

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 24 July 2006
 Subject: 06-279r3 SAS-2 Allow more than one ZSDS in a SAS domain

Revision history

Revision 0 (13 June 2006) First revision
 Revision 1 (4 July 2006) Incorporated comments from 20 June 2006 SAS zoning WG. Updated to a base of sas2r04a.
 Revision 2 (9 July 2006) Incorporated comments from 6 July 2006 SAS zoning WG.
 Revision 3 (24 July 2006) Incorporated comments from 13 July 2006 SAS zoning WG.

Related documents

sas2r04a - Serial Attached SCSI - 2 (SAS-2) revision 4
 06-281 SAS-2 Enable and disable zoning (Rob Elliott, HP)

Overview

SAS-2 currently mandates that there be no more than one zoned portion of the service delivery subsystem (ZPSDS) per SAS domain. However, there is no way to enforce this and no reason to do so - a user can easily attach two zoned JBODs together, and they will communicate fine. Each ZPSDS has its own zone manager that continues to manage that ZPSDS. The zone group numbers in one ZPSDS has no relationship to the zone group numbers in the other ZPSDS.

The user is not likely to leave the SAS domain configured in this manner, however, since they will default to considering all the SAS addresses in the other ZPSDS as being in the same zone group, probably ZG 0 with no access to anything. 06-281 describes merging ZPSDSes.

Figure 1 shows two ZPSDSes. As long as the ZONE PARTICIPATING bits are set to zero at the boundaries, they can coexist - devices on the left would not be able to talk to devices on the right (but all the devices can all see all the expanders, since expanders are in ZG 1).

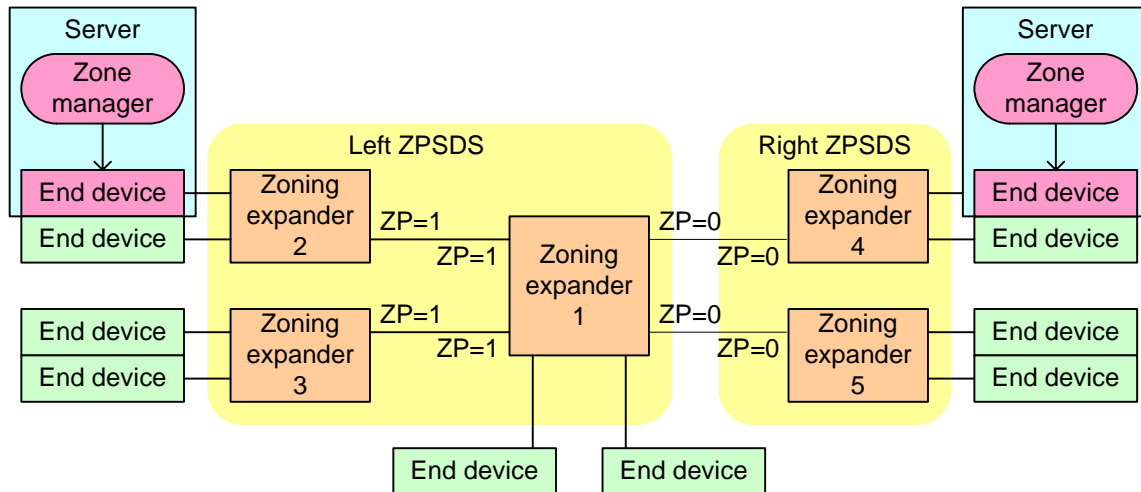


Figure 1 — Two ZPSDSes introduction

This proposal documents this state, and also introduces the concept of a “zone manager” which is the entity that configures a ZPSDS (either in-band using SMP functions or out-of-band using some mechanism outside the scope of this standard).

Suggested changes

3.1.265 zone group: A set of phys in a ZPSDS (see 3.1.267) that all have the same access permission. See 4.8.

3.1.xx zone manager: The entity responsible for configuring a ZPSDS. See 4.8.1.

3.1.266 zone permission table: The table that defines access permission between the zone group of a source phy and the zone group of the destination phy. See 4.8.3.3.

3.1.267 zoned portion of a service delivery subsystem (ZPSDS): A group of [zoning](#) expander devices that cooperate to control access between phys. The ZPSDS may include all or part of a service delivery subsystem. See 4.8.

