

ENDL TEXAS

Date: 6 March 2006
 To: T10 Technical Committee
 From: Ralph O. Weber
 Subject: SAS-2 Zoning - Phy Features

Summary

During the February 23 SAS Supervisor meeting (minutes in 06-114) it was suggested that 06-107 and 06-048 be combined in a single document covering all SAS-2 Zoning changes related to the physical level. This document is that combination.

To simplify the presentation in this document, all historical text not proposed for addition to SAS-2 by this document has been removed.

Referenced Documents

SAS-2 revision 2	(http://www.t10.org/ftp/t10/drafts/sas2/sas2r02.pdf)
SAS-2 Zoning [06-019r4]	(http://www.t10.org/ftp/t10/document.06/06-019r4.pdf)
SAS-2 zoning enhancements	(http://www.t10.org/ftp/t10/document.06/06-048r3.pdf)
SAS-2 Zoning sans Mgmt	(http://www.t10.org/ftp/t10/document.06/06-107r1.pdf)

Revision History

- r0 Initial revision
- r1 Revised as per the 2 March teleconference
 - Add the 4.6.2 text from 06-048r3 but generalize to cover all zone phy information,
 - Update 4.9.5 to state requirements in terms of ports, not phy, and
 - Clarify notation description.
- r2 Revised as per the 6 March SAS Protocol working group (minutes in 06-124)

Changes between r0 and r1/r2 are marked with change bars

Notation

- Additions to SAS-2 r02 are shown in **bright red** (or black when whole sections are added);
- Deletions from SAS-2 r02 are shown in **blue strikeouts**; and
- Notes (i.e., text that is not intended for inclusion in SAS-2) are shown in **green**.

Proposed SAS-2 Changes

3 Definitions, symbols, abbreviations, keywords, and conventions

3.1 Definitions

[Start of new definitions; additions and deletions not marked]

3.1.260 zone group: A set of expander phys in a ZPSDS all of which have the same access permission.

3.1.261 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.

3.1.262 zoned portion of a service delivery subsystem (ZPSDS): A group of zoning expanders that cooperate to control phy access. The ZPSDS may include all or part of the service delivery subsystem (see 4.1.6).

[End of new definitions]

3.2 Symbols and abbreviations

[Start of new abbreviations; additions and deletions not marked]

Abbreviation	Meaning
ZP[s,d]	Zone Permission bit for source, s, and destination, d, in the zone permission table (see 4.9.3.4)
ZPSDS	Zoned Portion of a Service Delivery Subsystem

[End of new abbreviations]

4. General

...

4.6.2 Expander ports

An external expander port contains one or more physical phys (see 4.1.2). Since each phy in the expander device has the same SAS address, expander ports are created based on the attached SAS addresses (see 4.1.3).

Each phy in an expander port shall have the same routing attribute (see 4.6.7.1), and the DISCOVER function (see 10.4.3.5) shall return the same value in the ROUTING ATTRIBUTE field for each phy in an expander port.

Each phy in an expander port containing phys with table routing attributes shall have the same number of routing table entries (see 4.6.7.3).

Each phy in an expander port shall have the same zone phy information (see 4.9.3.1). The zone phy information associated with the phys in an expander port may be treated as zoning properties of the expander port.

A set of phys with table routing attributes using the same external connector (see 5.2.3.3) is called an enclosure out port. A set of phys with subtractive routing attributes using the same external connector is called an enclosure in port.

[no other changes in 4.6.2]

4.6.5 Broadcast primitive processor (BPP)

The BPP receives broadcast primitive requests from each expander phy and requests transmission of those requests on all expander ports except the expander port from which the broadcast primitive request was received.

In a self-configuring expander device (see 4.1.5), the BPP requests transmission of a BROADCAST (CHANGE) when it completes configuration (see 10.4.3.3).

In a zoning expander device (see 4.9.5), receipt of a broadcast primitive may be translated into a SMP ORIGINATE BROADCAST PRIMITIVES request or vice versa.

