

To: T10 Technical Committee
From: Steven Fairchild, HP (steve.Fairchild@hp.com)
Date: 16 July 2002
Subject: Proposal to allow an initiator based configuration of the SAS topology
Doc #: 02-279

Revision r1:

- Identify address frame now includes two SAS addresses, one for this phy, and one for the far end phy, placed in the same positions as they occur in the Open address frame. The Discover, Report Route Information, and Configure Route Information has the same change.
- Identify address frame has wording that indicates [prog] min/max phy link rate field(s) may be reported as rate unknown.

This is a multi-part proposal that:

- Modifies the Identify Frame format.
- Modifies the SMP Request/Response format.
- Modifies the SMP Discover function to re-use the Identify Frame payload and support a shift in discovery from a “by cloud” mechanism to a “by Phy” mechanism.
- Modifies the SMP Report General function to remove the bit mask fields and add route fields.
- Deletes the SMP functions; Report Phy and Report Phy Device Names, because their information is provided in the modified Discover payload.
- Adds the SMP functions; Report Route Information and Configure Route Information, to facilitate the implementation of an initiator based topology discovery.
- Add support for a new device type, an Edge Router, which can be used to simplify the building of an edge cloud.
- Add discussion on guidelines for an initiator to discover and configure the topology.

The purpose for the recommended changes is to reduce the complexity of the expander devices by shifting the burden of topology discovery and configuration to one or more initiators. While not all the changes are required to accomplish the initiator-based configuration, in total they simplify the topology discovery. The details of the changes follow.

Modify the Identify Frame format

from:

Byte	7	6	5	4	3	2	1	0
0	Reserved			ADDRESS FRAME TYPE (0h)				
1	PHY IDENTIFIER							
2	Reserved			MAXIMUM PHYSICAL LINK RATE				
3	DEVICE TYPE	STP INITIATOR	STP TARGET	SSP INITIATOR	SSP TARGET	SMP INITIATOR	SMP TARGET	
4	DEVICE NAME							
11								
12	(MSB)	Reserved						(LSB)
27								
28	(MSB)	CRC						(LSB)
31								

to:

Byte	7	6	5	4	3	2	1	0
0	Reserved			ADDRESS FRAME TYPE (0h)				
1	Reserved							
2	PHY IDENTIFIER							
3	Reserved							
4								
11	CONNECTED TO SAS ADDRESS							
12								
19	PHY SAS ADDRESS							
20	ROUTE	ADDRESS DECODE		Reserved		DEVICE TYPE		
21	Reserved	LINK STATE		CURRENT PHY LINK RATE				
22	PROG MIN PHY LINK RATE			MIN PHY LINK RATE				
23	PROG MAX PHY LINK RATE			MAX PHY LINK RATE				
24	Reserved			SSP INITIATOR	STP INITIATOR	SMP INITIATOR	Reserved	
25	Reserved			SSP TARGET	STP TARGET	SMP TARGET	SATA TARGET	
26	(MSB)	Reserved						(LSB)
27								
28	(MSB)	CRC						(LSB)
31								

The CONNECTED TO SAS ADDRESS field is a placeholder for the SAS ADDRESS of the device connected to this phy. This field will be 0 filled for an Identify Frame exchange.

The PHY SAS ADDRESS field contains the device name of this phy.

The DEVICE TYPE field indicates the type of device containing the phy, and is defined in Table 1. A device, which is capable of being both an end device and an expander device, shall only report its expander device type in this field.

Table 1. Device Types

DEVICE TYPE	Device type
000b	End device only
001b	Edge expander device
010b	Fanout expander device
011b	Edge route expander device
100b-111b	Reserved

The ADDRESS DECODE field indicates the method of address decode support by this phy and is defined in Table 2. A device that is capable of supporting different decode mechanisms internally. A device which is capable of supporting multiple decode mechanisms should report the most capable method.

Table 2. Address Decode

ADDRESS DECODE	Address Decode
000b	None
001b	Subtractive
010b	Table
011b	Auto, self-discovery
100b-111b	Reserved

The ROUTE field indicates the content of the frame is valid for routing. The value will always be 1b for an Identify frame.

The CURRENT PHY LINK RATE field indicates the current physical link rate negotiated on this phy and is defined in Table 3.

