

## T10/08-267 Revision 0

Date: July 14, 2008

To: T10 Committee

From Brad Besmer, LSI

Subject: SAS-2 Domain Management Zoning Interaction

### Overview

A specific SAS Initiator (SSP, STP & SMP) may have several roles:

- a) Base software (normal SCSI I/O, Task management, etc) -> should only be exposed to end devices that zoning allows access to.
- b) Zone Management - Configures expander zoning tables -> Needs full access to the zoning SMP requests.
- c) Domain management - Monitors/controls various aspects of the entire SAS domain -> Does not need access to the end device itself, but does need access to all REPORT SMP functions, PHY CONTROL, PHY TEST, etc.

These roles may be independent (ie. only Base software usage), or combination (ie. Base Software and Domain Management, but NOT Zone Manager)

SAS 2.0 (r14c) use of IGNORE ZONE GROUP in SMP DISCOVER and SMP DISCOVER LIST allows a domain management application accessing the domain via an end device, to discover the full topology regardless of zoning constraints. Additionally, zone permission access to Zone Group 2 provides for additional aspects of a domain manager to send a few limited domain management functions to zoning expanders.

However allowances have not been provided which allow a domain management application (accessing via end device even with access to Zone Group 2) to query information pertinent to the COMPLETE status of the domain independent of zoning (Base Software + Domain Management).

SAS 2.0 currently has overloaded Zone Group 2 access for 2 of these roles:

- 1) Access to Zone Management SMP functions (ie. ZONE LOCK, CONFIGURE ZONE PHY INFORMATION)
- 2) Access to Domain Management functions (ie. PHY CONTROL, PHY TEST, CONFIGURE PHY EVENT), however the REPORT SMP functions are missing.

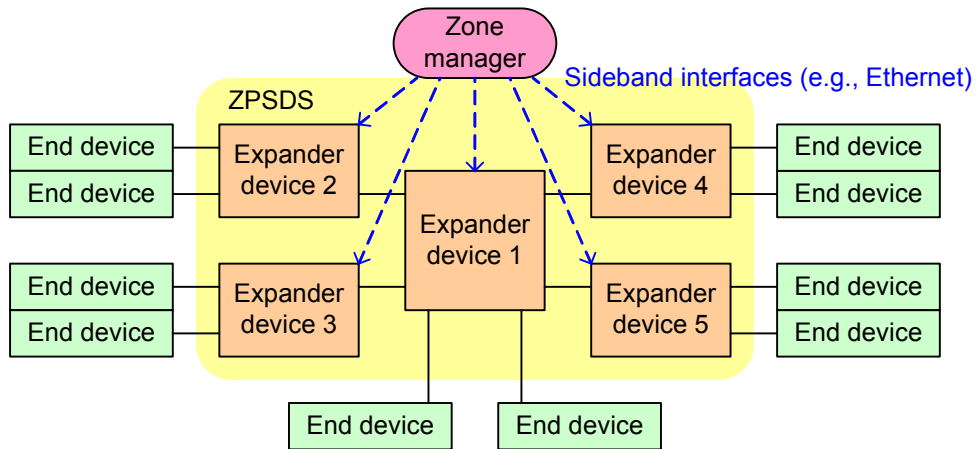
One method would be to add the missing SMP requests to the Domain Management Functions:

REPORT PHY ERROR LOG  
REPORT PHY SATA  
REPORT ROUTE INFORMATION  
REPORT PHY EVENT

This proposal however seeks to split this functionality into these 2 roles, thereby allowing one management server to be the ZONE MANAGER, and another management server be the DOMAIN MANAGER. An example usage of this would be one of the following example topologies (Figure 50 from SAS2r14c), where the Zone Manager is a side-band interface to the expanders, and the Domain Manager is running on one of the Initiator End Devices.

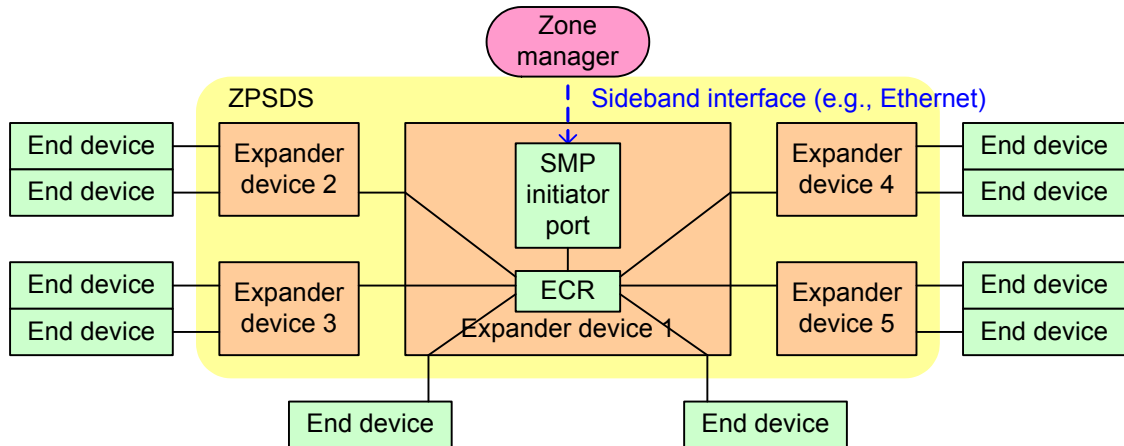
---

Zone manager attached directly to the expander devices in the ZPSDS



---

Zone manager attached directly to one expander device in the ZPSDS



Proposed Changes:

4.9.3.2 Zone groups

The zone groups are defined in table 30.

