

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 21 April 2006
 Subject: 06-166r4 SAS-2 Restrict access to SMP ZONED BROADCAST function

Revision history

Revision 0 (22 March 2006) First revision. Proposed allowing access to SMP initiator ports with access to zone group 2.
 Revision 1 (29 March 2006) Incorporated comments from 23 March 2006 SAS zoning teleconference - change to allow access only to SMP initiator ports with access to a new special zone group 3.
 Revision 2 (31 March 2006) Incorporated comments from 30 March 2006 SAS zoning teleconference.
 Revision 3 (19 April 2006) Incorporated comments from 13 April 2006 SAS zoning teleconference - remove the Link layer tested column from the zone group table (as mentioned by an editor’s note in sas2r03), since all zone groups are link layer tested.
 Revision 3 (21 April 2006) Incorporated comments from 20 April 2006 SAS zoning WG.

Related documents

sas2r03 - Serial Attached SCSI - 2 (SAS-2) revision 3
 06-082r2 SAS-2 zoning-related SMP functions (Ralph Weber, ENDL and Steve Johnson, LSI Logic)
 06-165 SAS-2 Change SMP ZONE VIOLATION code value (Rob Elliott, HP)
 06-168 SAS-2 OPEN address frame SOURCE ZONE GROUP field definition (Rob Elliott, HP)

Overview

The SMP ZONED BROADCAST function causes the recipient expander device to transmit:

- a) BROADCAST primitives to attached to end devices belonging to certain zones
- b) SMP ZONED BROADCAST functions to neighboring expander devices.

Issue #1

In sas2r03, the function can be sent by any SMP initiator port that has access to the expander device. This means an SMP initiator port in one zone group can introduce BROADCAST traffic into other zone groups (see figure 1), which somewhat defeats the purpose of zoning. It cannot introduce an OPEN to another zone group, because the expander to return OPEN_REJECT (ZONE VIOLATION) using the address resolved method.

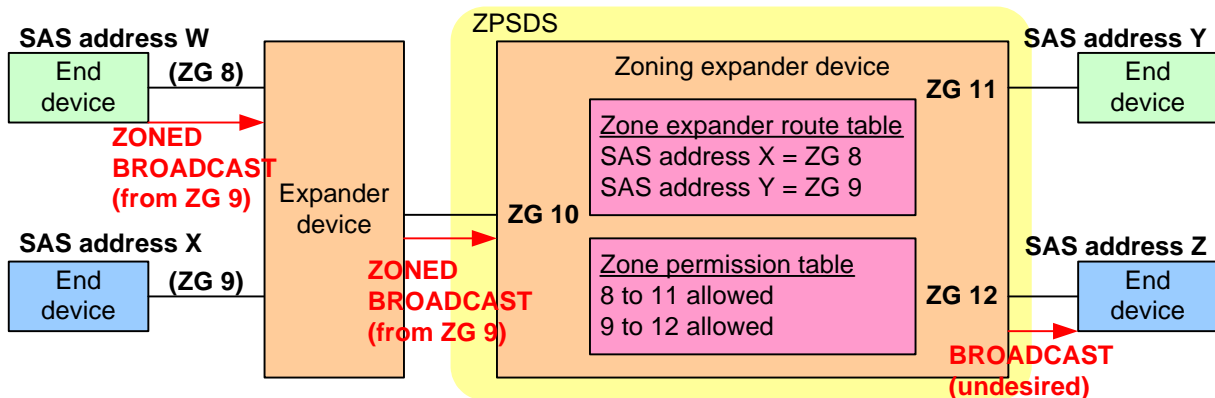


Figure 1 — ZONED BROADCAST from a rogue initiator port

Access needs to be restricted so SMP initiator ports cannot sneak BROADCAST traffic into other zone groups unless they have permission, making it the same as OPEN traffic.

Issue #2

With the “zone address resolved” mode, an expander phy on the boundary of the ZPSDS that is attached to another expander device may be the entry point for traffic from devices in more than one zone group. However, since BROADCAST primitives do not carry source zone group information, an incoming BROADCAST primitive must be considered to have come from any of those zone groups. The expander device must treat it as coming from all the possible source zone groups and forward it to all destination zone

groups to which all the source zone groups on that expander phy have permission (see figure 2). This is undesirable.