Table 3. Current Physical Link Rate

CURRENT PHY LINK RATE	Current physical link rate
0000b – 0010b	Reserved
0011b	1,5 Gbps
0100b	3,0 Gbps
0101b – 1111b	Reserved

The LINK STATE field indicates the current state of the physical link on this phy and will always be 000b for an Identify frame.

The MIN PHY LINK RATE field indicates the minimum physical link rate supported on this phy and is defined in Table 4. For an Identify Frame this field may contain 0001b, indicating the rate is unknown.

Table 4. Physical Link Rate

PHY LINK RATE	Physical Link Rate
0000b	Reserved for SMP frames
0001b	Rate unknown
0010b	Reserved for SMP frames
0011b	1,5 Gbps
0100b	3,0 Gbps
0101b – 1111b	Reserved

The PROG MIN PHY LINK RATE field indicates the minimum physical link rate programmed on this phy and is defined in Table 4. For an Identify Frame this field may contain 0001b, indicating the rate is unknown.

The MAX PHY LINK RATE field indicates the maximum physical link rate supported on this phy and is defined in Table 4. For an Identify Frame this field may contain 0001b, indicating the rate is unknown.

The PROG MAX PHY LINK RATE field indicates the maximum physical link rate programmed on this phy and is defined in Table 4. For an Identify Frame this field may contain 0001b, indicating the rate is unknown.

The SMP INITIATOR bit indicates the device is an SMP initiator device or SMP target/initiator device.

The STP INITIATOR bit indicates the device is an STP initiator device or STP target/initiator device.

The SSP INITIATOR bit indicates the device is an SSP initiator device or SSP target/initiator device.

The SATA TARGET bit indicates the device is an SATA target device or SATA target/initiator device.

The SMP TARGET bit indicates the device is an SMP target device or SMP target/initiator device.

The STP TARGET bit indicates the device is an STP target device or STP target/initiator device.

The SSP TARGET bit indicates the device is an SSP target device or SSP target/initiator device.

Modify the SMP Request/Response Format

from:

Byte	7	6	5	4	3	2	1	0	
0	INFORMATION UNIT TYPE (40h)								
1	Reserved								
23	Reserved								
24	FUNCTION								
25	ADDITIONAL REQUEST BYTES								
m	ADDITIONAL REQUEST BYTES								
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

to:

Byte	7	6	5	4	3	2	1	0	
0	INFORMATION UNIT TYPE (40h)								
1	FUNCTION								
2	Reserved								
3	Reserved								
4	ADDITIONAL REQUEST BYTES								
m	ADDITIONAL REQUEST BYTES								
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

and from:

Byte	7	6	5	4	3	2	1	0	
0	INFORMATION UNIT TYPE (41h)								
1	Reserved								
23	Reserved								
24	FUNCTION RESULT								
25	ADDITIONAL RESPONSE BYTES								
m	ADDITIONAL RESPONSE BYTES								
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

to:

Byte	7	6	5	4	3	2	1	0	
0	INFORMATION UNIT TYPE (41h)								
1	FUNCTION RESULT								
2	Reserved								
3	Reserved								
4	ADDITIONAL RESPONSE BYTES								
m	ADDITIONAL RESPONSE BYTES								
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

Note: The reason for the change is to reduce the overall size of the SMP payload. This change is not critical to the overall proposal. If accepted, then all SMP request and response frames will need to be adjusted. The editing involved is not reflected in this proposal.

Modify the SMP Discover Format

Editorial note: The current version of SAS drops the header bytes in the presentation of the functions. This makes it look like the fields in the frames are not on appropriate word boundaries. Would like to suggest that the single byte that is the information unit type be included in the descriptions so that boundaries are presented properly.

from:

Table x defines the request format.

Table x. DISCOVER request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (00h)							
1	Reserved							
3	Reserved							

Table y defines the response format.

Table y. DISCOVER response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
3	Reserved							
4	(MSB)	DEVICE NAME VALID BITMASK						(LSB)
11	Reserved							
12	(MSB)	ATTACHED FANOUT EXPANDER BITMASK						(LSB)
19	Reserved							
20	Reserved							
31	Reserved							
32	(MSB)	DEVICE NAME 0						(LSB)
39	Reserved							
...	...							
536	(MSB)	DEVICE NAME 63						(LSB)
543	Reserved							
544	(MSB)	CRC						(LSB)
547	Reserved							

to:

Table x defines the request format.

