



SAS Expander Initiator Based Configuration

Steven Fairchild
Senior Member Technical Staff

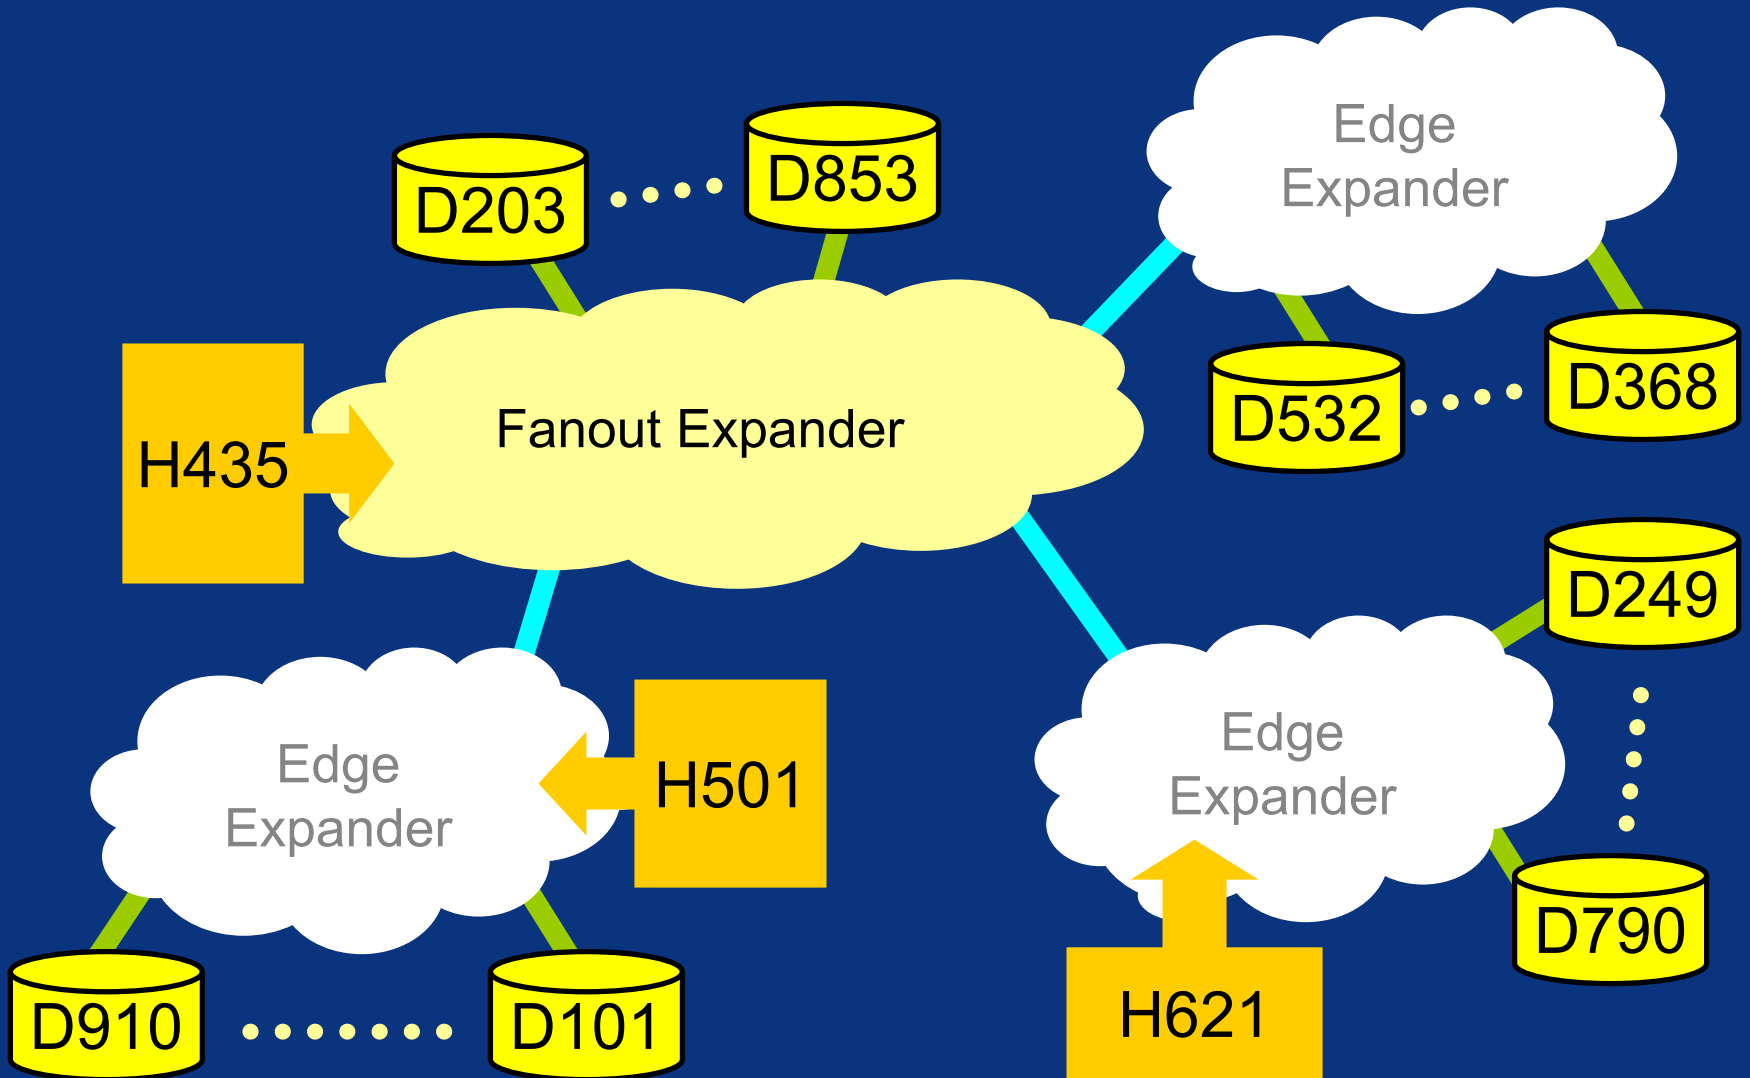
Industry Standard Servers

July 12, 2002

SAS expander initiator based configuration objectives

- use an initiator to:
 - discover topology
 - establish routing
 - manage topology
- allows expander device complexity to be reduced
- provides flexibility to solve topology issues with firmware in the initiator
- convert to a “by phy” discovery rather than the “by cloud” discovery defined in sas-r00c