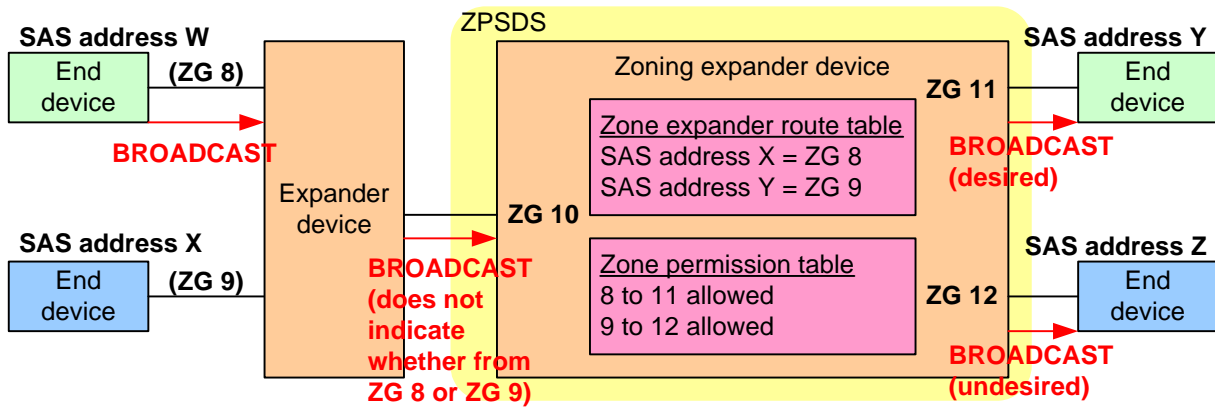


Figure 2 — BROADCAST escaping zone

A mechanism should be defined to allow the expander device outside the ZPSDS to identify the source zone group of the BROADCAST so it can be forwarded only to the proper destination zone groups. This cannot be perfect (since the expander device is not trusted), but it can limit the extent of the BROADCASTs if it is behaving (and anyway is restricted only to the zone groups that are sharing the door into the ZPSDS).

Proposal

In sas2r03, the zone permission table is used to restrict access to (yet to be defined) SMP functions performing zoning configuration was defined. SMP initiator ports that are in zone groups that have access to zone group 2 are allowed to perform SMP zoning configuration functions; SMP initiator ports without that access are refused with an SMP response of ZONE VIOLATION.

The ZPSDS consists of expander devices that the system vendor has determined (somehow) comprise a trusted environment (they trust each other to set the SOURCE ZONE GROUP field in the OPEN address frames correctly). Everything coming from outside the ZPSDS is trusted only to the level dictated by the zone permission table.

Expander devices outside the ZPSDS could have knowledge of zoning, even though they are not totally trusted. To reduce the excess BROADCAST traffic, these expander devices should be allowed to send SMP ZONED BROADCAST functions but not perform zoning configuration. The expander device inside the ZPSDS will only honor requests for source zone groups supported by the expander phy.

A new zone group 3 is proposed for this purpose. The zone permission table must indicate the SMP initiator port has access to zone group 3 for the SMP ZONED BROADCAST function to be allowed. This allows an external expander device, if its own source zone group has access, to use SMP ZONED BROADCAST (see figure 3).