Table 30 — Zone groups

Zone group	Configurable in the zone permission table <sup>a</sup>	Description
0	No	Phys in zone group 0 have access to phys in zone group 1 and do not have access to phys in other zone groups.
1	No	Phys in zone group 1 have access to phys in all zone groups.
2	Yes	<p>Phys in zone group 2 have access to phys in the zone groups indicated by the zone permission table.</p> <p>A management device server in a zoning expander device with zoning enabled only allows management application clients using phys in zone groups with access to zone group 2 to perform the following SMP functions:</p> <ul style="list-style-type: none"> <li>a) <del>CONFIGURE GENERAL (see 10.4.3.18);</del></li> <li>b) ZONE LOCK (see 10.4.3.21); and</li> <li>c) SMP zone configuration functions (see 4.9.6.1) performed while the zoning expander device is locked.</li> </ul> <p>A management device server in a zoning expander device with zoning enabled only allows management application clients to perform certain SMP phy-based control and configuration functions (e.g., PHY CONTROL, <del>PHY TEST FUNCTION, and CONFIGURE PHY EVENT</del>) if the zone group of the management application client's phy has access to zone group 2 or the zone group of the specified phy.</p>
3	Yes	<p>Phys in zone group 3 have access to phys in the zone groups indicated by the zone permission table.</p> <p>A management device server in a zoning expander device with zoning enabled only allows management application clients using a phy in a zone group with access to zone group 3 to perform certain SMP zoning-related functions (i.e., ZONED BROADCAST (see 10.4.3.20)).</p>
4	<u>Yes</u>	<p><u>Phys in zone group 4 have access to phys in the zone groups indicated by the zone permission table</u></p> <p><u>A management device server in a zoning expander device with zoning enabled only allows management application clients to perform certain SMP phy-based control and configuration functions (e.g., PHY CONTROL, PHY TEST FUNCTION, and CONFIGURE PHY EVENT) if the zone group of the management application client's phy has access to zone group 4 or the zone group of the specified phy.</u></p>
5 to 7	Reserved	
8 to 255	Yes	Phys in zone groups 8 through 255 have access to phys in the zone groups indicated by the zone permission table.
<p><sup>a</sup> A zone group defined as configurable is able to be changed with the SMP CONFIGURE ZONE PERMISSION TABLE function (see 10.4.3.26).</p>		

#### 10.4.3.11 REPORT PHY ERROR LOG function

The REPORT PHY ERROR LOG function returns error logging information about the specified phy. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

...

#### 10.4.3.12 REPORT PHY SATA function

The REPORT PHY SATA function returns information about the SATA state for a specified phy. This SMP function shall be implemented by management device servers behind SMP target ports that share SAS addresses with STP target ports and by management device servers in expander devices with STP/SATA bridges. This SMP function shall not be implemented by any other type of management device server. In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

...

#### 10.4.3.13 REPORT ROUTE INFORMATION function

The REPORT ROUTE INFORMATION function returns an expander route entry from a phy-based expander route table within an expander device (see 4.6.7.4). This SMP function shall be supported by management device servers in expander devices if the EXPANDER ROUTE INDEXES field is set to a non-zero value in the SMP REPORT GENERAL response (see 10.4.3.4). This SMP function may be used as a diagnostic tool to resolve topology issues. In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

...

#### 10.4.3.14.1 REPORT PHY EVENT function overview

The REPORT PHY EVENT function returns phy events (see 4.11) concerning the specified phy. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

...

#### 10.4.3.18 CONFIGURE GENERAL function

The CONFIGURE GENERAL function requests actions by the device containing the management device server. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled, then this function ~~shall only~~ should be processed from SMP initiator ports that have access to zone group ~~24~~ (see 4.9.3.2).

...

#### 10.4.3.28 PHY CONTROL function

The PHY CONTROL function requests actions by the specified phy. This SMP function may be implemented by any management device server. ~~In zoning expander devices, if zoning is enabled, then this function shall only be processed from SMP initiator ports that have access to zone group 24 or the zone group of the specified phy (see 4.9.3.2).~~ In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

...

10.4.3.29 PHY TEST FUNCTION function

The PHY TEST FUNCTION function requests actions by the specified phy. This SMP function may be implemented by any management device server. ~~In zoning expander devices, if zoning is enabled, then this function shall only be processed from SMP initiator ports that have access to zone group 24 or the zone group of the specified phy (see 4.9.3.2).~~

In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4...