the topology



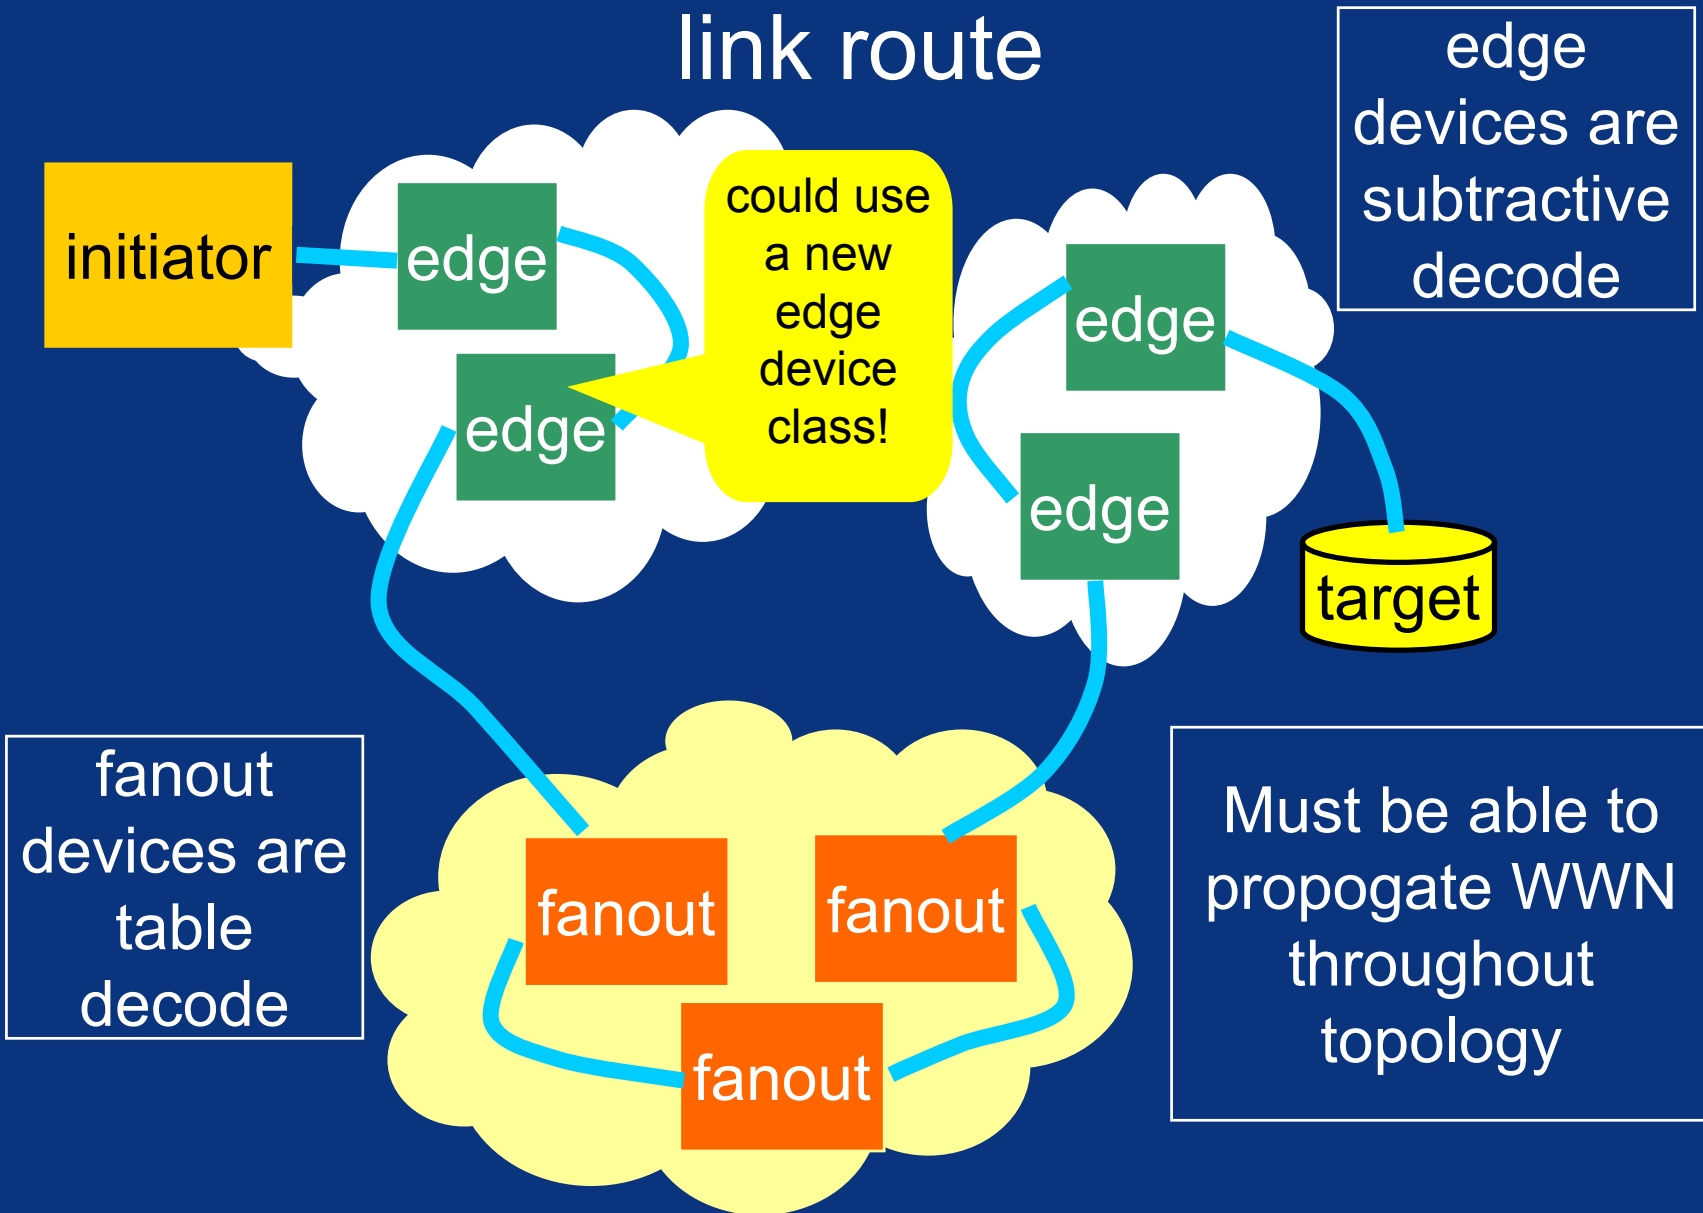
the challenge

- simple ASIC that can be used as a building block for the SAS clouds
 - no processor core
 - minimal memory requirements
 - fixed small number of Phys (~12 to 16)
 - multiple ASICs exist in an expander cloud
 - reasonable complexity
 - minimize proprietary considerations
 - maximize interoperability
 - low cost