4.6.6.5 BPP interface

Table 15 describes the requests from an expander phy to the BPP. See 7.11 for more information on broadcasts. See 4.9.5 for more information on how zoning expander devices handle broadcasts.

[No other changes in 4.6.6.5]

[Begin New Section (i.e., subclause); additions and deletions not marked]

4.9 Zoning

4.9.1 Zoning overview

SAS zoning is implemented by a set of zoning expander devices that define the zoned portion of a service delivery subsystem (ZPSDS). The zoning expander devices control whether a phy is permitted to open a connection to another phy.

There shall be at most one ZPSDS in a SAS domain. The ZPSDS may encompass some or all of the SAS domain.

Every phy in a ZPSDS belongs to a zone group. All phys in a wide port shall belong to the same zone group. Each zoning expander contains a zone permission table that controls whether a connection is allowed between the zone group of a source port and the zone group of a target port. There are 128 zone groups.

A requested connection shall only be opened 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 zone route table is an extended version of the expander route table (see 4.6.7.3) that includes zone phy information.

Expander devices that do not support zoning may be part of the service delivery subsystem, but they are treated as being outside the boundary of the ZPSDS (i.e., non-zoning expander devices are treated in the same manner as end devices). All phys in a non-zoning expander device belong to the zone group specified for the phy that attaches the non-zoning expander device to the ZPSDS.

Figure n1 shows an example of a ZPSDS in which the zoning expander devices all have the ZONE ADDRESS RESOLVED bit set to zero.

