To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 16 July 2008
Subject: 08-183r1 SAS-2 Add device slot numbering fields to DISCOVER

Revision history
Revision 0 (28 March 2008) First revision
Revision 1 (16 July 2008) Incorporated comments from July SAS Protocol WG

Related documents
sas2r14 - Serial Attached SCSI - 2 (SAS-2) revision 14

Overview
This proposal is a SAS-2 letter ballot comment.

Initiator discovery software needs to retrieve enough information about an attached drive to describe the drive to a user, and would prefer to do this based solely on information retrieved by the discover process (using SMP functions). Extended information is available through SCSI Enclosure Services (SES), but that requires many steps:

1.complete the SMP discover process

1) locate the SES logical unit
    1) poll each discovered SAS target device with REPORT LUNS
    2) poll each discovered logical unit with INQUIRY
2) retrieve the Configuration diagnostic page (with a RECEIVE DIAGNOSTIC RESULTS command)
3) retrieve the Enclosure Status diagnostic page (with a RECEIVE DIAGNOSTIC RESULTS command)
4) retrieve the Additional Element Status diagnostic page (with a RECEIVE DIAGNOSTIC RESULTS command)
5) tie the SES information back to the SMP information based on SAS addresses and

just to present a disk drive.

Instead, a few more SES-type fields can be included into the DISCOVER response:

a) Device Slot Number - a number used to identify the bay (i.e., device slot) within the enclosure. Binary number 0-254 (255 means no number is available). If the device slot is managed by SES, this is the value reported in the DEVICE SLOT NUMBER field in the Additional Element Status diagnostic page. If SGPIO is being used, this is reported over SGPIO.
b) Device Slot Group Number - a number used to identify the group of device slots. Binary number 0-254 (255 means no number is available). If SGPIO is being used, this is reported over SGPIO.
c) Device Slot Group Output Connector - a 6-byte ASCII string used to identify the path to the group of device slots (e.g., a switch port number). All ASCII spaces mean no path identifier.

This defers the need to consult with SES for simple topologies (or for topologies that lack SES altogether), allowing SMP management software to identify device slots to users.

These fields are not important enough to expand DISCOVER LIST - they don't contain information needed for expanders performing self-configuration, which is the main customer for DISCOVER LIST.

Different systems will use different terminology for the levels; the key request is that 3 levels be supported.

Figure 1 shows an example of device slot numbering in a topology using SGPIO.

