

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 29 June 2007
 Subject: 07-300r0 SAS-2 Configurable expander table-to-table error handling

Revision history

Revision 0 (29 June 2007) First revision

Related documents

sas2r10 - Serial Attached SCSI - 2 (SAS-2) revision 10

Overview

SAS-2 is unclear whether or not discover processes are expected to enter into an externally configurable expander's routing table the SAS address of an expander attached incorrectly with a table routing phy attached to another table routing phy. SAS-1.1 declared that as an error to be reported "in a vendor-specific manner," but wasn't clear about the state in which the routing table was expected to be left. Since table-to-table routing was added (for self-configuring expanders), the figure has shown the address not being added.

Addresses beyond that expander are clearly not added.

Without standardizing this behavior, discover processes may handle this error case inconsistently, with some adding the address and some not doing so. If that were the only impact, it wouldn't matter much; unfortunately, this leads to uncertainty about where to place subsequent addresses in the routing table.

This revision of this proposal proposes that the expander device's SAS address **NOT** be added.

Suggested changes

4.8.3 Route table optimization

The management application client shall support a route table optimization that reduces the number of entries required in an expander route table in an externally configurable expander device. The method used to enable and disable the route table optimization is vendor specific.

If the route table optimization is enabled, then the management application client shall exclude discovered SAS addresses from the expander route table when any of the following conditions are met:

- a) in the SMP DISCOVER response (see 10.4.3.8) for the discovered phy:
 - A) the FUNCTION RESULT field is set to a non-zero value (i.e., not SMP FUNCTION ACCEPTED);
- b) in the SMP DISCOVER response for the discovered phy:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table); and
 - C) the ATTACHED DEVICE TYPE field is set to zero (i.e., no device attached);
- c) in the SMP DISCOVER response for the discovered phy:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);
 - C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device or expander device); and
 - D) the ATTACHED SAS ADDRESS field contains the SAS address of the expander device being configured (i.e. a self-referencing address);
- d) in the SMP DISCOVER response for the discovered phy:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);
 - C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device or expander device); and
 - D) the ATTACHED SAS ADDRESS field contains the SAS address of a device directly attached to the expander device being configured;

or

- e) in the SMP DISCOVER response for the discovered phy:
 - A) the FUNCTION RESULT field is set to zero (i.e., SMP FUNCTION ACCEPTED);
 - B) the ROUTING ATTRIBUTE field is set to 1h (i.e., subtractive) or 2h (i.e., table);
 - C) the ATTACHED DEVICE TYPE field is set to a non-zero value (e.g., end device or expander device); and
 - D) the ATTACHED SAS ADDRESS field contains a SAS address that already exists in the expander route table.

If the discovered SAS address being included in the expander route table is for a device that is currently not attached (i.e., the ATTACHED DEVICE TYPE field is set to zero (i.e., no device attached) and the ROUTE ATTRIBUTE field is set to 0h (i.e., direct)), then the entry shall be inserted with the ROUTED SAS ADDRESS field set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit set to one (see 10.4.3.25).

If route table optimization is disabled, then all SAS addresses shall be qualified for insertion in the expander route table.

If the management application client supports route table optimization, then the management application client should provide a vendor-specific method for initiating a check of the resulting expander route tables. The check should be performed under the following situations:

- a) when an I_T nexus loss occurs for a destination port that is expected to be present;
- b) when a discover process has been completed;
- c) when another SMP initiator port is discovered in the SAS domain; or
- d) when a self-configuring expander device is discovered in the SAS domain.

If the management application client detects an inconsistency in the expander route tables when route table optimization is enabled (e.g., detects entries that appear to have been filled in by a discover process with route table optimization disabled), then the management application client shall report an error in a vendor-specific manner and shall disable route table optimization. The management application client should then re-initiate a discover process with route table optimization disabled.

4.8.4 Expander route index order

The expander route table shall be configured for each expander phy that has a table routing attribute.

If the phy is not attached to an expander device, every expander route entry for that phy shall be disabled (i.e., the ROUTED SAS ADDRESS field shall be set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit shall be set to one).

If the phy is attached to an expander device, the expander route table shall be configured for that phy as follows. For purposes of configuring the expander route table for that phy, the expander devices attached to the expander phy are assigned levels:

- 1) the expander device in which the expander route table is being configured is level 0;
- 2) the attached expander device is considered level 1;
- 3) devices attached to the level 1 expander device, except for the level 0 expander device, are considered level 2;
- 4) devices attached to level 2 expander devices, except for level 1 expander devices, are considered level 3; and
- 5) for each n greater than 3, devices attached to level n-1 expander devices, except for level n-2 expander devices, are considered level n.

