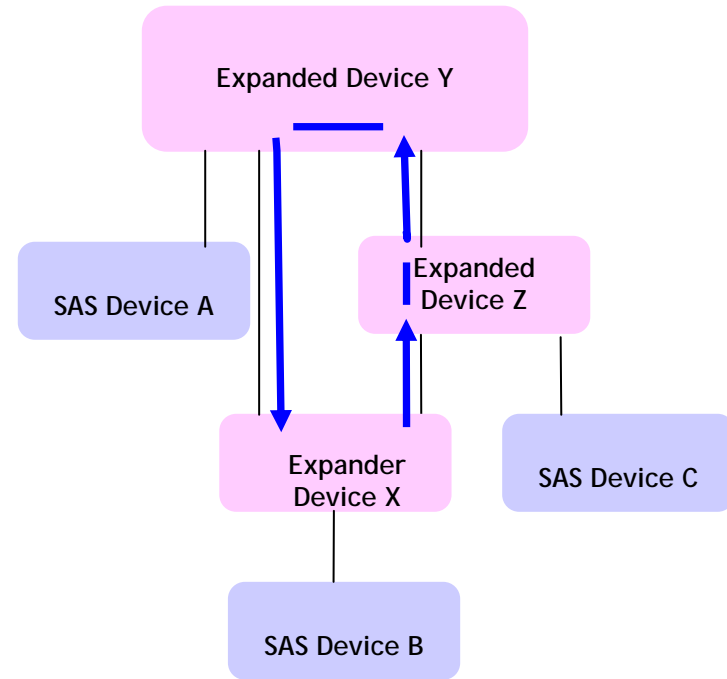
- Contains SAS Devices and Expander Devices



SAS Domain - Rules




No Multipaths



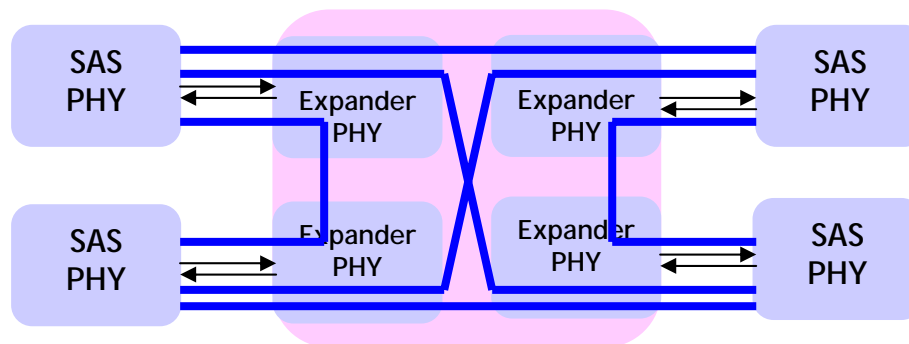
No Loops

Pathways & Connections

- Potential Pathway – set of physical links between SAS Initiator PHY and SAS Target PHY