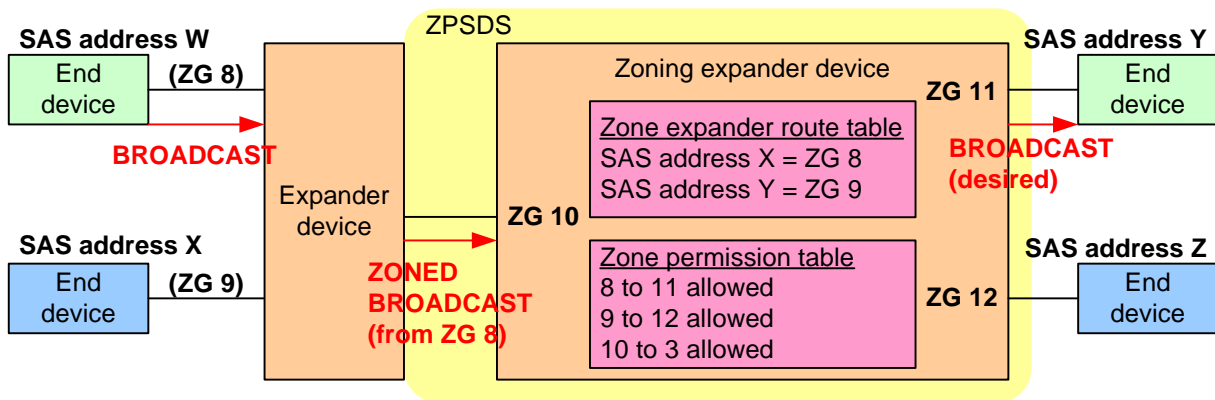


Figure 3 — BROADCAST kept in zone

Other options considered but discarded:

- a) Restrict access to SMP initiator ports with access to zone group 2 (used to prohibit access to zoning configuration related SMP functions). This is overly restrictive.
- b) Only allow ZONED BROADCAST if it arrives on an expander port with the ZONE PARTICIPATING bit set to one. This would prevent all SMP initiator ports outside the ZPSDS from sending the function, but allow all SMP initiator ports inside the ZPSDS to send the function. This is overly restrictive, and would require expanders at the boundaries to filter incoming SMP function requests (or communicate the incoming expander phy's ZONE PARTICIPATING bit in the OPEN address frame).

Suggested changes

4.8.3.2 Zone group types

4.8.3.2.1 Overview

Zone groups function in different ways based on zone group number as shown in table 22.

Table 22 — Zone group types

Zone group	Configurable <u>zone permission table entries</u>	<u>Link-layer-tested</u>	Description
0	No	Yes	Phys in zone group 0 have access to no other phys.
1	No	Yes	Phys in zone group 1 have access to other phys in all zone groups.
2	Yes	No	Phys in zone group 2 have access to other phys based on the zone permission table. Phys in zone groups with access to zone group 2 are allowed to perform SMP zoning configuration functions. Phys in zone groups without access to zone group 2 are prohibited from performing SMP zoning configuration functions.
<u>3</u>	<u>Yes</u>		<u>Phys in zone group 3 have access to other phys based on the zone permission table.</u> <u>Phys in zone groups with access to zone group 3 are allowed to perform certain SMP zoning related functions (i.e., ZONED BROADCAST).</u> <u>Phys in zone groups without access to zone group 3 are prohibited from performing certain SMP zoning related functions (i.e., ZONED BROADCAST).</u>
3 4 to 7	Reserved		
8 to 127	Yes	Yes	Phys in zone groups 8 through 127 have access to other phys based on the zone permission table.

Editor's Note 1: Wording for ZG 2 will remain vague until a zoning configuration function is defined. Probably a column will be added to the functions table indicating which zone groups are required.

4.8.3.2.2 Link layer tested zone groups

Link layer tested zone groups have their zone permission bits validated as part of processing and forwarding OPEN address frames. Violations of the zone permission are reported using OPEN_REJECT (ZONE VIOLATION).

Editor's Note 2: 4.8.3.3 and 4.8.3.5 define OPEN_REJECT (ZONE VIOLATION) by the link layer; no need to state it three times

4.8.3.2.3 SMP zoning configuration zone group testing

Use of the SMP CONFIGURE PHY ZONE function and SMP CONFIGURE ZONE PERMISSION function shall be restricted based on the zone group 2 permission bits. An SMP target port shall retain the contents of the SOURCE ZONE GROUP field from the OPEN address frame until the SMP request is received. If the SMP request is CONFIGURE PHY ZONE or CONFIGURE ZONE PERMISSION and the ZP[s, 2] bit is set to one, then the SMP target port shall process the function request. If the SMP request is CONFIGURE PHY ZONE or SMP CONFIGURE ZONE PERMISSION and the ZP[source zone group, 2] bit is set to zero, then the SMP target port shall not process the function request ~~shall not be processed~~ and shall return an SMP response frame with the FUNCTION RESULT field set to SMP ZONE VIOLATION ~~shall be returned~~.