The expander route table for each expander phy shall be configured starting from expander route index 0 by level (e.g., if there are three levels, then all level 1 entries first, then all level 2 entries, then all level 3 entries) up to the value of the EXPANDER ROUTE INDEXES field reported by the SMP REPORT GENERAL function (see 10.4.3.3).

Assuming the level 1 expander device has expander phys attached to N phys with qualified SAS addresses (see 4.8.3), the first N entries shall be used for those SAS addresses in expander phy order (i.e., the addresses attached to lower expander phy numbers first).

For each of the level 2 devices that:

- a) is an expander device attached to M phys with qualified SAS addresses; and
- b) is attached to an expander phy in the level 1 expander device with the table routing attribute,

the next M entries shall be used for the level 2 expander device's qualified SAS addresses in expander phy order (i.e., lower phy numbers first).

This process shall repeat for all levels of expander devices.

SAS addresses of devices attached beyond expander phys that are attached table-to-table shall not be included in the expander route table. This includes the SAS address of the first expander device that is attached table-to-table and the SAS address of every device attached beyond that expander device.

NOTE 1 - Not including those SAS addresses provides compatibility with management application clients compliant with previous versions of this standard. End devices in SAS domains containing externally configurable expander devices and table-to-table attachments may not be able to establish connections to each other. End devices in SAS domains containing only self-configuring expander devices compliant with this standard (i.e., supporting table-to-table attachments) are able to establish connections to any other end device.

Figure 48 shows an example of a route table that does not include SAS addresses beyond a table-to-table attachment.

