

To: T10 Technical Committee
From: Tim Symons, PMC-Sierra (Tim_Symons@pmc-sierra.com)
Date: 20 June 2006
Subject: 06-286r0 SAS-2 Zone Management lock.

Revision Information

- Revision 0: First Revision

Referenced Documents

sas2r04 Serial Attached SCSI – 2 (SAS-2) revision 4
06-201r1 SAS-2 SMP Configure phy zone (Tim Symons, PMC-Sierra)
06-202r2 SAS-2 SMP Configure zone permission (Tim Symons, PMC-Sierra)
06-213r2 SAS-2 REPORT GENERAL additions for zoning and self configuration (Steve Johnson, LSI Logic)
06-285r0 SAS-2 Zone Configuration Management models (Tim Symons, PMC-Sierra)
06-288r1 SAS-2 SMP ACTIVATE ZONE UPDATE function (Tim Symons, PMC-Sierra)
06-289r1 SAS-2 SMP ZONE RELEASE LOCK (Tim Symons, PMC-Sierra)

Overview

For a ZPSDS to function correctly all zoning expander devices must have identical zone permission tables. A Change to the zone permission tables shall be originated by a zone management client and applied to each zoning expander device individually. If there is any interruption to the process then there is a potential for the zone route tables in zoning expander devices to become inconsistent with each other and for the ZPSDS to be corrupted. The zone management lock, activate and unlock function minimizes the risk of corruption, and provides a positive mechanism for confirmation of completion.

Examples of causes of inconsistent zone permission tables are:

- a) Two or more zone management clients attempt to update the ZPSDS at the same time with different information
- b) A device failure causes the process to be aborted part way through an update

This proposal for a lock, load, activate and release procedure minimizes the risk of a ZPSDS becoming inconsistent and ensures identification of completion and error conditions:

1. LOCK: writes the SAS address of the zone management client to all zoning expanders to ensure that only one zone management client updates the zone configuration at any time.
2. LOAD: Only the device that locked the ZPSDS may issue configure requests. Changes are applied to shadow registers and shall not be activated until requested.
3. ACTIVATE: When all zoning configuration changes have been successfully loaded to the shadow registers, then the zone management client shall issue an activate request to all zoning expander devices and the shadow data then becomes operational.
4. RELEASE: When the shadow registers of all zoning expander devices have been activated then the zone expanders are unlocked and the configuration process is complete.

A zoning expander device shall support a shadow zone permission table. When a management application client locks the zone, all changes shall be written to shadow registers until an SMP ACTIVATE ZONE UPDATE request is received.

[Suggested addition to SAS-2 existing text (included in black), new additional text (included in blue) and changes between revisions shown in red]

3.1 Definitions

3.1.x active zone management client: The zone management client that locked the zoning expander devices in a ZPSDS.

3.1.x SMP zone locked function: a function that may only be originated by the active zone management client when the zoning expander devices are locked, indicated by the ZONE LOCKED bit set to one in the REPORT GENERAL response.

4.8.2 Zoning expander device requirements

In addition to the requirements for expander devices described in 4.6, a zoning expander device shall:

- ...
- h) contain a shadow zone permission table that supports 128 zone groups; and
- i) contain a shadow phy zone configuration descriptor for each phy.
- ...

[Editors note these relate to proposals 06-201 and 06-202, zoning related SMP functions, still to be defined in SAS-2]

10.4.3.1 SMP function request frame format

Table 196 – SMP functions (FUNCTION field)

Code	SMP function	Description	Reference
84h	ZONE LOCK	A zone management client sends the ZONE LOCK request to all zoning expanders in the ZPSDS. The zoning expander device shall register the ZONE LOCK and set the zone lock bit to one.	10.4.3.xx
85h	ACTIVATE ZONE LOCKED UPDATE	The zoning expander device shall make the shadow register data operational.	10.4.3.xx
86h	ZONE LOCK RELEASE	The zoning expander device shall set the zone lock bit to zero.	10.4.3.xx
...			

10.4.3.2 SMP function response frame format

Table 198 – FUNCTION RESULT field

Code	Name	SMP function(s)	Description
21h	ZONE LOCK VIOLATION	ZONE LOCK	A zoning expander device that is zone locked, received an SMP request function from a source that is not the ZONE LOCK device.
...			

[Editors note: ZONE LOCK VIOLATION shall also be used by SMP CONFIGURE ZONE PERMISSION (06-202) and SMP CONFIGURE PHY ZONE (06-201). These are zoning related SMP functions, still to be defined in SAS-2]

10.4.3.3 REPORT GENERAL function

Table 199 defines the response format.