3.1.268 zoning expander device: An expander device that supports zoning. See 4.8.

4.8 Zoning

4.8.1 Zoning overview

SAS zoning is implemented by a set of zoning expander devices with zoning enabled that define [the](#) zoned portion of a service delivery subsystem (ZPSDS). The zoning expander devices control whether a phy is permitted to participate in a connection to another phy.

[Figure 2 shows an example of one ZPSDS in a SAS domain.](#)

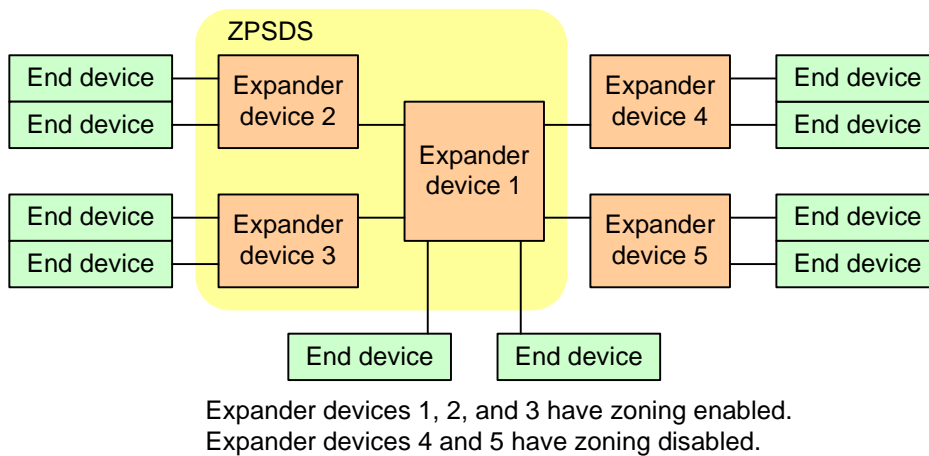



Figure 2 — One ZPSDS example

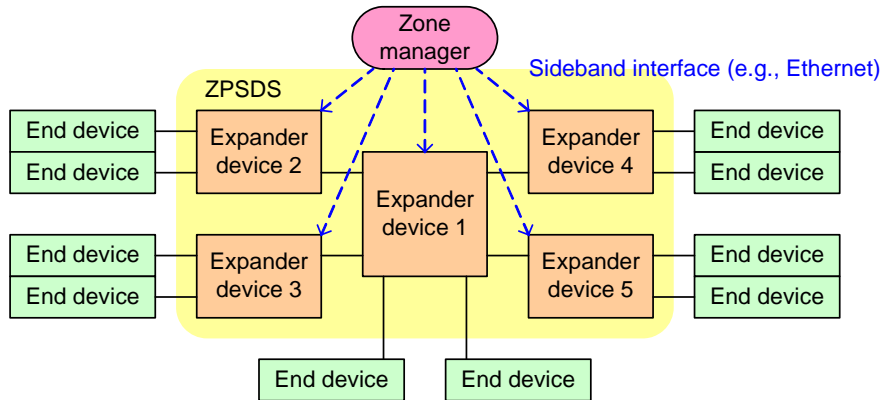
[A ZPSDS has a zone manager responsible for its configuration. The zone manager may have access to an end device with a SAS port whose zone group has access to zone group 2, or may have access to one or more zoning expander devices through a sideband interface \(e.g., Ethernet\) outside the scope of this standard.](#)

Figure 3 shows examples of zone manager locations in a SAS domain.

Zone manager attached to an end device with a SAS port whose zone group has access to zone group 2



Zone manager attached directly to the expander devices in the ZPSDS



Zone manager attached directly to one expander device in the ZPSDS




Figure 3 — Zone manager location examples

There shall be at most one may be any number of non-overlapping ZPSDSes in a service delivery subsystem, particularly as a SAS domain is being reconfigured (e.g., as a user is attaching enclosures together). The A ZPSDS may encompass some or all of the service delivery subsystem.

Figure 4 shows an example of three ZPSDSes in a SAS domain.