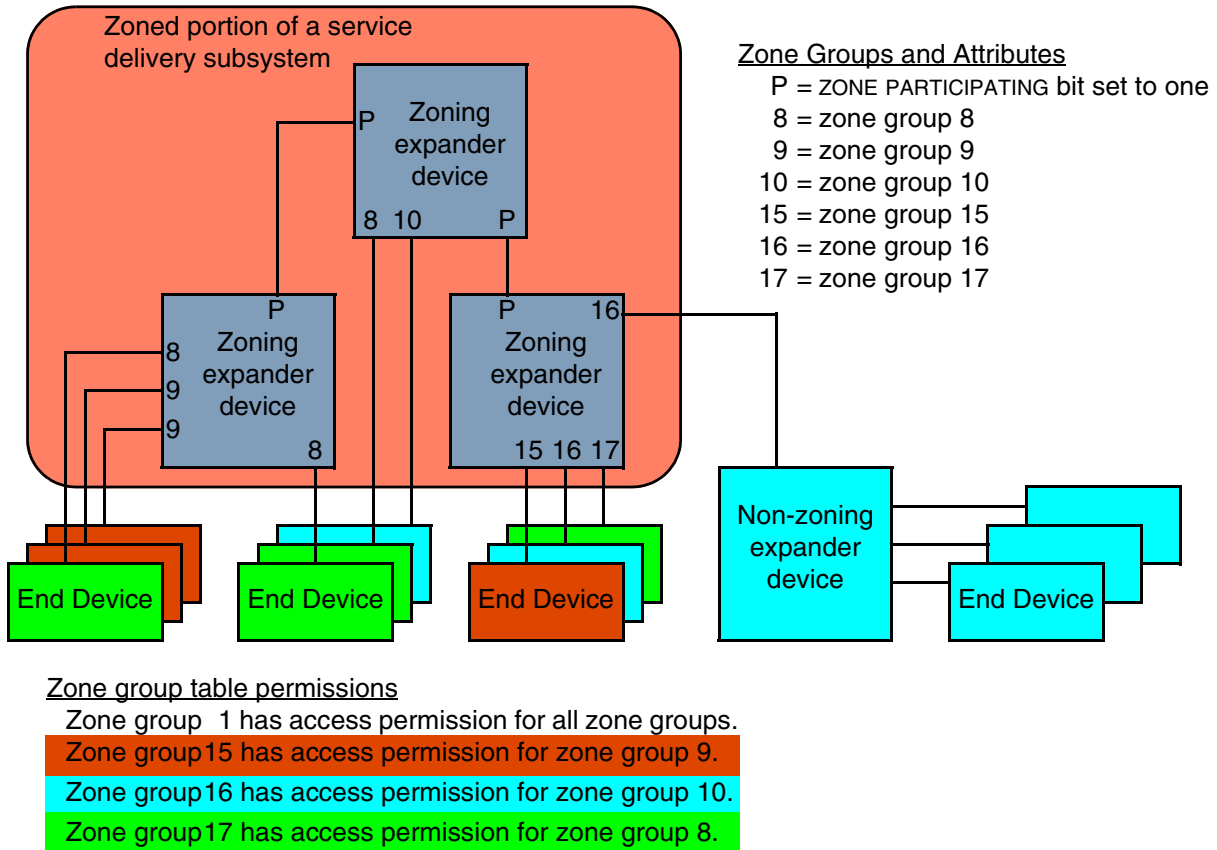


Figure n1 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to zero

Figure n2 shows an example of a ZPSDS in which the zoning expander device connected to the non-zoning expander device has the ZONE ADDRESS RESOLVED bit set to one.