10.4.3.30.1 CONFIGURE PHY EVENT function overview

The CONFIGURE PHY EVENT function configures phy events (see 4.11) for the specified phy. This SMP function may be implemented by any management device server. ~~In zoning expander devices, if zoning is enabled, then this function shall only be processed from SMP initiator ports that have access to zone group 24 or the zone group of the specified phy (see 4.9.3.2).~~In zoning expander devices, if zoning is enabled, then this function shall be processed from SMP initiator ports that have access to the zone group of the specified phy (see 4.9.3.2) and should be processed from SMP initiator ports that have access to zone group 4.

10.4.3.26.3 Zone permission configuration descriptor

....

Table 332 defines how the zone permission descriptor bits shall be set by the management application client and processed by the management device server.

Table 332 — Zone permission configuration descriptor bit requirements

Source zone group (i.e., s)	Management application client requirement(s) <sup>a</sup>	Management device server requirement(s) <sup>a</sup>
0	ZP[s, 0] shall be set to zero. ZP[s, 1] shall be set to one. ZP[s, 2 through (z-1)] shall be set to zero.	ZP[s, 0 through (z-1)] shall be ignored.
1	ZP[s, 0 through (z-1)] shall be set to one.	ZP[s, 0 through (z-1)] shall be ignored.
<del>4</del> , 5, 6 or 7	ZP[s, 0] shall be set to zero. ZP[s, 1] shall be set to one. ZP[s, <del>4</del> 5 through (z-1)] shall be set to zero.	ZP[s, 0 through (z-1)] shall be ignored.
2, 3, <del>4</del> , or 8 <sup>a</sup>	ZP[s, 0] shall be set to zero. ZP[s, 1] shall be set to one. ZP[s, 2 through <del>4</del> 3] may be set to zero or one. ZP[s, <del>4</del> 5 through 7] shall be set to zero. ZP[s, 8 through (z-1)] may be set to zero or one.	ZP[s, 0 through 1] shall be ignored. ZP[s, 2 through 3] shall <del>be</del> processed. <u>ZP[s, 4] should be processed.</u> ZP[s, <del>5</del> 4 through 7] shall be ignored. ZP[s, 8 through (z-1)] shall processed. For each source zone group t other than s, ZP[t, s] shall be set to ZP[s, t].
<sup>a</sup> The number of zone groups (i.e., z) is specified in NUMBER OF ZONE GROUPS field.		

4.9.3.3 Zone permission table

The zone permission table structure is shown in table 31.

**Table 31 — Zone permission table**

Destination zone group (i.e., d)	Source zone group (i.e., s) <sup>a b</sup>				
	0	1	2 to <del>4</del> 3	<del>5</del> 4 to 7	8 to (z-1) <sup>c</sup>
0	0	1	0	0	0
1	1	1	1	1	1
2 to <del>4</del> 3	0	1	ZP[s = 2 to <del>4</del> 3, d = 2 to <del>4</del> 3]	Reserved	ZP[s = 8 to (z-1), d = 2 to <del>4</del> 3]
<del>5</del> 4 to 7	0	1	Reserved	Reserved	Reserved
8 to (z-1) <sup>c</sup>	0	1	ZP[s = 2 to <del>3</del> 4, d = 8 to (z-1)]	Reserved	ZP[s = 8 to (z-1), d = 8 to (z-1)]

<sup>a</sup> Shading identifies configurable zone groups.  
<sup>b</sup> All reserved ZP bits shall be set to zero (e.g., bits ZP[~~5~~4 to 7, ~~5~~4 to (z-1)] are set to zero).  
<sup>c</sup> The number of zone groups (i.e., z) is reported in NUMBER OF ZONE GROUPS field in the REPORT GENERAL response (see 10.4.3.4).

...

If the zoning expander device preserves whether or not zoning is enabled and does not preserve the zone permission table, then it shall set the zone permission table to grant minimal permissions after power on or expander device reduced functionality as specified in table 32.

**Table 32 — Zone permission table granting minimal permissions**

Destination zone group (i.e., d)	Source zone group (i.e., s) <sup>a b</sup>				
	0	1	2 to <del>4</del> 3	<del>5</del> 4 to 7	8 to (z-1) <sup>c</sup>
0	0	1	0	0	0
1	1	1	1	1	1
2 to <del>4</del> 3	0	1	0	Reserved	0
<del>5</del> 4 to 7	0	1	Reserved	Reserved	Reserved
8 to (z-1) <sup>c</sup>	0	1	0	Reserved	0

<sup>a</sup> Shading identifies configurable zone groups.  
<sup>b</sup> All reserved ZP bits shall be set to zero (e.g., bits ZP[~~5~~4 to 7, ~~5~~4 to (z-1)] are set to zero).  
<sup>c</sup> The number of zone groups (i.e., z) is reported in NUMBER OF ZONE GROUPS field in the REPORT GENERAL response (see 10.4.3.4).