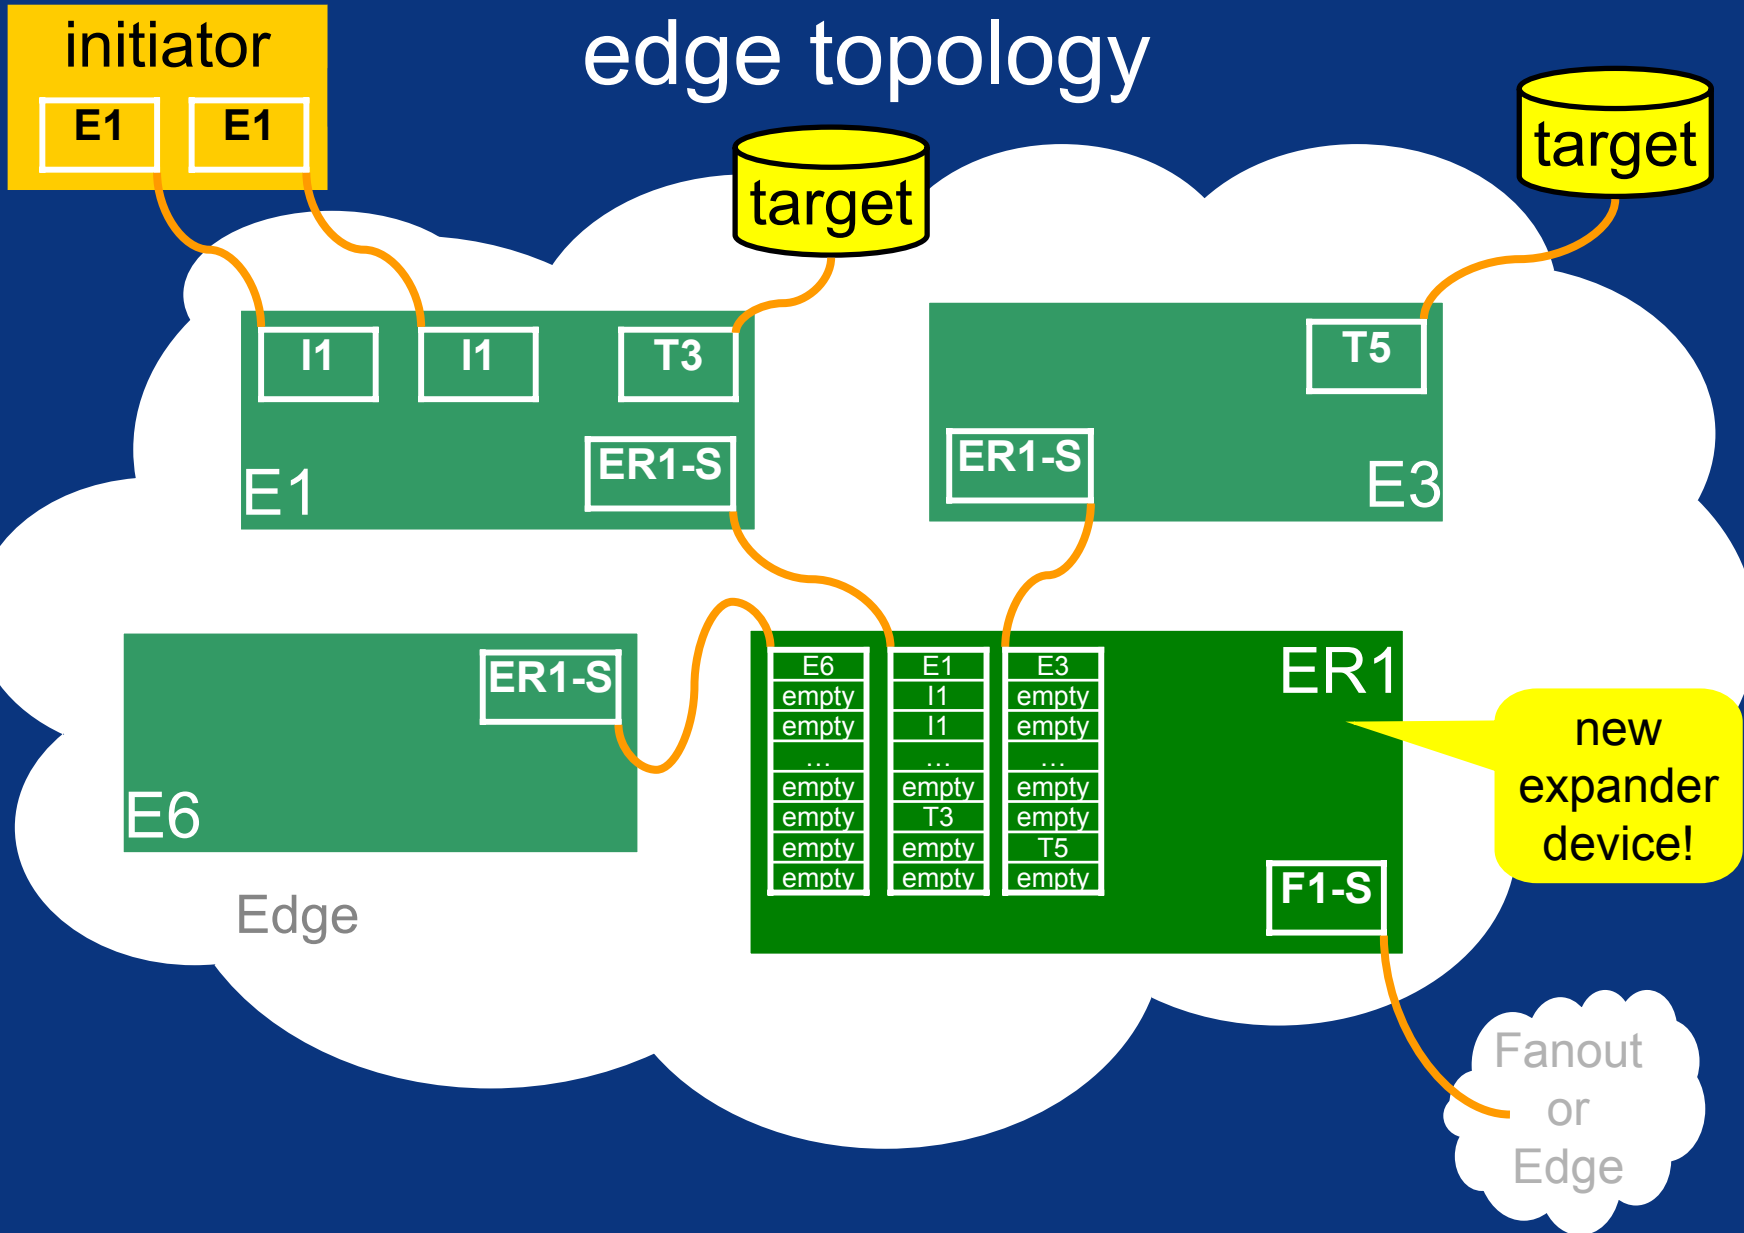
assumptions

- ASIC shall be configured via some non-volatile mechanism to indicate Phy characteristics
 - minimum/maximum speed supported
 - type of address decode
 - Phy activation
- ASIC shall support a WWN identifying itself and a WWN for each SATA device it can support
 - this implies support for 12-16 WWNs per ASIC
- ASIC shall support SMP protocol

link route



edge topology



device types

new edge
device
type!