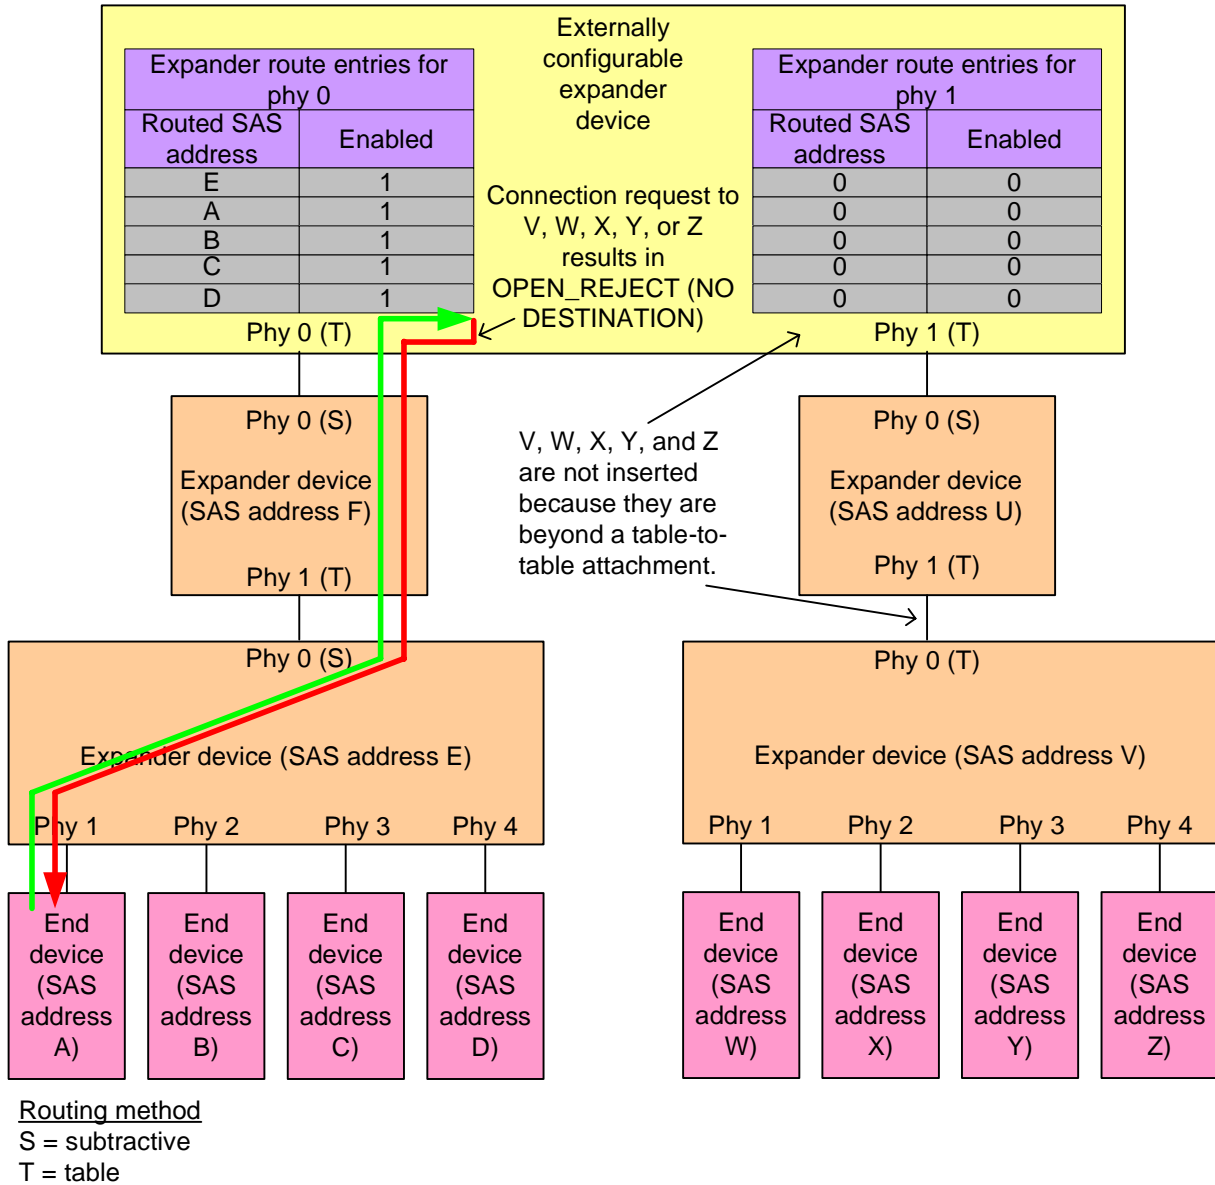


Figure 48 — Externally configurable expander device and table-to-table attachment [\[no change\]](#)

After the expander route table has been configured with entries for all levels of expander devices, all remaining expander route entries, if any, shall be disabled (i.e., the ROUTED SAS ADDRESS field shall be set to 00000000 00000000h and the DISABLE EXPANDER ROUTE ENTRY bit shall be set to one). The management application client is not required to disable entries if the topology of expander devices has not changed.

...

Figure 49 shows a portion of a SAS domain, where phy A in expander device R is being configured.