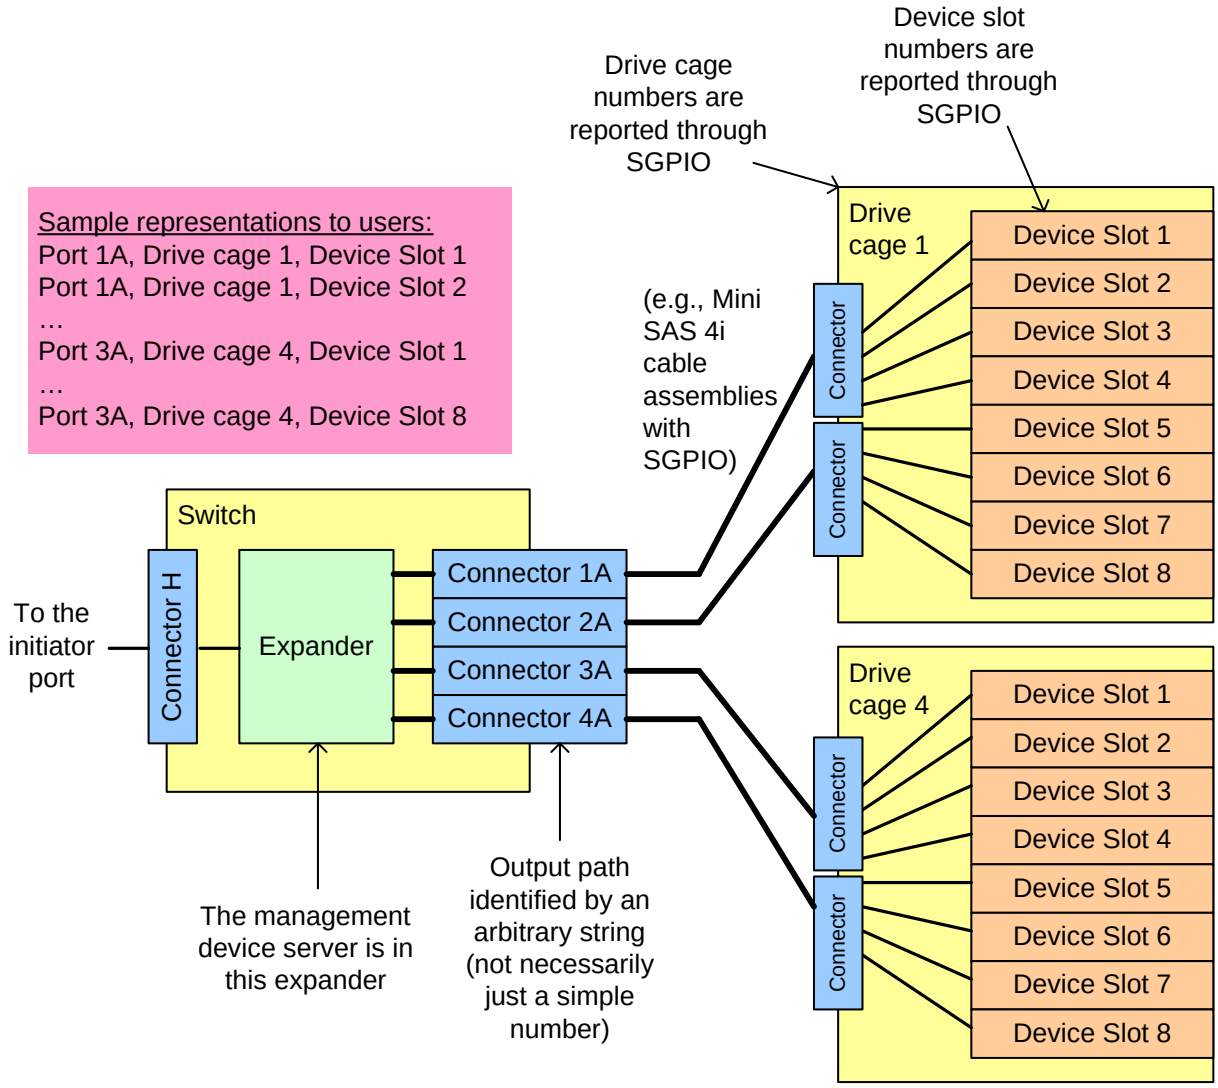

Figure 1 — Device slot number example using SGPIO

Figure 2 shows an example of device slot numbering in a topology not using SGPIO.