Table x. DISCOVER request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (00h)							
1	Reserved							
2	Reserved							
3	PHY IDENTIFIER							
4	Reserved							
6	Reserved							

The PHY IDENTIFIER field indicates the physical link for which the Discover information is being requested.

Table y defines the response format.

Table y. DISCOVER response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
2	Reserved							
3	Reserved				ADDRESS FRAME TYPE (0h)			
4	Reserved							
5	PHY IDENTIFIER							
6	Reserved							
7	Reserved							
14	CONNECTED TO SAS ADDRESS							
15	Reserved							
22	PHY SAS ADDRESS							
23	ROUTE	ADDRESS DECODE			Reserved	DEVICE TYPE		
24	Reserved	LINK STATE			CURRENT PHY LINK RATE			
25	PROG MIN PHY LINK RATE				MIN PHY LINK RATE			
26	PROG MAX PHY LINK RATE				MAX PHY LINK RATE			
27	Reserved				SSP INITIATOR	STP INITIATOR	SMP INITIATOR	Reserved
28	Reserved				SSP TARGET	STP TARGET	SMP TARGET	SATA TARGET
29	(MSB)	Reserved						(LSB)
30	Reserved							
31	(MSB)	CRC						(LSB)
34	Reserved							

The ADDRESS FRAME TYPE field is retained for commonality with the Identify frame.

The PHY IDENTIFIER field indicates the physical link for which the Discover information is being requested.

The CONNECTED TO SAS ADDRESS field contains the device name for the far end phy.

The PHY SAS ADDRESS field contains the device name of this phy.

The DEVICE TYPE field indicates the type of device containing the phy, and is defined in Table 1. A device, which is capable of being both an end device and an expander device, shall only report its expander device type in this field.

Table 1. Device Types

DEVICE TYPE	Device type
000b	End device only
001b	Edge expander device
010b	Fanout expander device
011b	Edge route expander device
100b-111b	Reserved

The ADDRESS DECODE field indicates the method of address decode support by this phy and is defined in Table 2. A device which is capable of supporting multiple decode mechanisms should report the most capable method. Auto decode is the most capable, Subtractive decode is the least capable.

Table 2. Address Decode

ADDRESS DECODE	Address Decode
000b	None
001b	Subtractive
010b	Table
011b	Auto, self-discovery
100b-111b	Reserved

The ROUTE field indicates the content of the frame is valid for routing. The 0b value indicates the route has been determined to be in violation of connection rules.

The CURRENT PHY LINK RATE field indicates the current physical link rate negotiated on this phy and is defined in Table 3.

Table 3. Current Physical Link Rate

CURRENT PHY LINK RATE	Current physical link rate
0000b	Phy does not exist
0001b	Rate unknown
0010b	Phy disabled
0011b	1,5 Gbps
0100b	3,0 Gbps
0101b – 1111b	Reserved

The LINK STATE field indicates the current state of the physical link on this phy.

Table 4. Link State

LINK STATE	Link State
000b	Active
001b	Inactive
010b	Failed
011b	OOB in Progress
100b	Spinup Hold OOB
101b	Release Spinup Hold OOB
110b – 111b	Reserved

The MIN PHY LINK RATE field indicates the minimum physical link rate supported on this phy and is defined in Table 5.

Table 5. Physical Link Rate

PHY LINK RATE	Physical Link Rate
0000b – 0010b	Reserved
0011b	1,5 Gbps
0100b	3,0 Gbps
0101b – 1111b	Reserved

The PROG MIN PHY LINK RATE field indicates the minimum physical link rate programmed on this phy and is defined in Table 5.

The MAX PHY LINK RATE field indicates the maximum physical link rate supported on this phy and is defined in Table 5.

The PROG MAX PHY LINK RATE field indicates the maximum physical link rate programmed on this phy and is defined in Table 5.

The SMP INITIATOR bit indicates the device is an SMP initiator device or SMP target/initiator device.

The STP INITIATOR bit indicates the device is an STP initiator device or STP target/initiator device.