Table 199 - REPORT GENERAL response

Byte\Bit	7	6	5	4	3	2	1	0	
...	...								
3	RESPONSE LENGTH (0bh)								
...	...								
36	OPTIMIZED ZONE BROADCAST	Reserved					ZONE LOCKED	ZONING ENABLED	
...	...								
40	ZONE LOCK SAS ADDRESS								
47									
48	(MSB)	CRC							
51							(LSB)		

...

The REQUEST LENGTH field shall be set to (0bh).

...

An OPTIMIZED ZONE BROADCAST bit set to one indicates that the zoning expander device issue a BROADCAST (Change) only to zone groups affected by the change (i.e. zone group members of a group that experiences an addition and removal due to any event). An OPTIMIZED BROADCAST bit set to zero indicates that the zoning expander device shall issue a BROACAST (Change) to all zone groups on any change event.

[Editors note: The OPTIMIZED ZONE BROADCAST operation affects the following functions SMP ACTIVATE ZONE LOCKED function (06-288), SMP CONFIGURE PHY ZONE function (06-201) and SMP CONFIGURE ZONE PERMISSION function (06-202)]

A ZONE LOCKED bit set to one indicates that the ZPSDS is locked by the zone management client device indicated by the SAS address in the ZONE LOCK SAS ADDRESS field. A ZONE LOCKED bit set to zero indicates that the ZPSDS is not locked and the SAS address in the ZONE LOCK SAS ADDRESS field indicates the zone management client that last updated the zoning expander device.

The ZONING ENABLED bit is defined in proposal 06-213r2 SAS-2 REPORT GENERAL additions for zoning and self configuration (Steve Johnson, LSI Logic)

...

The ZONE LOCK SAS ADDRESS field indicates the SAS address of the zone management client device. This field should be persistent though power cycle events or may be loaded to a user defined default value at power on.

...

 Start of new definitions. Changes between revisions are shown in red

10.4.3.x ZONE LOCK function

This function shall be supported by all zone management client devices. The ZONE LOCK function is used by a zone management client to register itself as the active zone management device. Each zoning expander device registers the zone management client address, and is locked from accepting any SMP zone locked functions from other zone management clients. The zoning expander device shall be locked during zone configuration changes. Zone configuration events shall be constrained by timers, and the lock shall be released as soon as the changes are complete or a time-out occurs.

[Editors Note: The timers are not defined in this document. Suggest a separate proposal to define the timers, and use this document to define the functions.]

To resolve any instance where more than one zone management client attempts to lock the zoning expanders in the ZPSDS, the following rules ensure that any concurrent requests are always resolved:

- a) If all zoning expander devices respond with SMP FUNCTION ACCEPTED then the ZPSDS is locked by the zone management client;
- b) If at least one zoning expander device responds with SMP FUNCTION ACCEPTED and any other response has a function result of ZONE LOCK VIOLATION, then this indicates that another zone management client device has locked the zoning expander device and the following actions shall proceed;
 - a. if the returned SAS address has a higher numeric value than the SAS address of the zone management client, then the zone management client shall send ZONE LOCK RELEASE requests to all zoning expander devices that returned SMP FUNCTION ACCEPTED; or
 - b. if the returned SAS address has a lower numeric value than the SAS address of the zone management client, then it shall send ZONE LOCK requests to all zoning expander devices that returned a function result of ZONE LOCK VIOLATION.
- c) If all zoning expander devices respond with a function result of ZONE LOCK VIOLATION then the ZPSDS is locked by another zone management client.

A ZONE LOCK request shall only be originated by a zone management client that has access permission for zone group 2.

Table 1 defines the ZONE LOCK request format.

Table 1 – ZONE LOCK request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (84h)								
2	Reserved								
3	REQUEST LENGTH (00h)								
4	(MSB)	CRC							
7								(LSB)	

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

The FUNCTION field shall be set to 84h.

The REQUEST LENGTH field shall be set to (00h).

The CRC field is defined in 10.4.3.2.

The ZONE LOCK response format is defined in Table 2.

Table 2 – ZONE LOCK response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (84h)								
2	FUNCTION RESULT								
3	RESPONSE LENGTH (02h)								
4	ZONE LOCK SAS ADDRESS								
11									
12	(MSB)	CRC							
15								(LSB)	

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

The FUNCTION field shall be set to 84h.

The RESPONSE LENGTH field shall be set to 02h.

The ZONE LOCK SAS ADDRESS field indicates the SAS address of the zone management client that has locked the zoning expander device.

The CRC field is defined in 10.4.3.2.