Editor's Note 3: 06-168 makes some other changes the above paragraph (not incorporated above). Those changes are incorporated into the new, derivative paragraph that follows.

Use of the SMP ZONED BROADCAST function shall be restricted based on the zone group 3 permission bits and the source zone group of the SMP initiator port requesting the function (see 4.8.3.5). If the SMP request is ZONED BROADCAST and the ZP[s, 3] bit is set one, then the SMP target port shall process the function request. If the SMP request is ZONED BROADCAST and ZP[s, 3] bit is set to zero, then the SMP target port shall not process the function request and shall return an SMP response frame with the FUNCTION RESULT field set to SMP ZONE VIOLATION.

4.8.3.3 Zone permission table

The zone permission table specifies access permission between zone groups. If a bit in the zone permission table is set to one then connection requests shall be permitted between phys in the zone groups. If a bit in the zone permission table is set to zero then connection requests between phys in the zone groups shall be rejected with OPEN_REJECT (ZONE VIOLATION) as described in 4.8.3.5.

The zone permission table structure is shown in table 23.

Table 23 — Zone permission table

Destination zone group d	Source zone group s				
	0	1	<u>2 to 3</u>	3 4 to 7	8 to 127
0	0	1	0	0	0
1	1	1	1	1	1
<u>2 to 3</u>	0	1	ZP[s = <u>2 to 3</u> , d = <u>2 to 3</u>]	Reserved	ZP[s = 8 to 127, d = <u>2 to 3</u>]
3 4 to 7	0	1	Reserved	Reserved	Reserved
8 to 127	0	1	ZP[s = <u>2 to 3</u> , d = 8 to 127]	Reserved	ZP[s = 8 to 127, d = 8 to 127]

Note: Shading identifies configurable zone groups.

A ZP[s, d] bit set to one specifies that zone group s has permission to access zone group d. A ZP[s, d] bit set to zero specifies that zone group s does not have permission to access zone group d.

If ZP[s, d] is set to a value, ZP[d,s] shall be set to the same value.

Zone group 0 shall not have access to any other group except zone group 1 (e.g., bits ZP[0, 0] and bits ZP[2 to 127, 0] are set to zero).

Zone group 1 shall have access to all other zone groups (e.g., bits ZP[0 to 127, 1] are set to one).

Regardless of the zoning expander phy attributes, OPEN address frames addressed to a zoning expander device shall be processed as if the destination is zone group 1.

Zone groups ~~3~~4 through 7 are reserved zone groups. All reserved ZP bits shall be set to zero (e.g., bits ZP[~~3~~4 to 7, ~~3~~4 to 127] are set to zero).

10.4.3.2 SMP function response frame format

...

The FUNCTION RESULT field is defined in table 197.

Table 197 — FUNCTION RESULT field

Code	Name	SMP function(s)	Description
...
04h	SMP ZONE VIOLATION	TBD <u>ZONED BROADCAST</u>	The SMP target port supports the function, but the application zone permission bit is set to zero (e.g., the ZP[s, 2 <u>3</u>] bit is set to zero <u>and the request was ZONED BROADCAST</u>)(see <u>4.8.3.2.3</u>).
...
All others	Reserved		

Editor's Note 4: 06-165 changes code 04h to 20h (not incorporated above)

10.4.3.11 ZONED BROADCAST function

The ZONED BROADCAST function requests that the specified BROADCAST be transmitted on all the ports that are in one or more specified zone groups, with the exception of the port on which the ZONED BROADCAST function was received (see 4.8.5). This SMP function shall be supported by SMP target ports in zoning expander devices (see 4.8). Other SMP target ports shall not support this SMP function. This SMP function shall only be processed from SMP initiator ports that have access to zone group 3 (see 4.8.3.2.3).

...