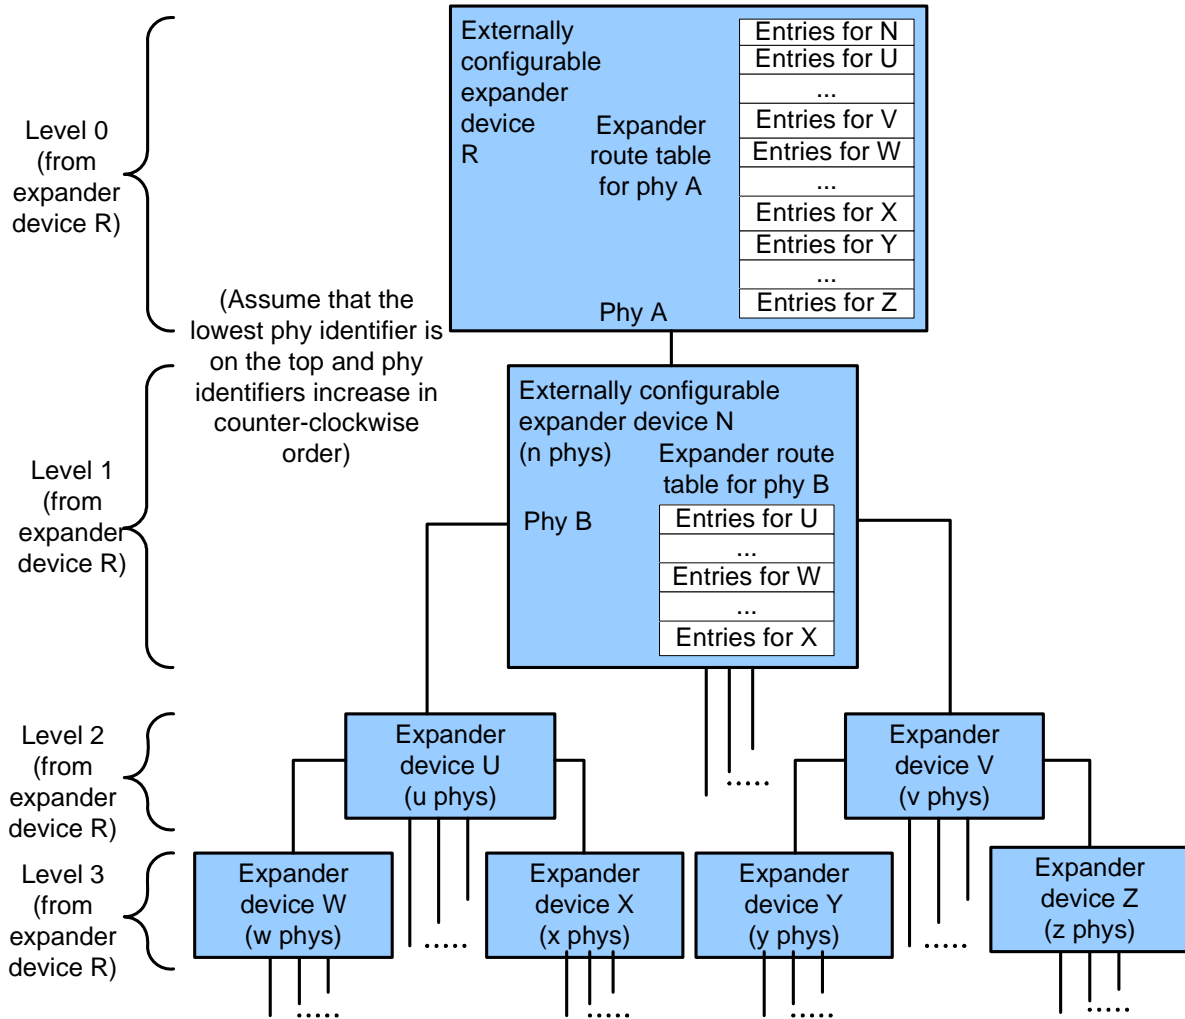


Figure 49 — Expander route index levels example

Table 1 shows how the expander route table is configured for externally configurable expander device R phy A in figure 49.

Table 1 — Expander route table levels for externally configurable expander device R phy A

Expander route index	Expander route entry contents
Level 1 (from device R) entries	
0 .. ($\leq n$ entries)	Qualified SAS addresses attached to expander device N
Level 2 (from device R) entries	
($\leq u$ entries)	Qualified SAS addresses attached to expander device U
...	...additional qualified SAS addresses for expander devices at level 2...
($\leq v$ entries)	Qualified SAS addresses attached to expander device V
Level 3 (from device R) entries	
($\leq w$ entries)	Qualified SAS addresses attached to expander device W
...	...additional qualified SAS addresses for expander devices at level 3...
($\leq x$ entries)	Qualified SAS addresses attached to expander device X
($\leq y$ entries)	Qualified SAS addresses attached to expander device Y
...	...additional qualified SAS addresses for expander devices at level 3...
($\leq z$ entries)	Qualified SAS addresses attached to expander device Z
Entries for additional levels	
...	...
Disabled entries	
...	...

Table 2 shows how the expander route table is configured for externally configurable expander device N phy B in figure 49.

Table 2 — Expander route table levels for externally configurable expander device N

Expander route index	Expander route entry contents
Level 1 (from device N) entries	
($\leq u$ entries)	Qualified SAS addresses attached to expander device U
Level 2 (from device N) entries	
($\leq w$ entries)	Qualified SAS addresses attached to expander device W
...	...additional qualified SAS addresses for expander devices at level 2...
($\leq x$ entries)	Qualified SAS addresses attached to expander device X
Entries for additional levels	
...	...
Disabled entries	
...	...

Figure 50 shows an example topology.