The SSP INITIATOR bit indicates the device is an SSP initiator device or SSP target/initiator device.
The SATA TARGET bit indicates the device is an SATA target device or SATA target/initiator device.
The SMP TARGET bit indicates the device is an SMP target device or SMP target/initiator device.
The STP TARGET bit indicates the device is an STP target device or STP target/initiator device.
The SSP TARGET bit indicates the device is an SSP target device or SSP target/initiator device.

Modify the SMP Report General format

from:

Table x defines the request format.

Table x. REPORT GENERAL request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (01h)							
1	Reserved							
3								

Table y defines the response format.

Table y. REPORT GENERAL response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
2	NUMBER OF PHYS							
3	INPUT PHY IDENTIFIER							
4	Reserved							
15								
16	(MSB)	ACTIVE PHY BITMASK						(LSB)
23								
24	(MSB)	ATTACHED FANOUT EXPANDER BITMASK						(LSB)
31								
32	(MSB)	ATTACHED EDGE EXPANDER BITMASK						(LSB)
39								
40	(MSB)	ATTACHED SAS INITIATOR BITMASK						(LSB)
47								
48	(MSB)	ATTACHED SAS TARGET BITMASK						(LSB)
53								
54	(MSB)	ATTACHED SATA BITMASK						(LSB)
75								
76	(MSB)	PHY RATE MULTIBITMASK						(LSB)
83								
84	(MSB)	FUNCTIONS SUPPORTED BITMASK						(LSB)
115								

to:

Table x defines the request format.

Table x. REPORT GENERAL request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (01h)							
1	Reserved							
6								

Table y defines the response format.

Table y. REPORT GENERAL response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
2	Reserved							
3	NUMBER OF PHYS							
4	INPUT PHY IDENTIFIER							
5	MAXIMUM ROUTE SLOT							
6	MAXIMIM ROUTE INDEX							
7	(MSB)	CRC						(LSB)
10								

The NUMBER OF PHYS field contains the number of phys in the device.

The INPUT PHY IDENTIFIER field contains the identifier of the phy through which the REPORT GENERAL request was received.

The MAXIMUM ROUTE SLOT field contains the maximum number of route slots for an edge router or fanout device. The number of route slots shall be at least equal to the number of phys on the device.

The MAXIMUM ROUTE INDEX field contains the maximum number of route indexes for an edge router or fanout device. The number of route indexes shall be greater than or equal to the number of phys on the device for an edge router. The number of route indexes shall be greater than or equal to the maximum supported devices in a single edge cloud for a fanout device (currently this is 64).

Delete the SMP functions; Report Phy and Report Phy Devices

Note: These functions are no longer needed; because their functionality has been consolidate into the Identify frame and Discover response.

Add the SMP function: Report Route Information

The REPORT ROUTE INFORMATION function returns the route table information for a specific route slot and route index within an edge route expander device or fanout expander device. Edge expander devices or end devices do not need to support this function. This function is used primarily as a diagnostic tool to resolve topology routing issues.

Table x defines the request format.

Table x. REPORT ROUTE INFORMATION request

Byte	7	6	5	4	3	2	1	0
0	FUNCTION (04h)							
1	Reserved							
2	Reserved							
3	ROUTE SLOT							
4	ROUTE INDEX							
5	Reserved							
6	Reserved							

The ROUTE SLOT field indicates the route slot for which the Report Route information is being requested. The value must be in the range of 0 to MAXIMUM ROUTE SLOT or a function reject response shall occur.

The ROUTE INDEX field indicates the route index for which the Report Route information is being requested. The value must be in the range of 0 to MAXIMUM ROUTE INDEX or a function reject response shall occur.

Table y defines the response format.

Table y. REPORT ROUTE INFORMATION response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
2	Reserved							
3	ROUTE SLOT							
4	ROUTE INDEX							
5	PHY IDENTIFIER							
6	Reserved							
7	CONNECTED TO SAS ADDRESS							
14	Reserved							
15	Reserved							
22	Reserved							
23	ROUTE	Reserved						
24	Reserved			CURRENT PHY LINK RATE				
25	Reserved							
26	Reserved							
27	Reserved			SSP INITIATOR	STP INITIATOR	SMP INITIATOR	Reserved	
28	Reserved			SSP TARGET	STP TARGET	SMP TARGET	SATA TARGET	
29	(MSB)	Reserved						(LSB)
30	Reserved							
31	(MSB)	CRC						(LSB)
34	Reserved							

The ROUTE SLOT field indicates which slot in a route table is being returned. This value shall match the value in the request.