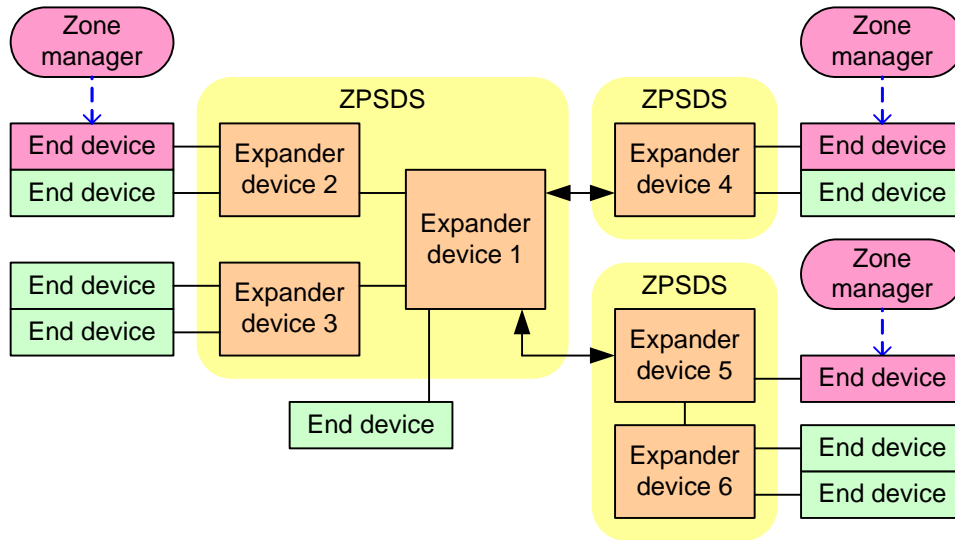


Figure 4 — Three ZPSDSes example

The zone manager assigns zone groups (see 4.8.3.2) to all phys that are capable of participating in connections through the ZPSDS. Every phy in a SAS domain may belong be assigned to a zone group. There are 128 zone groups numbered 0 through 127. All phys in a wide port shall belong be assigned to the same zone group. Zone group numbers are assigned to zoning expander phys as part of the zone phy information (see 4.8.3.1) and assigned to SAS addresses in the zoning expander route table (see 4.8.3.4). An expander phy inside the ZPSDS belongs shall be assigned to zone group 1. A phy attached to the boundary of the ZPSDS (i.e., attached to an expander phy with the ZONE PARTICIPATING bit set to zero) belongs to the zone group indicated by the ZONE GROUP field of the expander phy to which it is attached. A phy attached one or more expander devices beyond the boundary of the ZPSDS belongs to a zone group if the ZONE ADDRESS RESOLVED bit is set to one in the boundary expander phy and the zoning expander route table in that expander device contains a zone group value for the phy's SAS address. The zone groups assigned in one ZPSDS have no relationship to the zone groups assigned in another ZPSDS.

Each zoning expander device contains a zone permission table (see 4.8.3.3) that controls whether a connection is allowed between phys based on their zone groups. the zone group of a source port and the zone group of a target port. As defined in 4.8.3.5, a requested connection shall only be established if the zone permission table indicates that access between the zone group of the source port and the zone group of the target port is allowed.

The zoning expander route table (see 4.8.3.4) is an extended version of the expander route table (see 4.6.7.3) that also includes zone phy information the zone group number of each SAS address.

The treatment of OPEN address frames received by expander phys on the boundary of the ZPSDS depends on:

- a) whether the OPEN address frame enters the ZPSDS via a table routing phy; and
- b) the value of the ZONE ADDRESS RESOLVED bit for the receiving expander phy (see 4.8.3.1).

If an OPEN address frame is received from a non-zoning expander device via a table routing phy for which the ZONE ADDRESS RESOLVED bit is set to one, then the source zone group is determined based on the source SAS address in the OPEN address frame. In this case, proper management of the zoning expander device extends the zoning capability to beyond the ZPSDS.

Otherwise, the source zone group is a fixed value that is associated with the receiving expander phy. In this case, non-zoning expander devices are treated in the same manner as end devices and all phys in the non-zoning expander device and all phys attached behind the non-zoning expander device belong to the zone group specified for the receiving expander phy.