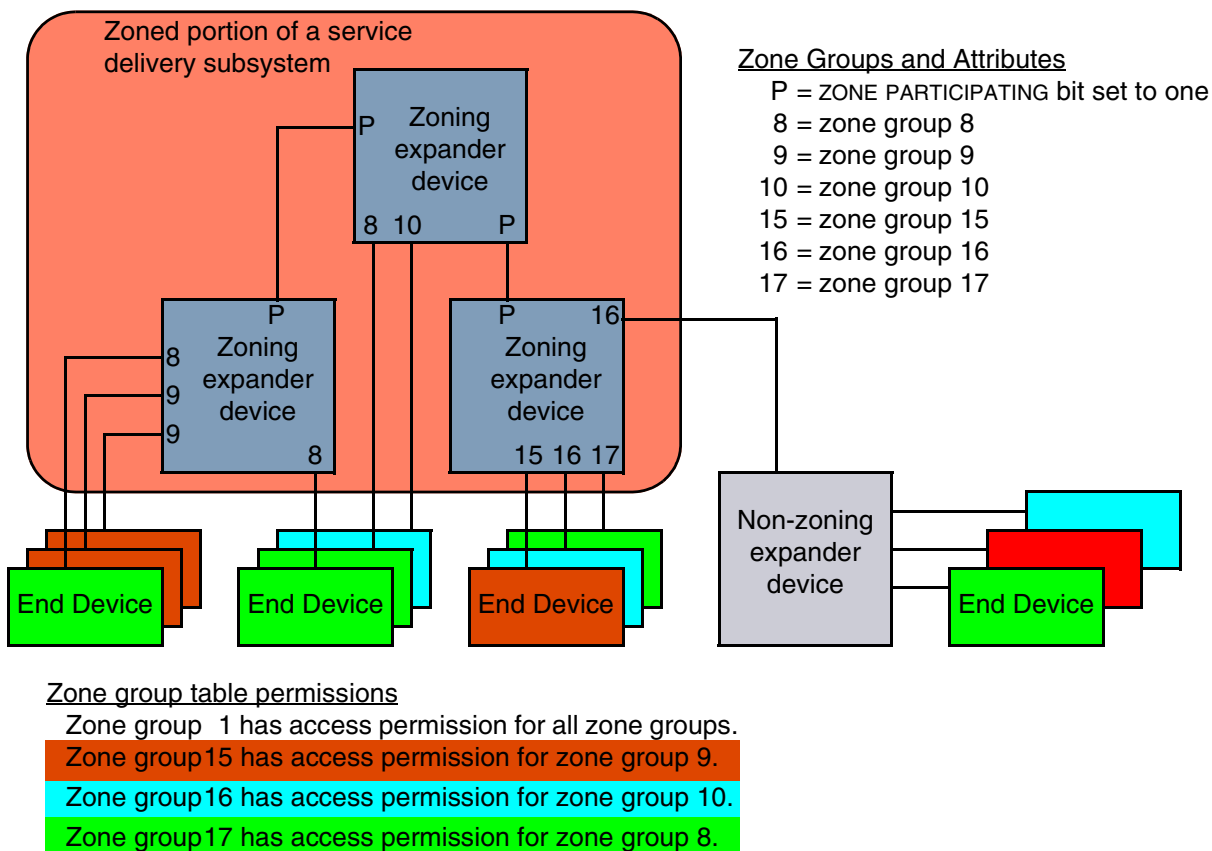


Figure n2 — Example of a ZPSDS with the ZONE ADDRESS RESOLVED bit set to one

4.9.2 Zoning expander device requirements

A zoning expander device shall have the following attributes:

- a) contain a zone expander zone route table (see 4.9.3.4);
- b) contain a zone permission table that supports 128 zone groups (see 4.9.3.3);
- c) allow or deny connection requests based on the zoning permissions;
- d) set the ZONE DEVICE bit to one in its SMP REPORT GENERAL response;
- e) be self-configuring;
- f) contain an SMP initiator and SMP target (see 4.6.1); and
- g) support all zoning related SMP functions (see TBD).

4.9.3 Zone Operation

4.9.3.1 Zone phy information

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

- a) ZONE PARTICIPATING bit;
- b) ZONE ADDRESS RESOLVED bit; and
- c) ZONE GROUP field.

The ZONE PARTICIPATING bit indicates the 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. The ZONE PARTICIPATING bit shall be set to one when the phy is attached to a zoning expander device. If the ZONE PARTICIPATING bit is set to zero, then zoning information shall not be sent on the phy and any zoning information received on the phy shall be ignored.

The ZONE ADDRESS RESOLVED bit specifies the method used to determine the source zone group for an OPEN request received by a phy at the boundary of the ZPSDS as specified in table n3.

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

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

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

- a) the phy is contained within a non-zoning expander device; or
- b) 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 0 to 127 that indicates the zone group to which the phy belongs.

All phys in an expander port have the same zone phy information (see 4.6.2).

4.9.3.2 Zone group types

4.9.3.2.1 Overview

Zone groups function in different ways based on zone group number as shown in table n1.

Table n1 — Zone group types

Zone group	Configurable	Link layer tested	Description
0, 1	No	Yes	This standard defines zone group 0 as prohibiting all accesses and zone group 1 as allowing all accesses (see 4.9.3.4).
2	Yes	No	Zone group 2 is used by the SMP application layer to restrict processing of SMP zoning configuration commands to the source phys specified by the zone permission table.
3 ... 7	Reserved		
8 ... 127	Yes	Yes	Zone groups 8 through 127 restrict access to phys.

4.9.3.3 Link layer tested zone groups

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

4.9.3.3.1 Zone group 2, SMP zoning configuration 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. A SMP target shall retain the contents of the SOURCE ZONE GROUP field from the OPEN frame until the SMP request is received. 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 function request shall not be processed and an SMP response frame with the FUNCTION RESULT field set to SMP ZONE VIOLATION shall be returned.

4.9.3.4 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 shall be rejected (see 4.9.3.6).

The zone permission table structure is shown in table n2.

Table n2 — Zone permission table

Zone group d	Zone group s								
	0	1	2	3	...	7	8	...	127
0	0	1	0	0	...	0	0	...	0
1	1	1	1	1	...	1	1	...	1
2	0	1	ZP[2,2]	Reserved	...	Reserved	ZP[8,2]	...	ZP[127,2]
3	0	1	Reserved	Reserved	...	Reserved	Reserved	...	Reserved
...
7	0		Reserved	Reserved	...	Reserved	Reserved	...	Reserved
8	0	1	ZP[2,8]	Reserved	...	Reserved	ZP[8,8]	...	ZP[127,8]
...
127	0	1	ZP[2,127]	Reserved	...	Reserved	ZP[8,127]	...	ZP[127,127]

Note: Gray 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, then 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...127,0] are set to zero).