The ROUTE INDEX field indicates which index in a route table is being returned. This value shall match the value in the request.

The PHY IDENTIFIER field indicates the physical link referenced by this route table entry.

The CONNECTED TO SAS ADDRESS field contains the device name reference by this route table entry.

The ROUTE field indicates the content of the frame referenced by this route table entry is valid for routing. The 0b value indicates the route has been determined to be in violation of connection rules.

The CURRENT PHY LINK RATE field indicates the current physical link rate referenced by this route table entry.

The SMP INITIATOR bit indicates the device referenced by this route table entry is an SMP initiator device or SMP target/initiator device.

The STP INITIATOR bit indicates the device referenced by this route table entry is an STP initiator device or STP target/initiator device.

The SSP INITIATOR bit indicates the device referenced by this route table entry is an SSP initiator device or SSP target/initiator device.

The SATA TARGET bit indicates the device referenced by this route table entry is an SATA target device or SATA target/initiator device.

The SMP TARGET bit indicates the device referenced by this route table entry is an SMP target device or SMP target/initiator device.

The STP TARGET bit indicates the device referenced by this route table entry is an STP target device or STP target/initiator device.

The SSP TARGET bit indicates the device referenced by this route table entry is an SSP target device or SSP target/initiator device.

The SSP TARGET bit reflects the value to be saved in this route table entry to indicate whether the device is an SSP target device or SSP target/initiator device.

Table y defines the response format.

Table y. CONFIGURE ROUTE INFORMATION response

Byte	7	6	5	4	3	2	1	0
0	FUNCTION RESULT							
1	Reserved							
2	Reserved							
3	ROUTE SLOT							
4	ROUTE INDEX							
5	PHY IDENTIFIER							
6	Reserved							
7	CONNECTED TO SAS ADDRESS							
14	Reserved							
15	Reserved							
22	Reserved							
23	ROUTE	Reserved						
24	Reserved			CURRENT PHY LINK RATE				
25	Reserved							
26	Reserved							
27	Reserved			SSP INITIATOR	STP INITIATOR	SMP INITIATOR	Reserved	
28	Reserved			SSP TARGET	STP TARGET	SMP TARGET	SATA TARGET	
29	(MSB)	Reserved						(LSB)
30	Reserved							
31	(MSB)	CRC						(LSB)
34	Reserved							

The ROUTE SLOT field indicates which slot in a route table is being returned. This value shall match the value in the request.

The ROUTE INDEX field indicates which index in a route table is being returned. This value shall match the value in the request.

The PHY IDENTIFIER field indicates the physical link referenced by this route table entry.

The CONNECTED TO SAS ADDRESS field contains the device name reference by this route table entry.

The ROUTE field indicates the content of the frame is valid for routing. The 0b value indicates the route has been determined to be in violation of connection rules.

The CURRENT PHY LINK RATE field indicates the current physical link rate negotiated on this phy and is defined in Table 1.

Table 1. Current Physical Link Rate

CURRENT PHY LINK RATE	Current physical link rate
0000b	Phy does not exist
0001b	Rate unknown
0010b	Phy disabled
0011b	1,5 Gbps
0100b	3,0 Gbps
0101b – 1111b	Reserved

The SMP INITIATOR bit indicates the device is an SMP initiator device or SMP target/initiator device.

The STP INITIATOR bit indicates the device is an STP initiator device or STP target/initiator device.

The SSP INITIATOR bit indicates the device is an SSP initiator device or SSP target/initiator device.

The SATA TARGET bit indicates the device is an SATA target device or SATA target/initiator device.

The SMP TARGET bit indicates the device is an SMP target device or SMP target/initiator device.

The STP TARGET bit indicates the device is an STP target device or STP target/initiator device.

The SSP TARGET bit indicates the device is an SSP target device or SSP target/initiator device.

Add a new device type: Edge Router

edge route expander device: An edge expander device that presents simple routing capability to a fanout expander device, but complex routing capability to an edge expander. This class of expander device may be used to connect edge expanders and create larger topologies.

change section 4.1.7 from:

There are two classes of expander devices:

- a) fanout expander devices: expanders that contain complex routing capability; and
- b) edge expander devices: expanders that only contain simple routing capability.

to:

There are three classes of expander devices:

- a) fanout expander devices: expanders that contain complex routing capability; and
- b) edge expander devices: expanders that only contain simple routing capability.
- c) edge route expander devices: expanders that contain a hybrid routing capability.

change section 4.1.9 from:

No more than one fanout expander device shall be included in a SAS domain. The fanout expander device may be attached to up to 64 edge expanders, initiator ports, or target ports.