Editor's Note 1: The above three paragraphs just duplicate material covered in the source zone group determination table 23 in 4.8.3.5.

Editor's Note 2: moved the following text/figures about the zone address resolved method to 4.8.3.5:

Figure 55 shows an example of a ZPSDS in which the ZONE ADDRESS RESOLVED bit of the expander port of zoning expander device C that is attached to non-zoning expander device D is set to zero.

Figure 52 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to zero

Figure 56 shows an example of a ZPSDS in which the ZONE ADDRESS RESOLVED bit of the expander port of zoning expander device C that is attached to non-zoning expander device D is set to one.

Figure 53 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to one

4.8.2 Zoning expander device requirements

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

- contain a zoning expander route table (see 4.8.3.4);
- contain a zone permission table that supports 128 zone groups (see 4.8.3.3);
- if zoning is enabled, allow or deny connection requests based on the zone permission table ([see 4.8.3.5](#));
- set the ZONE DEVICE bit to one in its SMP REPORT GENERAL response (see 10.4.3.3);
- be self-configuring;
- contain an SMP initiator port (see 4.6.1); and
- support all zoning-related SMP functions.

4.8.3 Zone operation

4.8.3.1 Zone phy information

Each phy of a zoning expander device shall support the following zone phy information fields:

- ZONE PARTICIPATING bit;
- ZONE ADDRESS RESOLVED bit;
- ZONE GROUP PERSISTENT bit; and
- ZONE GROUP field.

The ZONE PARTICIPATING bit specifies a boundary of the ZPSDS. The ZONE PARTICIPATING bit shall be set to zero when the phy is attached to ~~an end device or~~ an expander device ~~that does not support zoning~~ [outside the ZPSDS or to an end device](#). The ZONE PARTICIPATING bit shall be set to one when the phy is attached to a zoning expander device with zoning enabled [inside the ZPSDS](#).

The ZONE ADDRESS RESOLVED bit specifies the method used to determine the source zone group for a connection request received by a phy at the boundary of the ZPSDS as specified in table 26 (see 4.8.3.5).

The ZONE ADDRESS RESOLVED bit may be set to one when:

- the phy is contained within a zoning expander device; and
- the ZONE PARTICIPATING bit for the phy is set to zero.

The ZONE ADDRESS RESOLVED bit shall be set to zero when:

- the phy is contained within a non-zoning expander device; or
- the phy is contained within a zoning expander device and the ZONE PARTICIPATING bit for the phy is set to one.

The ZONE GROUP field has a value in the range of 0 to 127 that specifies the zone group to which the phy belongs. If the ZONE PARTICIPATING bit of the phy is set to one, the ZONE GROUP field shall be set to 1. The zone group of the SMP initiator port and SMP target port in a zoning expander device shall be 1. [4.8.3.2 defines more about zone groups.](#)

The ZONE GROUP PERSISTENT bit specifies the method of determining the zone group value of the phy after a link reset sequence (see 4.8.4). If the ZONE PARTICIPATING bit is set to one, the ZONE GROUP PERSISTENT bit shall be set to one.

All phys in an expander port shall have the same zone phy information (see 4.8.3.1). The default (e.g., power on) values for the zone phy information fields are vendor-specific.

4.8.3.2 Zone groups

Zone groups function in different ways based on zone group number as shown in table 24. [A zone group with configurable zone permission table entries is able to be changed with the SMP CONFIGURE ZONE PERMISSION function.](#)

Table 24 — Zone groups

Zone group	Configurable zone permission table entries	Description
0	No	Phys in zone group 0 only have access to other phys in zone group 1.
1	No	Phys in zone group 1 have access to other phys in all zone groups.
2	Yes	<p>Phys in zone group 2 have access to other phys based on 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"> a) SMP zoning configuration functions (e.g., CONFIGURE ZONE PERMISSION); and b) SMP expander configuration functions (i.e., CONFIGURE GENERAL). <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 (i.e., PHY CONTROL, PHY TEST FUNCTION, and CONFIGURE PHY EVENT INFORMATION) 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 other phys based on 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).</p>
4 to 7	Reserved	
8 to 127	Yes	Phys in zone groups 8 through 127 have access to other phys based on the zone permission table.