Directly Attached

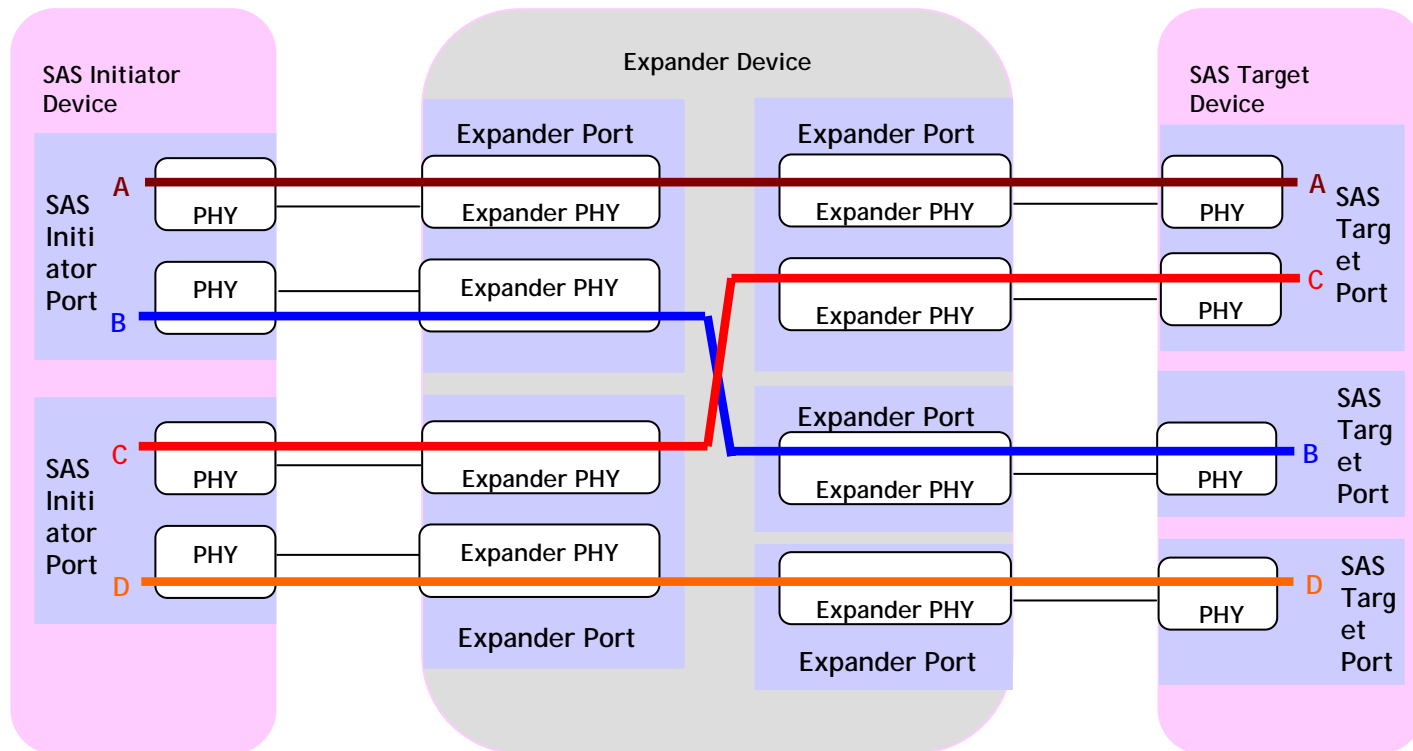


Attached through Expander

Pathways & Connections

- Pathway – set of physical link between SAS Initiator PHY and SAS Target PHY being used by connection
- Partial Pathway – set of physical links participating in connection request (not yet established)
- Connection – temporary association between SAS Initiator Port and SAS Target Port
 - Only one connection can be active on physical link at a given time
 - Number of connection established by a SAS Port is limited to no of PHYs in given SAS Port
 - SAS Initiator Wide Port can have multiple connection with different SAS Target Ports
 - SAS Target Wide Port can have multiple connection with different SAS Initiator Ports

Connection Example



SAS Management Protocol - SMP

- SMP – Protocol used by SMP Initiator port to communicate with SMP target in SAS Domain
- SMP Request Frame – Initiator to Target
- SMP Response Frame – Target to Initiator

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION							
2	RESERVED							
3	REQUEST LENGTH							
4 to (n-3)	ADDITIONAL REQUEST BYTES							
(n-3) to n	CRC							

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
4 to (n-3)	ADDITIONAL RESPONSE BYTES							
(n-3) to n	CRC							