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

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

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

4.9.3.5 Zoning expander zone route table

The zone route table in a zoning expander device shall contain zone phy information in addition to the expander phy information.

Figure 2 shows a representation of a zoning expander route table.

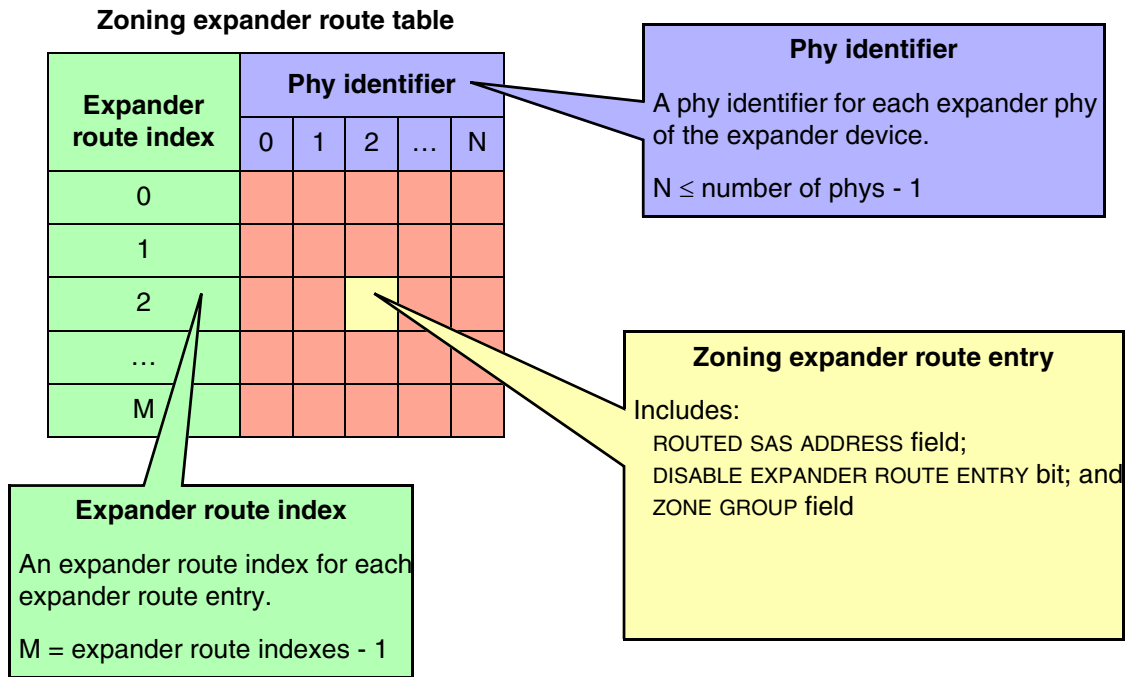


Figure n3 — Zoning expander route table example

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

4.9.3.6 Zone routing

When a zoning expander device receives an OPEN request:

- a) the zone group of the source port is identified as defined in table n3; and
- b) the zone group of the destination port is identified as defined in table n4.

If the ZP[source port zone group, destination port zone group] bit is set to one then access between the phys shall be permitted and the zoning expander shall process the ECM arbitration procedure. If the ZP[source port zone group, destination port zone group] bit is set to zero then access between the phys is not permitted and the OPEN_REJECT (ZONE VIOLATION) response shall be sent to the device that originated the OPEN request (see 7.8.3).

Zoning expander devices shall follow the rules in table n3 to determine the zone group of the source port.