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 25.

Table 25 — Zone permission table

Destination zone group d	Source zone group s				
	0	1	2 to 3	4 to 7	8 to 127
0	0	1	0	0	0
1	1	1	1	1	1
2 to 3	0	1	ZP[s = 2 to 3, d = 2 to 3]	Reserved	ZP[s = 8 to 127, d = 2 to 3]
4 to 7	0	1	Reserved	Reserved	Reserved
8 to 127	0	1	ZP[s = 2 to 3, 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).

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

If an expander phy has the ZONE ADDRESS RESOLVED METHOD bit set to one, then the zone group of that phy shall be configured to have permission to access all the zone groups to which the SAS addresses in the zoning expander route table for that expander phy have access.

4.8.3.4 Zoning expander route table

A zoning expander device contains a zoning expander route table, which is an expander-based expander route table (see 4.6.7.3) that is able to hold the zone group of each routed SAS address.

Figure 54 shows a representation of the zoning expander route table.

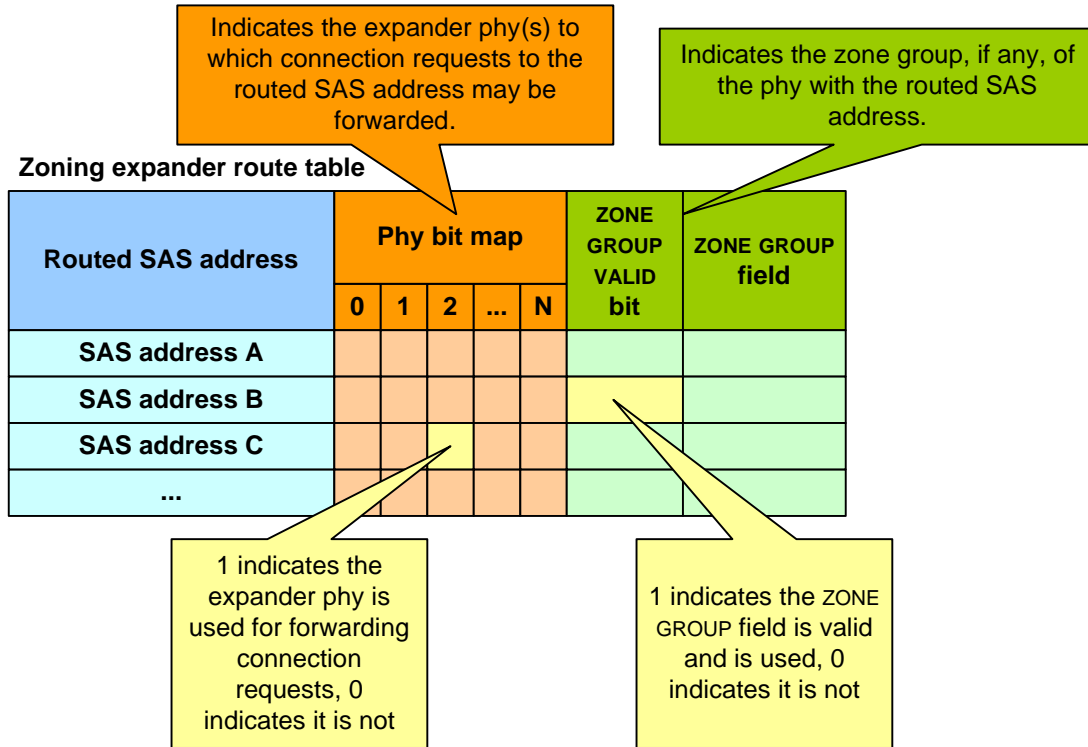


Figure 54 — Zoning expander route table

Some SMP functions may combine information from the zoning expander route table and information from other sources in a single report or configure function.

4.8.3.5 Source zone group and destination zone group determination

When a zoning expander device with zoning enabled receives an OPEN address frame (see 7.8.3):

- a) the zone group of the source port (i.e., s) is identified as defined in table 26; and
- b) the zone group of the destination port (i.e., d) is identified as defined in table 27.