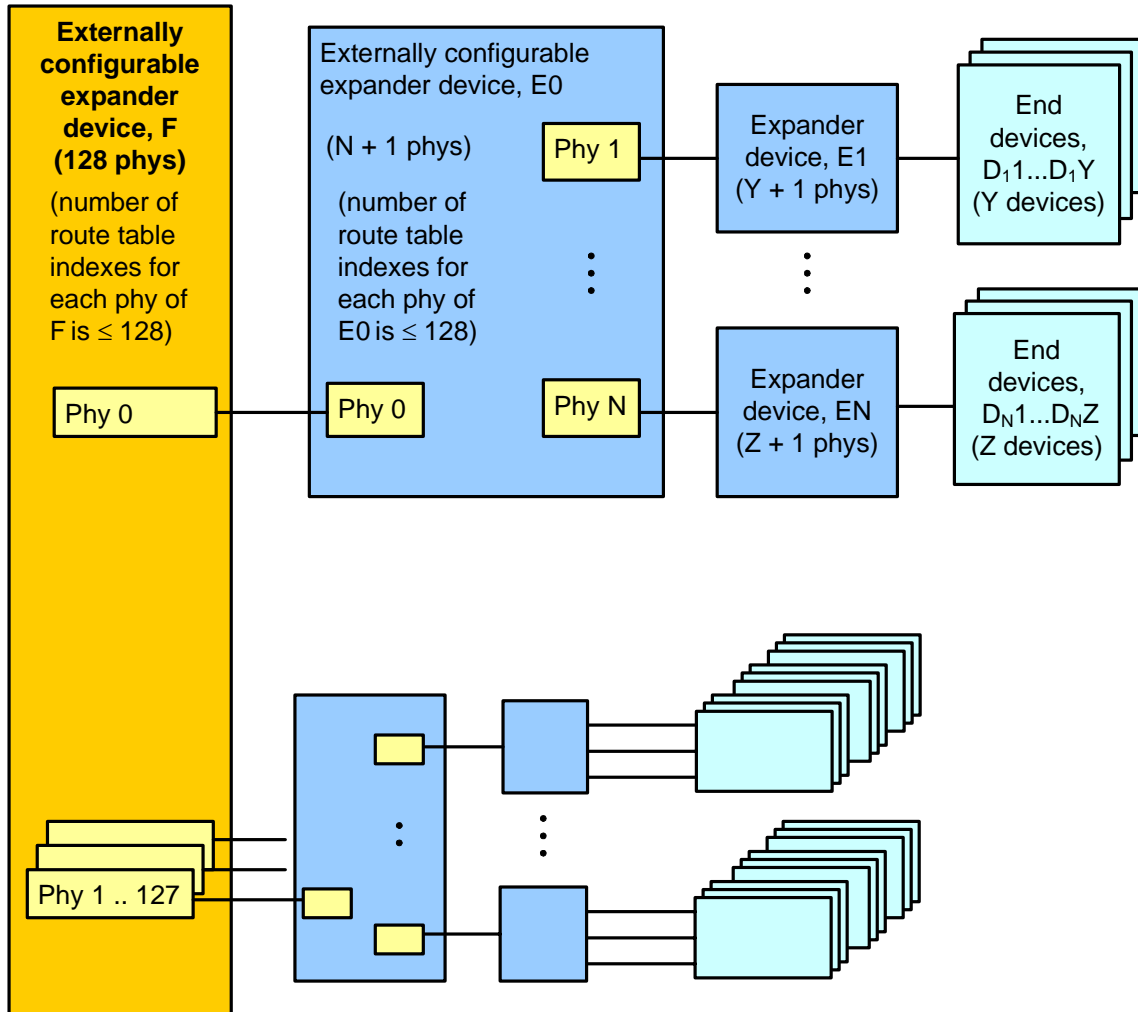


Figure 50 — Expander route index order example

Table 3 shows the expander route index order for externally configurable expander device E0 phy 1 in figure 50.

Table 3 — Expander route entries for externally configurable expander device E0 phy 1

Expander route index	Expander route entry contents
Level 1 entries	
0	SAS address (e.g., D ₁ 1) of the device attached to phy 1 of expander device E1
1	SAS address (e.g., D ₁ 2) of the device attached to phy 2 of expander device E1
...	...
Y - 1	SAS address (e.g., D ₁ Y) of the device attached to phy Y of expander device E1
Level 2 and beyond: no entries	
Disabled entries	

Table 4 shows the expander route index order for externally configurable expander device F phy 0 in figure 50.

Table 4 — Expander route entries for externally configurable expander device F phy 0

Expander route index	Expander route entry contents
Level 1 entries	
0	SAS address (e.g., E1) of the device attached to phy 1 of expander device E0
...	...additional qualified SAS addresses for expander device E0...
N - 1	SAS address (e.g., EN) of the device attached to phy N of expander device E0
Level 2 entries	
N	SAS address (e.g., D ₁ 1) of the device attached to phy 1 of expander device E1
...	...additional qualified SAS addresses for expander device E1...
	SAS address (e.g., D ₁ Y) of the device attached to phy Y of expander device E1
...	...additional qualified SAS addresses for expander devices E2 through EN-1...
	SAS address (e.g., D _N 1) of the device attached to phy 1 of expander device EN
...	...additional qualified SAS addresses for expander device EN...
	SAS address (e.g., D _N Z) of the device attached to phy Z of expander device EN
Level 3 and beyond: no entries since all devices attached to E1 through EN, except for E0, are end devices	
Disabled entries	