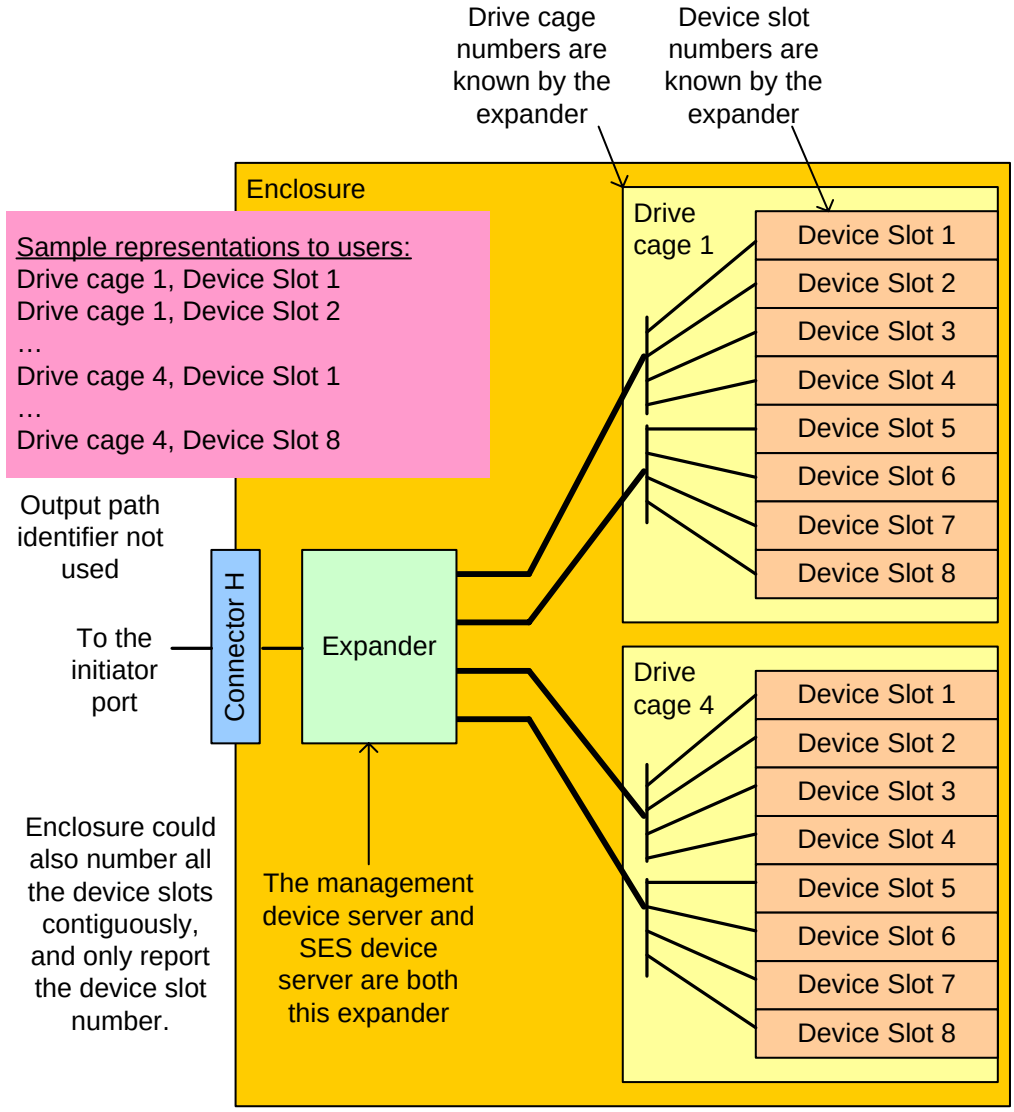

Figure 2 — Device slot number example not using SGPIO

Suggested changes

3.1.xx left-aligned: A type of field containing ASCII data in which unused bytes are placed at the end of the field (highest offset) and are filled with ASCII space (20h) characters. See SPC-4.

10.4.3.9 DISCOVER function

The DISCOVER function returns information about the specified phy. This SMP function provides information from the IDENTIFY address frame received by the phy and additional phy-specific information. This SMP function shall be implemented by all management device servers.

NOTE 102 - The DISCOVER LIST function (see 10.4.3.14) returns information about one or more phys.

Table 388 defines the request format.

Table 388 — DISCOVER request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (10h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
7								
8	Reserved							IGNORE ZONE GROUP
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)	CRC						
15								(LSB)

The SMP FRAME TYPE field shall be set to 40h.

The FUNCTION field shall be set to 10h.

The REQUEST LENGTH field shall be set to 02h. For compatibility with previous versions of this standard, a REQUEST LENGTH field set to 00h specifies that there are 2 dwords before the CRC field.

An IGNORE ZONE GROUP bit set to one specifies that the management device server shall return information about the specified phy (i.e., the phy specified by the PHY IDENTIFIER field) regardless of the zone permission table.

An IGNORE ZONE GROUP bit set to zero specifies that the management device server shall:

a) if the SMP initiator port has access to the specified phy based on the zone permission table, return the requested information; and
b) if the SMP initiator port does not have access to the specified phy, return a function result of PHY VACANT in the response frame (see table 315 in 10.4.3.2).

If the management device server is not in a zoning expander device with zoning enabled, it shall ignore the IGNORE ZONE GROUP bit.

The PHY IDENTIFIER field specifies the phy (see 4.2.8) for which the information is being requested.

The CRC field is defined in 10.4.3.1.

Table 389 defines the response format.

Table 389 — DISCOVER response (part 1 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
Header								
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (1Ah)							
4	(MSB) EXPANDER CHANGE COUNT							
5								(LSB)
6	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
Received IDENTIFY address frame fields								
12	Reserved	ATTACHED DEVICE TYPE			ATTACHED REASON			
13	Reserved				NEGOTIATED LOGICAL LINK RATE			
14	Reserved				ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15	ATTACHED SATA PORT SELECTOR	Reserved			ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA DEVICE
16	SAS ADDRESS							
23								
24	ATTACHED SAS ADDRESS							
31								
32	ATTACHED PHY IDENTIFIER							
33	Reserved					ATTACHED INSIDE ZPSDS PERSISTENT	ATTACHED REQUESTED INSIDE ZPSDS	ATTACHED BREAK_REPLY CAPABLE
34	Reserved for IDENTIFY address frame-related fields							
39								
Other SAS-1.1 fields (and SAS-2 fields implemented by SAS-1.1 expanders)								
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							

