

To: T10 Technical Committee
From: Steve Johnson, LSI Logic (sjohnson@lsil.com)
Date: 25 February, 2005
Subject: 05-089r0 SAS-1.1 Discovery Process Missing Route Table Rules

Revision history

Revision 0 (25 February 2005) First revision

Related documents

sas1r07 - Serial Attached SCSI 1.1 revision 7

Overview

Currently there is a discrepancy on how to program route tables between the Discovery Process defined in section 4.6.7.4 of the SAS 1.1 spec revision 7 and Annex K (informative) Discover process example implementation overview. The informative section is correct and as discussed an intended and the spec needs to be updated to match.

Background

One of the goals when creating the optimized route table algorithm was to minimize the disruption of individual phys going up and down and end devices coming and going. Expanders coming and going are considered a major change to the topology and typically require major table rebuilds. We agreed this is unavoidable behavior.

Currently as the spec is written in section 4.6.7.4, If individual links within a wide port go up and down they change the order of the route table by inserting or deleting entries within the table. This causes valid (enabled) entries to move within the route table which can be very disruptive to current IO and something we were trying to avoid.

This can happen with any wide link from expander to end device or expander to expander.

A simple example topology (see below) shows how the table changes when some links (red) go up and down. When the table shrinks (red links came up) entries will not get trashed until they have been previously re-written in a lower index. This case is not disruptive. When the table grows (red links drop) entries will get trashed as new entries overwrite them. This case is disruptive.

Section 4.6.7.4 Discover Process, The rules for what goes into the route table:

“If the discover process optimization is enabled, then discovered SAS addresses shall be qualified for insertion into the expander device route table by the following:

- a) Discovered SAS addresses referencing the expander device being configured (i.e., self-referencing addresses) shall be excluded from the expander device route table;
- b) Discovered SAS addresses directly attached to the expander device being configured shall be excluded from the expander device route table;

- c) The first instance of a discovered SAS address shall be inserted in the configured expander device route table. Subsequent instances of a discovered SAS address shall be excluded from the configured expander device route table; and
- d) Discovered phys shall be inserted in the expander device route table even if the end device is currently not attached and a SAS address is not available. If a SAS address is not available an 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 10.4.3.9)."

The idea behind rule "d" was that hot plug of end devices and resets would not cause the table to change size and therefore not disrupt other connections in the topology. In addition, that individual links changes of wide expander connections would not cause the tables to change.

Virtual PHY Considerations

It is not specified that a Virtual PHY have a direct routing attribute. If it has a subtractive or table attribute than it could change the table size depending on PHY state.

Should we specify in section 4.6.7 "Expander Virtual phys shall have a direct routing attribute."?

Suggested Changes To SAS 1.1 Spec

Update rules in section 4.6.7.4

"If the discover process optimization is enabled, then discovered SAS addresses shall be qualified for insertion into the expander device route table by the following:

- a) Discovered SAS addresses referencing the expander device being configured (i.e., self-referencing addresses) shall be excluded from the expander device route table;
- b) Discovered SAS addresses directly attached to the expander device being configured shall be excluded from the expander device route table;
- c) The first instance of a discovered SAS address shall be inserted in the configured expander device route table. Subsequent instances of a discovered SAS address shall be excluded from the configured expander device route table; and
- d) Discovered phys [with the direct routing attribute](#) shall be inserted in the expander device route table even if the end device is currently not attached and a SAS address is not available. If a SAS address is not available an 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 10.4.3.9)."
- e) [Discovered phys with subtractive and table route attributes that do not have an attached SAS address shall be excluded from the route table.](#)

Example Topology and Route Table

Route table for E1(4-7) All Links Up SAS 1.1 Revision 7 rules				
Route Index		Route Address	PHY Identifier	PHY Owner
0	E	E3	4-7	E2
1	E	T4	8	E2
2	E	T5	9	E2
3	E	T6	10	E2
4	D	0	11	E2
5	E	T7	8	E3
6	E	T8	9	E3
7	D	0	10	E3
8	E	T9	11	E3
9	E	SEP	12	E3

Route table for E1(4-7) Red Links Down SAS 1.1 Revision 7 rules				
Route Index		Route Address	PHY Identifier	PHY Owner
0	D	0	4	E2
1	E	E3	5-7	E2
2	E	T4	8	E2
3	E	T5	9	E2
4	D	0	10	E2
5	D	0	11	E2
6	E	T7	8	E3
7	E	T8	9	E3
8	D	0	10	E3
9	E	T9	11	E3
10	E	SEP	12	E3

Route table for E1(4-7) Red links Up Or Down Updated Rules				
Route Index		Route Address	PHY Identifier	PHY Owner
0	E	E3	4-7	E2
1	E	T4	8	E2
2	E	T5	9	E2
3	E/D	T6/0	10	E2
4	D	0	11	E2
5	E	T7	8	E3
6	E	T8	9	E3
7	D	0	10	E3
8	E	T9	11	E3
9	E	SEP	12	E3

This PHY going down causes table to grow by one entry

Table remains same size with this PHY up or down