- End
 - target or initiator SATA device
 - target or initiator SAS device
- Edge
 - expander device
 - primary use is to connect End devices to the Edge cloud
 - uses Subtractive routing for addressing
- Edge Route
 - expander device
 - primary use is to combine multiple Edge devices into a single Edge cloud
 - uses Subtractive routing when attached to another Edge Router or a Fanout device.
 - uses Table routing when attached to Edge devices.
- Fanout
 - expander device
 - primary use is to connect Edge clouds
 - uses Table routing for all addressing

classes of routing

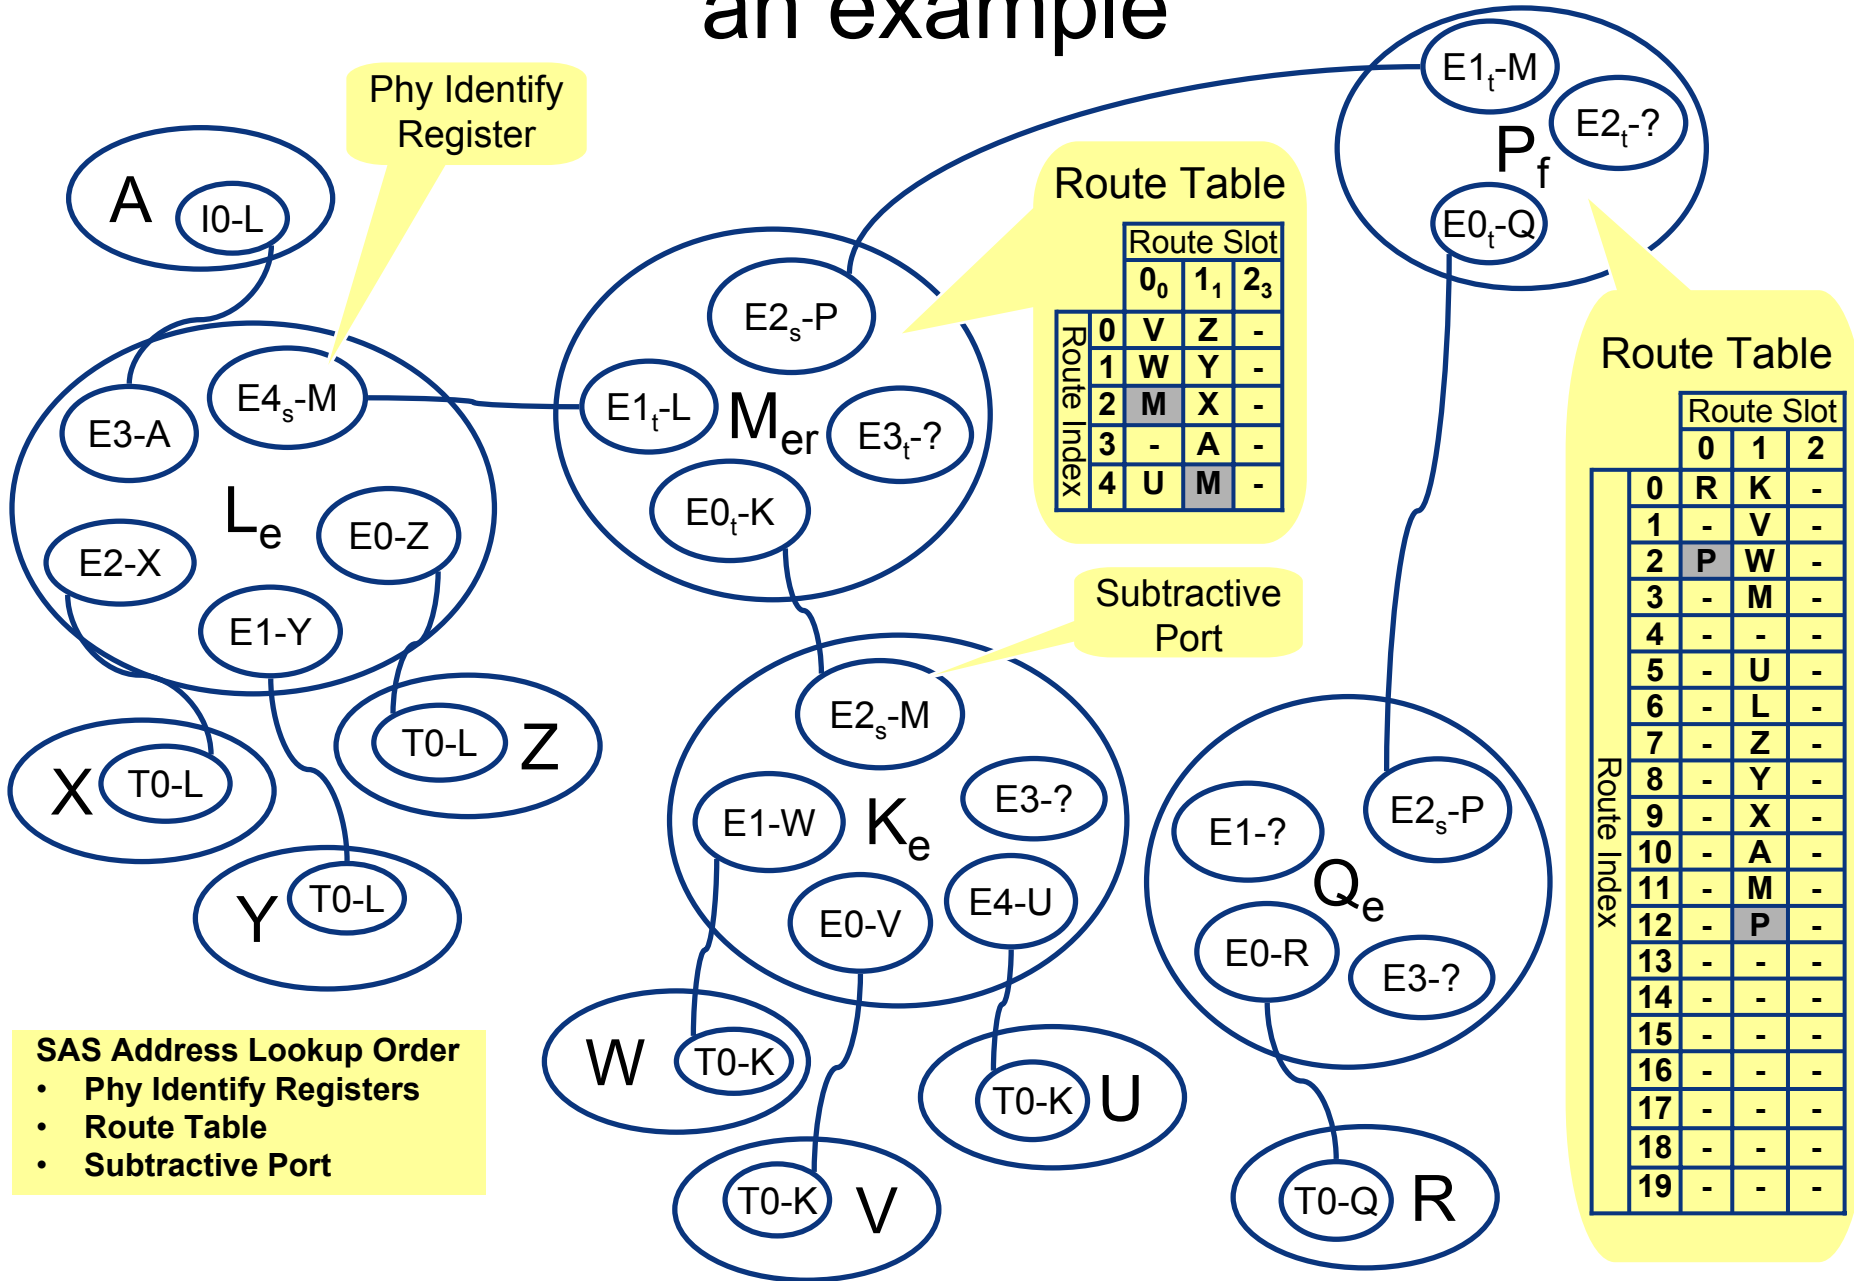
- No Routing
 - SATA device
 - SAS device
- Subtractive Routing
 - storage for a single WWN at each Phy and a single port (can be wide) designated as the subtractive port
 - the arbiter will compare the OPEN WWN with the WWN at each Phy
 - if the WWN is found, then the OPEN is routed to the appropriate Phy
 - if the WWN is not found, then the OPEN is routed to the subtractive port
- Table Routing
 - storage for N WWNs at each Phy
 - N is based on the arbitrary topology limiting rules in SAS, currently defined as 256
 - the arbiter will compare the OPEN WWN with the WWNs at each Phy
 - if the WWN is found, then the OPEN is routed to the appropriate Phy
 - if the WWN is not found, then the OPEN is rejected with a not found error
- Auto Routing
 - capable of automatically building all necessary route tables without initiator intervention

initiator based discovery

Initiator traverses the topology, configuring expanders with table routing requirements:

- the initiator will begin the topology walk by recognizing that it is connected to an expander.
- the initiator sequences through each of the Phys on the near expander.
- a subtractive route Phy does not need to be configured by the initiator.
- a table route Phy must be configured by the initiator.
- the route information is obtained by walking through each of the Phys, issuing the SMP Discover request, descending into each new expander Phy as it is encountered.
- the size of the table in the expander being configured must be sufficient to contain all of the Phy route entries including duplicates encountered due to traversing the topology.
- duplicate table entries are not collapsed and are positional relevant.
- self-referencing table entries or loops can be marked as unroutable.
- all initiators will walk and configure the topology, since table entries should be viewed identically by each initiator, no master or hold off is required.

an example



SAS Address Lookup Order

- Phy Identify Registers
- Route Table
- Subtractive Port

identify frame exchange

	7	6	5	4	3	2	1	0
0-3	SOAF							
0	Reserved				Identify Frame (0h)			
1	Phy Identifier							
2	Reserved				Maximum Phy Link Rate			
3	Device Type	STPi	STPt	SSPi	SSPt	SMPi	SMPt	
4-11	Device Name							
12-27	Reserved							
28-31	CRC							
0-3	EOAF							



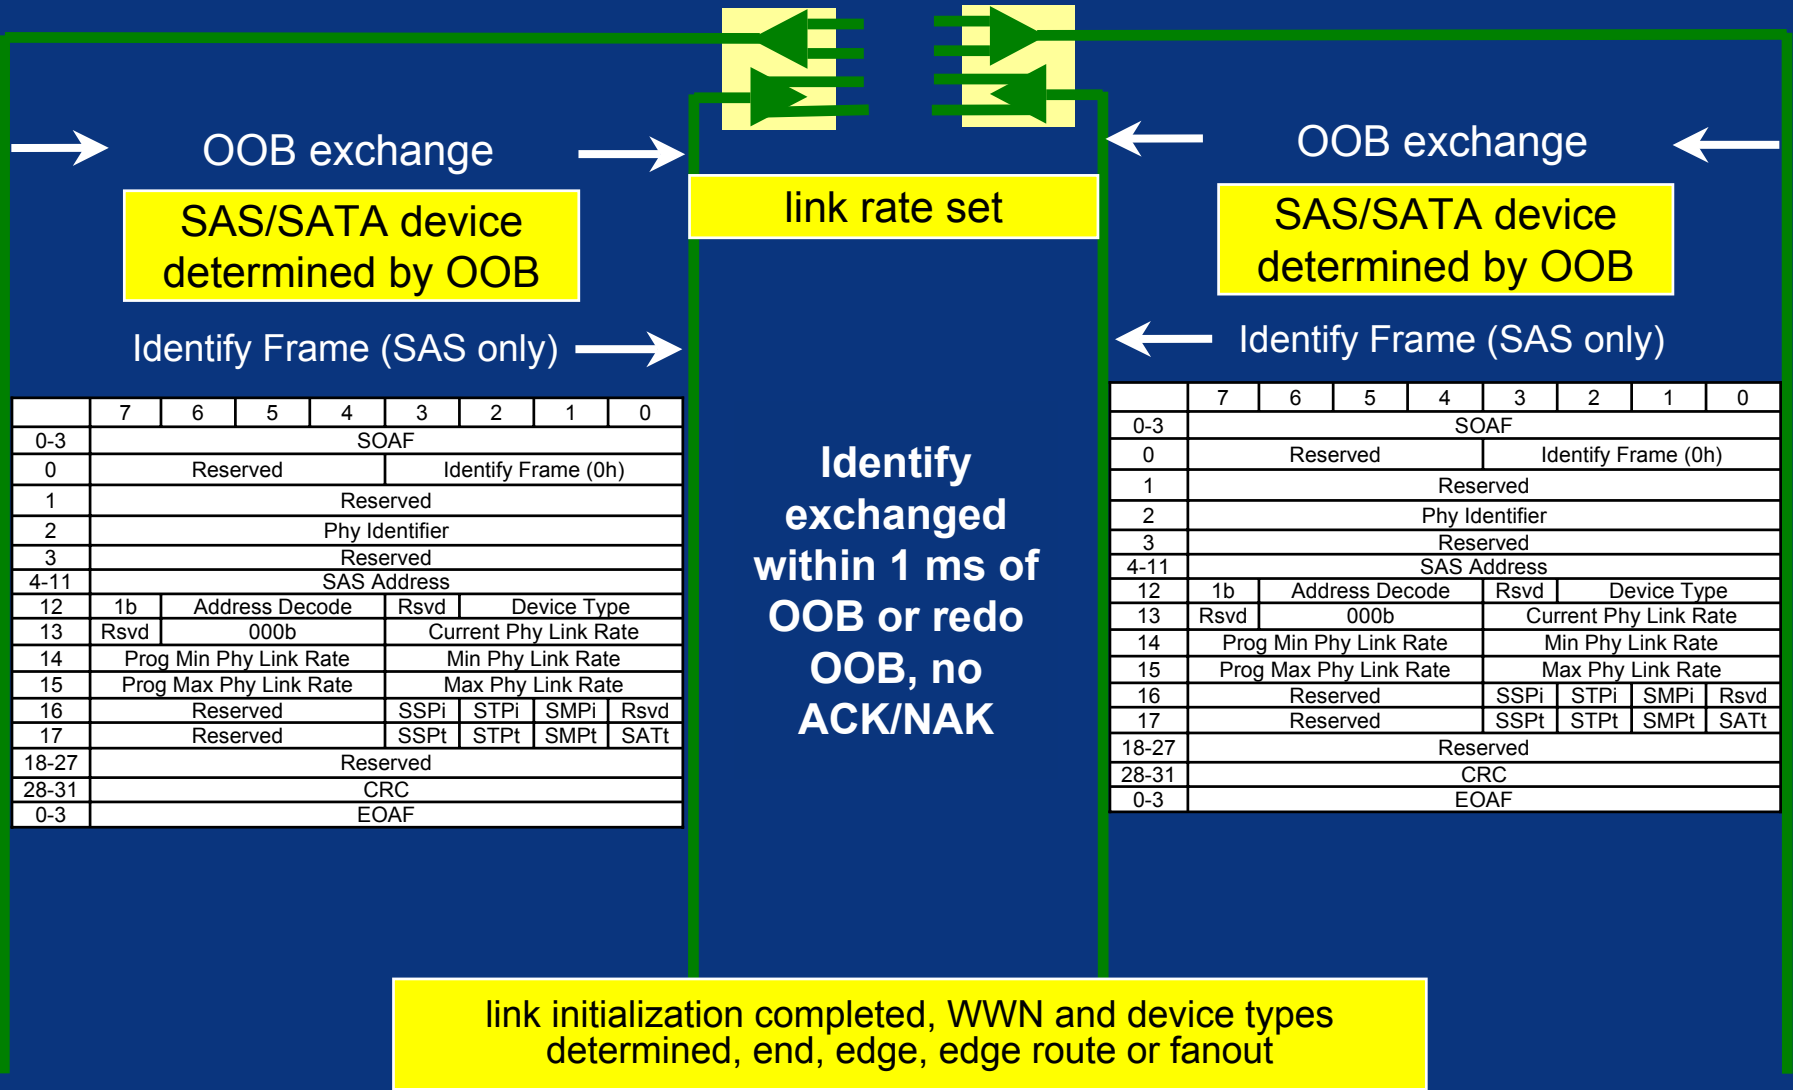
- used to exchange Phy specific information between connected Phys following OOB
- current sas-r00c definition does not provide enough information
- proposed definition:

	7	6	5	4	3	2	1	0
0-3	SOAF							
0	Reserved				Identify Frame (0h)			
1	Reserved							
2	Phy Identifier							
3	Reserved							
4-11	SAS Address							
12	Route	Address Decode			Rsvd	Device Type		
13	Rsvd	Link State			Current Phy Link Rate			
14	Prog Min Phy Link Rate				Min Phy Link Rate			
15	Prog Max Phy Link Rate				Max Phy Link Rate			
16	Reserved				SSPi	STPi	SMPi	Rsvd
17	Reserved				SSPt	STPt	SMPt	SATt
18-27	Reserved							
28-31	CRC							
0-3	EOAF							

identify frame definitions

- Phy Identifier
 - 00h – FFh Phy number
- SAS Address (WWN)
- Device Type
 - 000b End
 - 001b Edge
 - 010b Fanout
 - 011b Edge Router
 - 100b – 111b Reserved
- Address Decode
 - 000b None
 - 001b Subtractive
 - 010b Table
 - 011b Auto, self discovery
 - 100b – 111b Reserved
- Route
 - 0b Data is invalid for routing
 - 1b Data is valid for routing
- Current Phy Link Rate
 - 0000b Phy does not exist
 - 0001b Rate unknown
 - 0010b Phy disabled
 - 0011b 1.5 Gb/s
 - 0100b 3.0 Gb/s
 - 0101b – 1111b Reserved
- Link State
 - 000b Active
 - 001b Inactive
 - 010b Failed
 - 011b OOB in Progress
 - 100b Spinup Hold OOB
 - 101b Release Spinup Hold OOB
 - 110b – 111b Reserved
- Minimum Phy Link Rate
- Programmed Minimum Phy Link Rate
- Maximum Phy Link Rate
- Programmed Maximum Phy Link Rate
- SATt, SATA Protocol Target
- SMPi/SMPt, Serial Management Protocol Initiator/Target
- STPi/STPt, Serial Tunneled Protocol Initiator/Target
- SSPi/SSPt, Serial SCSI Protocol Initiator/Target

link initialization



	7	6	5	4	3	2	1	0
0-3	SOAF							
0	Reserved				Identify Frame (0h)			
1	Reserved							
2	Phy Identifier							
3	Reserved							
4-11	SAS Address							
12	1b	Address Decode			Rsvd	Device Type		
13	Rsvd	000b			Current Phy Link Rate			
14	Prog Min Phy Link Rate				Min Phy Link Rate			
15	Prog Max Phy Link Rate				Max Phy Link Rate			
16	Reserved				SSPi	STPi	SMPi	Rsvd
17	Reserved				SSPt	STPt	SMPt	SATt
18-27	Reserved							
28-31	CRC							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOAF							
0	Reserved				Identify Frame (0h)			
1	Reserved							
2	Phy Identifier							
3	Reserved							
4-11	SAS Address							
12	1b	Address Decode			Rsvd	Device Type		
13	Rsvd	000b			Current Phy Link Rate			
14	Prog Min Phy Link Rate				Min Phy Link Rate			
15	Prog Max Phy Link Rate				Max Phy Link Rate			
16	Reserved				SSPi	STPi	SMPi	Rsvd
17	Reserved				SSPt	STPt	SMPt	SATt
18-27	Reserved							
28-31	CRC							
0-3	EOAF							

**Identify
exchanged
within 1 ms of
OOB or redo
OOB, no
ACK/NAK**

link initialization completed, WWN and device types determined, end, edge, edge route or fanout

SMP functions

	7	6	5	4	3	2	1	0
0-3	SOAF							
0	Init	Protocol			Open Frame (1h)			
1	Features			Link Rate				
2-3	Initiator Connection Tag							
4-11	Destination Device Name							
12-19	Source Device Name							
20-21	Reserved							
22	Scale							
23	Arbitration Wait Time							
24-27	Reserved							
28-31	CRC							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1-23	Reserved							
24	Function							
25-m	Additional Request Bytes							
Fill Bytes (as needed to begin CRC on dword boundary)								
(n-3)-n	CRC							
0-3	EOF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1-23	Reserved							
24	Function Result							
25-m	Additional Response Bytes							
Fill Bytes (as needed to begin CRC on dword boundary)								
(n-3)-n	CRC							
0-3	EOF							

- **Discover (00h) - modified**
- **Report General (01h) - modified**
- Report SATA Capabilities (02h)
- Report Manufacturer Information (03h)
- *Report Route Information (04h) - added*
- Report Phy (10h) - deleted
- Report Phy Error Log (11h)
- Report Phy SATA (12h)
- Report Phy Device Names (13h) - deleted
- *Configure Route Information (80h) - added*
- Phy Control (90h)
- Phy Margin Control (91h)

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	Function							
2-3	Reserved							
4-m	Additional Request Bytes							
Fill Bytes (as needed to begin CRC on dword boundary)								
(n-3)-n	CRC							
0-3	EOF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	Function Result							
2-3	Reserved							
4-m	Additional Response Bytes							
Fill Bytes (as needed to begin CRC on dword boundary)								
(n-3)-n	CRC							
0-3	EOF							

discover

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (00h) – Discover							
2-3	Reserved							
4	Phy Identifier							
5-7	Reserved							
8-11	CRC							
0-3	EOF							



- definition presented is different from current sas-r00c specification, consolidated information from multiple SMP requests
- used to retrieve the expander configuration parameters for a specific Phy Identifier
- returns the identify frame as the SMP response payload

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Reserved				Identify Frame (0h)			
5	Reserved							
6	Phy Identifier							
7	Reserved							
8-15	SAS Address							
16	Route	Address Decode		Rsvd	Device Type			
17	Rsvd	Link State		Current Phy Link Rate				
18	Prog Min Phy Link Rate				Min Phy Link Rate			
19	Prog Max Phy Link Rate				Max Phy Link Rate			
20	Reserved			SSPi	STPi	SMPi	Rsvd	
21	Reserved			SSPt	STPt	SMPt	SATt	
22-27	Reserved							
28-31	CRC							
0-3	EOF							

initiator based discovery

initiator device

expander device

link initialization completed

Open Edge Device



	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

Discover Request



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (00h) – Discover							
2-3	Reserved							
4	Phy Identifier							
5-7	Reserved							
8-11	CRC							
0-3	EOF							

Initiator uses SMP Discover requests to determine Phy details for expander, repeat this sequence until all Phy information on the expander is obtained