If the ZP[s, d] bit is set to one then access between the phys shall be permitted and the zoning expander device shall perform the ECM arbitration procedure. If the ZP[s, d] bit is set to zero then access between the phys is not permitted and the zoning expander device shall transmit OPEN_REJECT (ZONE VIOLATION) in response to the connection request.

Zoning expander devices with zoning enabled shall follow the rules in table 26 to determine the zone group of the source port.

Table 26 — Source zone group determination

Attribute of the expander phy that received the OPEN address frame			Source zone group
ZONE ADDRESS RESOLVED bit	Routing method	ZONE PARTICIPATING bit	
0	Any	0	Zone group of the receiving expander phy
		1	Source zone group specified by the SOURCE ZONE GROUP field in the received OPEN address frame
1	Direct	0	Zone group of the receiving expander phy
		1	Not applicable
	Subtractive	0	Zone group of the receiving expander phy
		1	Not applicable
	Table	0	Zone group stored in the zoning expander route table for the source SAS address. If the source SAS address is not found in the zoning expander route table then the zone group of the receiving expander phy.
		1	Not applicable

Zoning expander devices with zoning enabled shall follow the rules in table 27 to determine the zone group of the destination port.

Table 27 — Destination zone group determination

Routing method of the destination expander phy	Destination zone group
Direct	Zone group of the destination expander phy
Subtractive	Zone group of the destination expander phy (i.e., the subtractive expander phy)
Table	Zone group stored in the zoning expander route table for the destination SAS address

Editor's Note 3: the following text and figures were moved from 4.8.1.

Figure 55 shows an example of a ZPSDS in which the ZONE ADDRESS RESOLVED bit of the expander port of **zoning** expander device C that is attached to **non-zoning** expander device D is set to zero.