Table n3 — Source zone group determination

ZONE PARTICIPATING bit	ZONE ADDRESS RESOLVED bit			
	zero	one		
		Expander phy routing method (see 4.6.7.1)		
		Direct	Subtractive	Table
zero	Zone group of the receiving phy		Zone group stored in the zone route table for the source SAS address. If the source SAS address is not found in the zone route table then zone group of the receiving phy.	
one	Source zone group in OPEN address frame.	Not applicable		

Zoning expander devices shall follow the rules in table n4 to check the zone group access permission of the DESTINATION SAS ADDRESS field in the OPEN address frame.

Table n4 — Routed zone group

Expander phy routing method (see 4.6.7.1)	Destination zone group determination
direct routing	zone group of the destination phy
table routing	zone group stored in the zone route table for the destination SAS address
subtractive routing	zone group of the subtractive phy

When an OPEN address frame is transmitted by a zoning expander device and the phy has the ZONE PARTICIPATING bit set to zero, the SOURCE ZONE GROUP field shall be set to zero.

4.9.4 Phy reset sequence

After a phy reset sequence (e.g., the hot swap of an end device) where the initial conditions shown in table n5 are met, the zone group to which the phy belongs shall be maintained according to the rules in table n5.

Table n5 — Phy reset sequence behavior

Initial condition	Event	New phy zone configuration
SAS device attached. Phy in the SP15: SAS_PHY_Ready state, with an initial phy zone configu- ration assigned	Phy exits the SP15: SAS_PHY_Ready state, and later re-enters the SP15: SAS_PHY_Ready state as the result of receiving an IDENTIFY frame with the same SAS address as prior to the exit out of SP15: SAS_PHY_Ready state.	The zoning expander device shall set the phy ZONE GROUP field to the value that it contained prior to the phy exiting the SP15: SAS_PHY_Ready state.
	Phy exits the SP15: SAS_PHY_Ready state, and later re-enters the SP15 state as the result of receiving an IDENTIFY frame with a different SAS address as prior to the exit out of SP15: SAS_PHY_Ready state or the phy enters the SP22: SATA_PHY_Ready state.	The zoning expander device shall set the phy ZONE GROUP field to zero.
SATA device attached. Phy in the SP22: SATA_PHY_Ready state, with an initial phy zone configu- ration assigned.	Phy exits the SP22: SATA_PHY_Ready state, and later re-enters the SP22: SATA_PHY_Ready state before an expander Hot-Plug Timeout timer expires.	The zoning expander device shall set the phy ZONE GROUP field to the value that it contained prior to the phy exiting the SP22: SATA_PHY_Ready state.
	Phy exits the SP22: SATA_PHY_Ready state, and later re-enters the SP22: SATA_PHY_Ready state after an expander Hot-Plug Timeout timer expires or the phy enters the SP15:SAS_PHY_Ready state.	The zoning expander device shall set the phy ZONE GROUP field to zero.

4.9.5 BROADCAST events

If a zoning expander device receives a BROADCAST primitive or an SMP ORIGINATE BROADCAST PRIMITIVES request, then the zoning expander device shall process the broadcast notification for all other ports by sending one of the following:

- a) an SMP ORIGINATE BROADCAST PRIMITIVES request on ports with the ZONE PARTICIPATING bit set to one whenever at least one zone permission bit is set to one in:
 - A) any of the zone groups listed in the received SMP ORIGINATE BROADCAST PRIMITIVES request; or
 - B) the zone group associated with port on which the BROADCAST primitive was received;
 or
- b) the relevant BROADCAST primitive as defined in 7.2.5.4 on ports with ZONE PARTICIPATING bit set to zero whenever one of the following conditions is met:
 - A) at least one zone permission bit is set to one in any of the zone groups listed in the received SMP ORIGINATE BROADCAST PRIMITIVES request; or
 - B) the zone group associated with the non-participating port matches the zone group associated with the port on which the BROADCAST primitive was received.