Open Acknowledge



Discover Response



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Response (00h) – Function Accepted							
2-3	Reserved							
4-27	Identify Frame							
28-31	CRC							
0-3	EOF							

Close Connection



Close Connection



~100 bytes exchanged per Phy to get discovery information,
@ 150 MB/s ~ = 1 us/Phy ignoring overhead

discovery complete when all expander phys have been traversed

report general

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (01h) – Report General							
2-7	Reserved							
8-11	CRC							
0-3	EOF							



- definition presented is different from current sas-r00c specification, reduced to eliminate bit mask information
- used to retrieve global expander configuration parameters
- added Maximum Route Slot and Maximum Route Index used to defined the size of the route table for an expander with table routing

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Number of Phys							
5	Maximum Route Slot							
6	Maximum Route Index							
7	Reserved							
8-11	CRC							
0-3	EOF							

report SATA capabilities

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (02h) – Report SATA capabilities							
2-7	Reserved							
8-11	CRC							
0-3	EOF							

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
- used to retrieve the SATA unique configuration parameters for a specific expander device

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Reserved						Que	SATA
5	SATA Version							
6	Number of Initiator Ports							
7	Reserved							
8-11	CRC							
0-3	EOF							

report manufacturer information

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
- used to retrieve the manufacturing data for a specific expander device

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (03h) – Report Manufacturer Information							
2-7	Reserved							
8-11	CRC							
0-3	EOF							



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Additional Length (33h)							
5-7	Reserved							
8-15	Vendor Identification							
16-31	Product Identification							
32-35	Product Revision Level							
36-55	Vendor Unique							
56-59	CRC							
0-3	EOF							

report route information

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (04h) – Report Route Information							
2-3	Reserved							
4	Route Slot							
5	Route Index							
6-7	Reserved							
8-11	CRC							
0-3	EOF							



- new function
- used to retrieve the route information resident in a specific expander device
- route slot and route index are used to address into a fixed memory array that contains the route information
- route array size is maximum route slot by maximum route index
- route slot is 0 based and its maximum shall correspond to the number of phys on the expander device (maximum route slot)
- route index is 0 based and its maximum shall correspond to the number of phys addressible behind each phy (maximum route index)
- route (1b) indicates the route information is routable

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Route Slot							
5	Route Index							
6	Phy Identifier							
7	Reserved							
8-15	SAS Address							
16	Route	Reserved						
17	Reserved				Current Phy Link Rate			
18-19	Reserved							
20	Reserved				SSPi	STPi	SMPi	Rsvd
21	Reserved				SSPt	STPt	SMPt	SATt
22-27	Reserved							
28-31	CRC							
0-3	EOF							

report Phy

- defined in current version of sas-r00c document
- recommend deleting in favor of consolidated identify frame
 - expander complexity is reduced
 - obtained in new definition of Discover

report Phy error log

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
- used to retrieve the Phy error counts

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (11h) – Report Phy Error Log							
2-3	Reserved							
4	Phy Identifier							
5-7	Reserved							
8-11	CRC							
0-3	EOF							



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Report Phy Error Log Result							
5-7	Reserved							
8-11	Invalid Character Count							
12-15	Disparity Error Count							
16-19	Loss of Bit Sync Count							
20-23	CRC							
0-3	EOF							

report Phy SATA

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
- used to retrieve the SATA FIS information for a specific Phy

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (12h) – Report Phy SATA							
2-3	Reserved							
4	Phy Identifier							
5-7	Reserved							
8-11	CRC							
0-3	EOF							



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Report Phy SATA Result							
5-7	Reserved							
8-27	Register Device to Host FIS							
28-31	CRC							
0-3	EOF							

report Phy device names

- defined in current version of sas-r00c document
- recommend deleting in favor of consolidated identify frame
 - expander complexity is reduced
 - obtained in new definition of Discover
- assumes transition to a “by Phy” discovery instead of a “by cloud” discovery
 - reduces expander complexity
 - eliminates barriers to larger topologies in the future

configure route information

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Response							
1	SMP Function (80h) – Configure Route Information							
2-3	Reserved							
4	Route Slot							
5	Route Index							
6	Phy Identifier							
7	Reserved							
8-15	SAS Address							
16	Route	Reserved						
17	Reserved	Current Phy Link Rate						
18-19	Reserved							
20	Reserved	SSPi	STPi	SMPi	Rsvd			
21	Reserved	SSPt	STPt	SMPt	SATt			
22-27	Reserved							
28-31	CRC							
0-3	EOF							


- new function
- used to configure the route information resident in a specific expander device
- route slot and route index are used to address into a fixed memory array that contains the route information
- route array size is maximum route slot by maximum route index
- route slot is 0 based and its maximum shall correspond to the number of phys on the expander device (maximum route slot)
- route index is 0 based and its maximum shall correspond to the number of phys addressible behind each phy (maximum route index)
- route (1b) indicates the route information is routable

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Route Slot							
5	Route Index							
6	Phy Identifier							
7	Reserved							
8-15	SAS Address							
16	Route	Reserved						
17	Reserved	Current Phy Link Rate						
18-19	Reserved							
20	Reserved	SSPi	STPi	SMPi	Rsvd			
21	Reserved	SSPt	STPt	SMPt	SATt			
22-27	Reserved							
28-31	CRC							
0-3	EOF							

initiator route entry configuration

initiator device

edge route/fanout device



discovery completed

Open Edge Device



	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

Configure Request



	7	6	5	4	3	2	1	0	
0-3	SOF								
0	Information Unit Type (40h) – SMP Request								
1	SMP Function (80h) – Configure Route Information								
2-3	Reserved								
4	Route Slot								
5	Route Index								
6	Phy Identifier								
7	Reserved								
8-15	SAS Address								
16	Route	Reserved							
17	Reserved				Current Phy Link Rate				
18-19	Reserved								
20	Reserved		SSPi	STPi	SMPi	Rsvd			
21	Reserved		SSPt	STPt	SMPt	SATt			
22-27	Reserved								
28-31	CRC								
0-3	EOF								

Close Connection



Open Acknowledge



Configure Response



	7	6	5	4	3	2	1	0	
0-3	SOF								
0	Information Unit Type (41h) – SMP Response								
1	SMP Result (00h) – Function Accepted								
2-3	Reserved								
4	Route Slot								
5	Route Index								
6	Phy Identifier								
7	Reserved								
8-15	SAS Address								
16	Route	Reserved							
17	Reserved				Current Phy Link Rate				
18-19	Reserved								
20	Reserved		SSPi	STPi	SMPi	Rsvd			
21	Reserved		SSPt	STPt	SMPt	SATt			
22-27	Reserved								
28-31	CRC								
0-3	EOF								