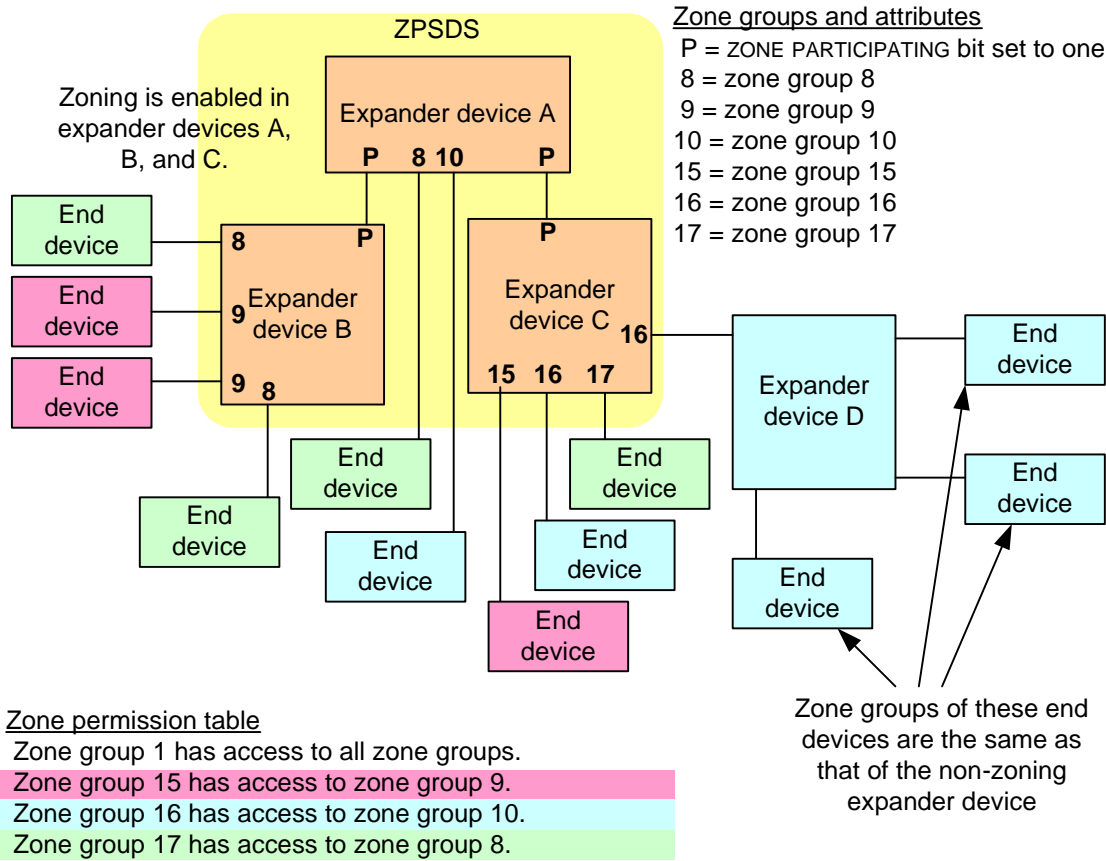


Figure 55 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to zero updated to remove “Zoning” and “Non-zoning” prefixes in front of “expander device”

Figure 56 shows an example of a ZPSDS in which the ZONE ADDRESS RESOLVED bit of the expander port of ~~zoning~~ expander device C that is attached to ~~non-zoning~~ expander device D is set to one.




Figure 56 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to one updated to remove “Zoning” and “Non-zoning” prefixes in front of “expander device”

Editor’s Note 4: End of moved figures and text

4.8.4 Zone groups and link reset sequences

If the ZONE GROUP PERSISTENT bit is set to one, then a link reset sequence (see 4.4) shall not cause the zone group value of an expander phy to change. If the ZONE GROUP PERSISTENT bit is set to zero, then table 28 specifies events based on the initial condition of an expander phy that shall cause a zoning expander device with zoning enabled to change the ZONE GROUP field of the expander phy to its default value (e.g., zero).

Table 28 — Events that cause the ZONE GROUP field to be set to its default value when the ZONE GROUP PERSISTENT bit set to zero

Initial condition	Event after the initial condition is established
Completed link reset sequence with a SAS device attached	A subsequent link reset sequence completes and: <ol style="list-style-type: none"> a) the SAS address received in the IDENTIFY address frame (see 7.8.2) during the identification sequence is different from the SAS address prior to the completion of the link reset sequence; or b) a SATA device is attached.
Completed link reset sequence with a SATA device attached	Either: <ol style="list-style-type: none"> a) A subsequent link reset sequence completes and: <ol style="list-style-type: none"> A) a hot-plug timeout (see 6.7.5) occurred between the time of the initial condition and the time the link reset sequence completed; B) the zoning expander device has detected the possibility that a new SATA device has been inserted. The method of detection is outside the scope of this standard (e.g., an enclosure services process reports a change in the ELEMENT STATUS CODE field in the Device or Array Device element (see SES-2), or a change in the WORLD WIDE NAME field in the attached SATA device's IDENTIFY DEVICE or IDENTIFY PACKET DEVICE data (see ATA8-ACS)); or C) a SAS device is attached; or b) The expander phy is disabled with the SMP PHY CONTROL function (see 10.4.3.17) DISABLE phy operation.

4.8.5 Broadcast processing in a zoning expander device with zoning enabled

The BPP determines the source zone group(s) of the Broadcast as follows:

- a) if the BPP receives a Broadcast Event Notify message from an expander phy (i.e., a zoning expander phy received a BROADCAST), the Broadcast has a single source zone group set to the zone group of that expander phy; and
- b) if the BPP receives a message from the management device server indicating that it received an SMP ZONED BROADCAST request from an SMP initiator port that has access to zone group 3, the Broadcast has each of the source zone groups specified in the SMP ZONED BROADCAST request.

The BPP forwards the Broadcast to each expander port other than the one on which the Broadcast was received (i.e., the expander port that received the BROADCAST or SMP ZONED BROADCAST request) if any of the source zone groups have access to the zone group of the expander port.

To forward a Broadcast to an expander port:

- a) if the expander port's ZONE PARTICIPATING bit is set to one, the BPP shall request that the SMP initiator port establish a connection on at least one phy in the expander port to the SMP target port of the attached expander device and transmit an SMP ZONED BROADCAST request specifying the source zone group(s); and
- b) if the expander port's ZONE PARTICIPATING bit is set to zero, the BPP shall send a Transmit Broadcast message to at least one phy in the expander port, causing it to transmit a BROADCAST.