Table 389 — DISCOVER response (part 2 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	Reserved				ROUTING ATTRIBUTE			
45	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved							
49								
50	Vendor specific							
51								
SAS-2 miscellaneous fields								
52	ATTACHED DEVICE NAME							
59								
60	Reserved	REQUESTED INSIDE ZPSDS CHANGED BY EXPANDER	INSIDE ZPSDS PERSISTENT	REQUESTED INSIDE ZPSDS	Reserved	ZONE GROUP PERSISTENT	INSIDE ZPSDS	ZONING ENABLED
61	Reserved for zoning-related fields							
62								
63	ZONE GROUP							
64	SELF-CONFIGURATION STATUS							
65	SELF-CONFIGURATION LEVELS COMPLETED							
66	Reserved for self-configuration related fields							
67								
68	SELF-CONFIGURATION SAS ADDRESS							
75								
SAS-2 link reset sequence related fields								
76	PROGRAMMED PHY CAPABILITIES							
79								
80	CURRENT PHY CAPABILITIES							
83								
84	ATTACHED PHY CAPABILITIES							
87								

Table 389 — DISCOVER response (part 3 of 3)

Byte\Bit	7	6	5	4	3	2	1	0
88	Reserved							
93								
94	REASON				NEGOTIATED PHYSICAL LINK RATE			
95	Reserved						NEGOTIATED SSC	HARDWARE MUXING SUPPORTED
Default, saved, and shadow zone phy information								
96	Reserved		DEFAULT INSIDE ZPSDS PERSISTENT	DEFAULT REQUESTED INSIDE ZPSDS	Reserved	DEFAULT ZONE GROUP PERSISTENT	Reserved	DEFAULT ZONING ENABLED
97	Reserved							
98	Reserved							
99	DEFAULT ZONE GROUP							
100	Reserved		SAVED INSIDE ZPSDS PERSISTENT	SAVED REQUESTED INSIDE ZPSDS	Reserved	SAVED ZONE GROUP PERSISTENT	Reserved	SAVED ZONING ENABLED
101	Reserved							
102	Reserved							
103	SAVED ZONE GROUP							
104	Reserved		SHADOW INSIDE ZPSDS PERSISTENT	SHADOW REQUESTED INSIDE ZPSDS	Reserved	SHADOW ZONE GROUP PERSISTENT	Reserved	
105	Reserved							
106	Reserved							
107	SHADOW ZONE GROUP							
Other SAS-2 fields								
108	DEVICE SLOT NUMBER							
109	DEVICE SLOT GROUP NUMBER							
110	DEVICE SLOT GROUP OUTPUT CONNECTOR							
115								
Footer								
~~108~~116	(MSB)			CRC				
~~111~~119								(LSB)

The SMP FRAME TYPE field shall be set to 41h.

The FUNCTION field shall be set to 10h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 1Ah. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h indicates that there are 12 dwords before the CRC field.

The EXPANDER CHANGE COUNT field is defined in the SMP REPORT GENERAL response (see 10.4.3.3).

The PHY IDENTIFIER field indicates the phy for which information is being returned.

...

The CONNECTOR TYPE field indicates the type of connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2). A CONNECTOR TYPE field set to 00h indicates no connector information is available and that the CONNECTOR ELEMENT INDEX field and the CONNECTOR PHYSICAL LINK fields are invalid and shall be ignored.

The CONNECTOR ELEMENT INDEX indicates the element index of the SAS Connector element representing the connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2).

The CONNECTOR PHYSICAL LINK field indicates the physical link in the connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2).

...

The SHADOW ZONE GROUP field contains the default value of the ZONE GROUP field in the zone phy information (see 4.9.3.1).

The DEVICE SLOT NUMBER field indicates the number of the enclosure device slot to which the phy provides access, as reported by the enclosure services process for the enclosure (see the Additional Element Status descriptor for Device Slot and Array Device Slot elements in SES-2). A DEVICE SLOT NUMBER field set to FFh indicates that no device slot number is available.

The DEVICE SLOT GROUP NUMBER field indicates the number of the group of device slots containing the device slot indicated by the DEVICE SLOT NUMBER field. A DEVICE SLOT GROUP NUMBER field set to FFh indicates that no device slot group number is available.

> NOTE 103 - This may be the same as the Group ID reported via the SGPIO input stream from the enclosure (see SFF-8485).

The DEVICE SLOT GROUP OUTPUT CONNECTOR field contains a left-aligned ASCII string describing the connector of the enclosure containing the management device server attached to the device slot group indicated by the DEVICE SLOT GROUP NUMBER field. A DEVICE SLOT GROUP OUTPUT CONNECTOR field set to 202020202020h (i.e., six space characters) indicates that no device slot group output connector information is available.

The CRC field is defined in 10.4.3.2.