Close Connection



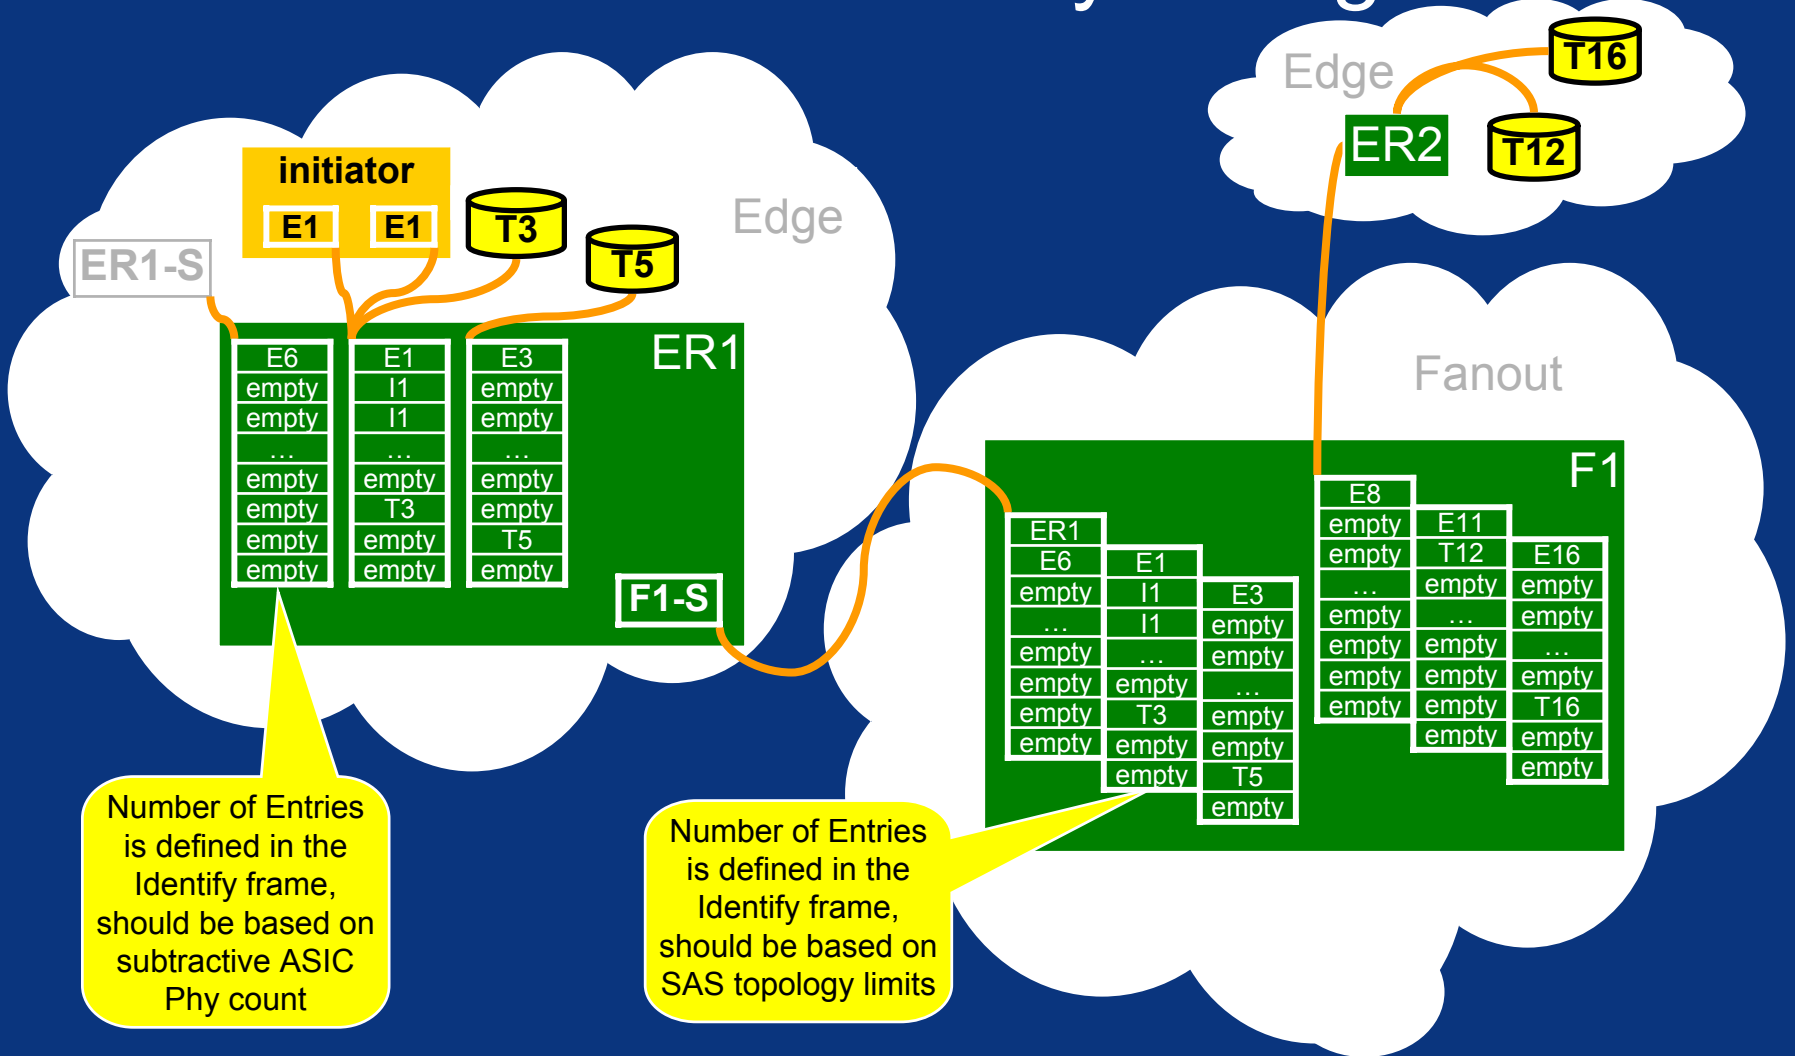
configuration complete when all phys have been traversed

Initiator uses SMP Configure requests to configure route for edge route or fanout device, repeat this sequence iterating by slot and index until all route information on the device is configured

fanout discovery and configuration

- fanout expander device is discovered and configured in the same manner as the edge route expander device
- major difference is the increased size of the route table

initiator fanout discovery/configuration



initiator discovery and configuration algorithm

- reference the psuedo 'C' code in:

ExpanderConfiguration.CPP

- code illustrates a representative algorithm for discovery and configuration
- to work in a multi-initiator environment, the route table entries must be uniquely identified, regardless of where in the topology the initiator is attached

Phy control

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
- used to control Phy specific functions

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (90h) – Phy Control							
2-3	Reserved							
4	Phy Identifier							
5	Phy Operation							
6	Min Phy Link Rate				Max Phy Link Rate			
7	Reserved							
8-11	CRC							
0-3	EOF							



	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Phy Control Result							
5-7	Reserved							
8-11	CRC							
0-3	EOF							

Phy margin control

	7	6	5	4	3	2	1	0
0-3	SOAF							
0-31	Open Frame							
0-3	EOAF							

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (40h) – SMP Request							
1	SMP Function (91h) – Phy Margin Control							
2-3	Reserved							
4	Phy Identifier							
5	Reserved							
6-7	Vendor Unique							
8-11	CRC							
0-3	EOF							

- definition presented is equivalent to sas-r00c specification, using new request/reponse format
 - added vendor unique field in response
- used to control Phy specific functions

	7	6	5	4	3	2	1	0
0-3	SOF							
0	Information Unit Type (41h) – SMP Response							
1	SMP Result (00h) – Function Accepted							
2-3	Reserved							
4	Phy Margin Control Result							
5	Reserved							
6-7	Vendor Unique Result							
8-11	CRC							
0-3	EOF							



i n v e n t