Each edge expander device may be attached to no more than one fanout expander device. Each edge expander device may be attached to up to 64 initiator ports or target ports.

An edge expander device may be attached to another edge expander device only if there are no other expanders in the SAS domain.

to:

No more than one fanout expander device shall be included in a SAS domain. The fanout expander device may be attached to up to 64 edge expanders, initiator ports, or target ports.

Each edge expander device may be attached to no more than one fanout expander device. Each edge expander device may be attached to up to 64 initiator ports or target ports.

An edge expander device may be attached to another edge expander device only if there are no other expanders in the SAS domain.

An edge route expander device may be attached to no more than one fanout expander device, and up to N edge expanders, initiator ports, or target ports. If connected to a fanout expander device and one or more edge expanders, the edge route expander will present itself as an edge expander to the fanout expander device and as a fanout expander device to each edge expander device.

change section 4.8.4.1:

Fanout expander devices must perform additional discovery for each attached edge expander device in order to establish a routing table.

to:

Fanout expander devices and edge route devices require additional discovery for each attached edge expander device in order to establish a routing table.

change section 7.6.3:

After learning that it is attached to an edge expander device, a fanout expander device shall use the SMP DISCOVER function (see ???) to retrieve the list of device names to which the edge expander device is attached.

After receiving a CHANGE primitive sequence from an edge expander device, the fanout expander device shall use the SMP DISCOVER function to obtain an updated list of device names from that edge expander device.

to:

A fanout expander device may require configuration by an initiator before it can perform its routing function. As an initiator traverses the topology it shall use the SMP DISCOVER function to build the necessary routing information, which shall then be used to configure the fanout expander device. The configuration information is presented to the fanout expander device using multiple SMP CONFIGURE ROUTE INFORMATION functions.

After receiving a CHANGE primitive sequence from an expander, the initiator rebuilds the topology to update the list of SAS ADDRESSes for the fanout expander device.

7.6.4 Edge route expander device specific rules

An edge route expander device may require configuration by an initiator before it can perform its routing function. As an initiator traverses the topology it shall use the SMP DISCOVER function to build the necessary routing information, which shall then be used to configure the edge route expander device. The configuration information is presented to the edge route expander device using multiple SMP CONFIGURE ROUTE INFORMATION functions.

After receiving a CHANGE primitive sequence from an expander, the initiator rebuilds the topology to update the list of SAS ADDRESSes for the edge route expander device.

Guidelines for an initiator to discover and configure the SAS topology

To simplify the edge expander devices, edge route expander devices and fanout expander devices within the SAS topology, it is desirable to provide a mechanism for initiators to configure the topology. The algorithm presented defines a method where any and all initiators in the topology may configure the topology without requiring any form of coordination between the initiators.

There are two major components of the algorithm; the method used to traverse the topology and the mechanism used to update the route information in the topology.

The initiator traverses the topology by entering each expander and accessing each phy using the SMP Discover function in an ascending order. When the last phy in the current expander is encountered, the initiator enters the next expander encountered on the next sequential phy and begins the process again. The result is to group end devices by expander. The SAS addresses compiled are then organized without collapsing empty or duplicate phys.

Once the SAS addresses are compiled, the initiator shall update each of the edge route expander devices and fanout expander devices within the topology. There is no configuration required for the edge expanders. The initiator shall update the near edge route expander device prior to exiting to configure the fanout expander device and far edge route expander devices.

Within each edge route expander device or fanout expander device an organized route table shall exist that the initiator will update to establish the topology route. To avoid issues with multiple initiators having to coordinate the update operation, the route table shall be updated identically, independent of which initiator could perform the operation. Since the time to complete a discovery and configuration cycle is relatively small, there is no need to prevent an initiator from participating.

The mechanism used to update the edge route expander device and fanout expander device is the SMP Configure Route Information function. The content of the request payload in SMP Configure Route Information function shall be consistent regardless of which initiator makes the request.