SMP Functions

Code	SMP Function
00h	REPORT GENERAL
01h	REPORT MANUFACTURER INFORMATION
02h	READ GPIO REGISTER
03h	REPORT SELF-CONFIGURATION STATUS
04h	REPORT ZONE PERMISSION
10h	DISCOVER
11h	REPORT PHY ERROR LOG
12h	REPORT PHY SATA
13h	REPORT ROUTE INFORMATION
14h	REPORT PHY EVENT INFORMATION
15h	REPORT PHY BROADCAST COUNTS
16h	DISCOVER LIST
17h	REPORT EXPANDER ROUTE TABLE
80h	CONFIGURE GENERAL
81h	ENABLE DISABLE ZONING
82h	WRITE GPIO REGISTER
85h	ZONED BROADCAST
86h	ZONE LOCK
87h	ZONE ACTIVATE
88h	ZONE UNLOCK
8Ah	CONFIGURE ZONE PHY INFORMATION
8Bh	CONFIGURE ZONE PERMISSION TABLE
90h	CONFIGURE ROUTE INFORMATION
91h	PHY CONTROL
92h	PHY TEST FUNCTION
93h	CONFIGURE PHY EVENT INFORMATION
C0h - FFh	VENDOR SPECIFIC

SMP – Response Frame

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (01h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
....							
12 - 19	VENDOR IDENTIFICATION							
20 - 35	PRODUCT IDENTIFICATION							
....							
60 - 63	CRC							

REPORT MANUFACTURER INFORMATION RESPONSE

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (00h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
....							
6 - 7	EXPANDER ROUTE INDEXES							
8							
9	NUMBER OF PHYs							
48 - 51	CRC							

REPORT GENERAL RESPONSE

SAS Overview

© 2007 Storage Networking Industry Association. All Rights Reserved.

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (13h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
....							
6 - 7	EXPANDER ROUTE INDEX							
8							
9	PHY IDENTIFIER							
....							
16 - 23	ROUTED SAS ADDRESS							
40 - 43	CRC							

REPORT ROUTE INFORMATION RESPONSE

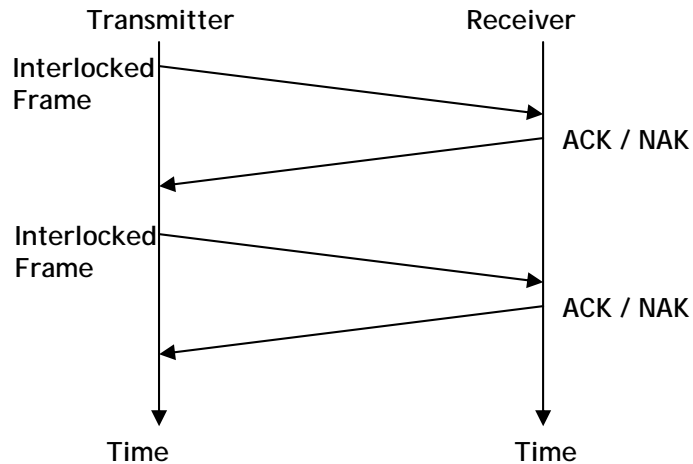
Serial SCSI Protocol - SSP

- Full Duplex Protocol
- All SSP Frames are acknowledged by ACK/NAK
- ACK/NAK must arrive within 1ms
- There is only one frame definition in SSP


Byte/Bit	7	6	5	4	3	2	1	0
0	FRAME TYPE							
1 – 3	HASHED DESTINATION SAS ADDRESS							
4	RESERVED							
5 – 7	HASHED SOURCE SAS ADDRESS							
8 – 9	RESERVED							
10	RESERVED			RETRY		RETRAN SMIT		CDP
11	RESERVED					NO OF FILL BYTES		
12 – 15	RESERVED							
16 – 17	TAG							
18 – 19	TARGET PORT TRANSFER TAG							
20 – 23	DATA OFFSET							
24 - m	INFORMATION UNIT							
m – (n-3)	FILL BYTES							
(n- 3) – n	CRC							