[End of new Section]

[Start of suggested changes to existing sections]

7.2.2 Primitive summary

Table 70 defines the primitives not specific to the type of connection.

Table 70 — Primitives not specific to type of connection ...

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
...								
OPEN_REJECT (PATHWAY BLOCKED)	NoConn		E		I		T	Single
OPEN_REJECT (PROTOCOL NOT SUPPORTED)	NoConn	I		T	I		T	Single
OPEN_REJECT (RESERVED ABANDON 0) OPEN_REJECT (ZONE VIOLATION)	NoConn		E		I		T	Single
...								

{no other changes in 7.2.2}

7.2.3 Primitive encodings

Table 73 defines the primitive encoding for primitives not specific to type of connection.

Table 73 — Primitive encoding for primitives not specific to type of connection ...

Primitive	Character			
	1 st	2 nd	3 rd	4 th (last)
...				
OPEN_REJECT (RESERVED ABANDON 0) OPEN_REJECT (ZONE VIOLATION)	K28.5	D31.4	D02.2	D27.4
...				

{no other changes in 7.2.3}

7.2.5.11 OPEN_REJECT

...

Table 85 — OPEN_REJECT abandon primitives

Primitive	Originator	Description
...		
OPEN_REJECT (RESERVED ABANDON 0) OPEN_REJECT (ZONE VIOLATION)	Zoning expander phy	Reserved. Process the same as OPEN_REJECT (WRONG DESTINATION). The connection request is from a zone group that does not have permission to access the zone group that contains the destination phy.
...		

[no other changes in 7.2.5.11]

7.8.3 OPEN address frame

Table 95 defines the OPEN address frame format used for connection requests.

Table 95 — OPEN address frame format

Byte/ Bit	7	6	5	4	3	2	1	0
0	INITIATOR PORT	PROTOCOL			ADDRESS FRAME TYPE (1h)			
1	FEATURES			CONNECTION RATE				
2	(MSB)	INITIATOR CONNECTION TAG						(LSB)
3								
4	DESTINATION SAS ADDRESS							
11								
12	SOURCE SAS ADDRESS							
19								
20	COMPATIBLE FEATURES SOURCE ZONE GROUP							
21	PATHWAY BLOCKED COUNT							
22	(MSB)	ARBITRATION WAIT TIME						(LSB)
23								
24	MORE COMPATIBLE FEATURES							
27								
28	(MSB)	CRC						(LSB)
31								

...

~~The COMPATIBLE FEATURES field shall be set to zero. The destination phy shall ignore the COMPATIBLE FEATURES field.~~

The SOURCE ZONE GROUP field identifies the zone group that contains the phy making the connection request. If the OPEN request is received on a zoning expander device phy with the ZONE PARTICIPATING bit set to zero (i.e., the source device is outside the ZPSDS), then the SOURCE ZONE GROUP field of the OPEN request shall be set to the zone group that is associated with the zoning expander device phy on which the OPEN request has been received. Zone source group values between 128 and 255, inclusive, are reserved.

10 Application layer ...

10.4.3.2 SMP function response frame format

...
 The FUNCTION RESULT field is defined in table 192.

Table 192 — FUNCTION RESULT field

Code	Name	SMP function(s)	Description
...			
04h	SMP ZONE VIOLATION	CONFIGURE PHY ZONE, CONFIGURE ZONE PERMISSION	The SMP target port supports the function, but the applicable zone permission bit is set to zero (e.g., the ZP[source device,2] bit is set to zero when a CONFIGURE PHY ZONE request is received).
...			

[no other changes in 10.4.3.2]

Annex J

...

Table J.1 — Primitives with Hamming distance of 8 ...

1 st	2 nd	3 rd	4 th	Assignment
...				
K28.5	D31.4	D02.2	D27.4	OPEN_REJECT (RESERVED ABANDON 0) OPEN_REJECT (ZONE VIOLATION)
...				

{no other changes in annex J}