- Information Unit

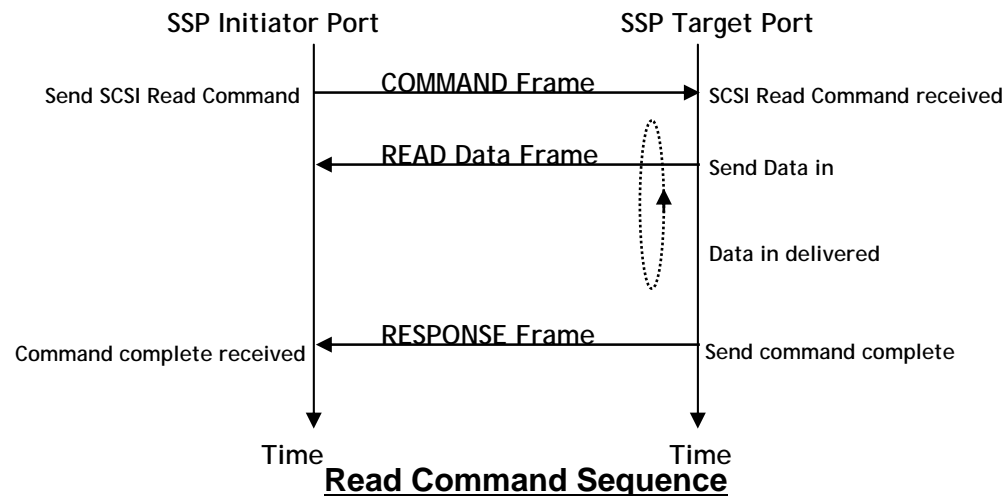
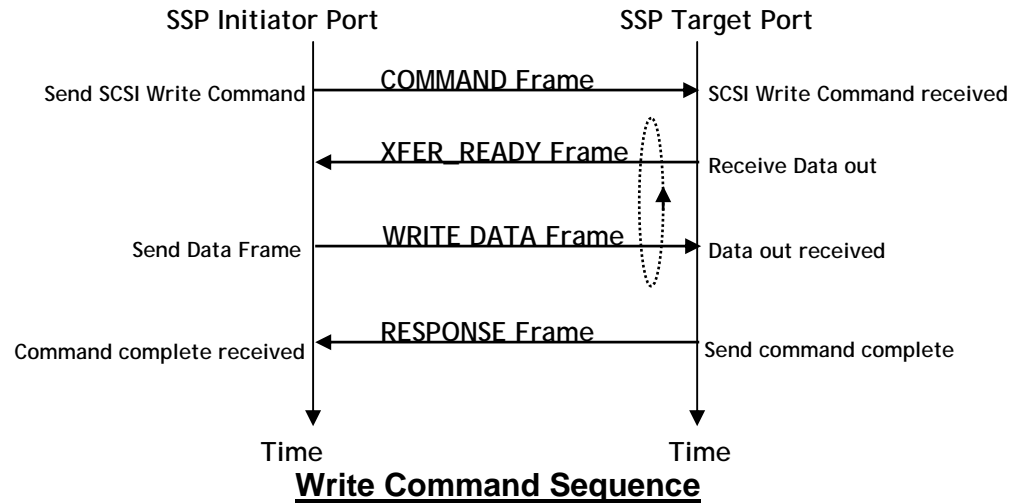
Code	Frame	Originator	Unit Size (Bytes)	Interlock Requirement
01h	DATA (READ/WRITE)	INITIATOR OR TARGET	1 to 1024	NON-INTELOCKED
05h	XFER_READY	TARGET	12	INTELOCKED
06h	COMMAND	INITIATOR	28 to 284	INTELOCKED
07h	RESPONSE	TARGET	24 to 1024	INTELOCKED
16h	TASK MANAGEMENT	INITIATOR	28	INTELOCKED
F0h – FFh	VENDOR SPECIFICS			



Interlocked Frame
Transmission Sequence



Non-Interlocked Frame
Transmission Sequence



SATA & STP

- SATA devices are not aware of multiple hosts
 - Flow Control is Host-to-device
- STP Frame transmission is defined by SATA (ATA/ATAPI-7 V3)
- STP encapsulates SATA Frames with connection management
- SATA devices tunnel SATA frames over SAS for transport as if they are directly connected
- STP flow control is point-to-point
- Expanders require SATA/STP bridge
- Set of primitives are defined for STP

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracknetworking@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

SNIA Education Committee

Howard Goldstein