

To: T10 Technical Committee
From: Tim Symons, PMC-Sierra (Tim_Symons@pmc-sierra.com)
Heng Liao, PMC-Sierra (liaoheng@pmc-sierra.com),
Tom Grieff, HP (thomas.grieff@hp.com)
Date: September 7, 2005
Subject: T10/05-144r8 SAS-2 zoning

Revision Information

- Revision 0 (9 May 2005) first revision
- Revision 1 (May 2005) second revision: move the new date field position added to REPORT GENERAL and DISCOVER that overlaps with fields that is already in use.
- Revision 2 (May 11, 2005) third revision: add description about the supervising expander election process, and permission table update through the supervising expander, changes made to SMP to support this.
- Revision 3 (June 12, 2005) fourth revision: clean up the document based on 05-195r0 SAS Ad Hoc group conference call minutes. Add the expander Supervising Allowed attribute to provide supervisor control over which expander can be a supervising expander. The SMP commands are changed/added to support this feature.
- Revision 4 (July 6, 2005): Change SUPERVISING ALLOWED to SUPERVISING PRIORITY. Editorial changes for numbered lists, tables, and reserved fields.
- Revision 5 (Aug 14,2005): Revised wording, but not context, of most sections to improve readability. Revised table and figure numbering and references to provide consistency throughout the document.
- Revision 6 (Aug 17, 2005): Added Check Source ID checking, added details on supervisor types and access for management.
- Revision 7 (Sept 7, 2005): Reorganized section ordering, Rationalized section numbers and references, Removed Source ID checking. Replaced the term "Trusted device" with "Participating Device". Added extra definition wording to discovery, election and broadcast sections.
-

Introduction

This standard defines the interfaces, functions and operations necessary to implement zoning for Serial Attached SCSI (SAS) systems. Conforming implementations may employ any design technique that does not violate interoperability. Where possible this standard uses existing SMP commands, employing reserved fields for zoning data requirements.

Table of contents

1	Scope	4
2	Normative References	4
3	Definitions	5
4	General	6
5	SAS Zoning	7
5.1	Zoned service delivery subsystem.....	7
5.2	Zoning Configuration.....	7
5.2.1	Per Expander Phy Configuration	7
5.2.2	Zoning permission table.....	8
5.3	Supervisors	9
5.3.1	Supervisor Election Process.....	9
5.3.2	Topology Discovery	9
5.3.3	REPORT GENERAL Response	10
5.3.4	DISCOVER Response	11
5.4	Device Group Reassignment	13
5.5	OPEN Address Frame	14
5.5.1	SOURCE GROUP ID.....	14
5.5.2	ACCESS ZONE MANAGEMENT	14
5.6	SMP Functions for zone management.....	17
5.6.1	Zone initialization	17
5.6.2	CONFIGURE PHY ZONE Function.....	17
5.6.3	CONFIGURE ZONE PERMISSION Function.....	19
5.6.4	REPORT ZONE PERMISSION Function	22
5.7	ZONE Management Updates.....	24
5.7.1	Expander Zone Route Table.....	24
5.7.2	REPORT ZONE ROUTE TABLE Function.....	25
6	Zone Operation	28
6.1	PHY and EXPANDER CHANGE COUNT reporting.	28
6.2	No Zone Masking	28
6.3	Broadcasts	28
	ZONED BROADCAST Address Frame.....	29
6.4	Zone Checks	30

Table of Figures

Figure 1 : Zoning Service Delivery Subsystem	6
Figure 2 Expander Zone Route Table Example	24

Table of Tables

Table 1- Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)	7
Table 2 - Zoning Permission Table.....	8
Table 3 - Zoning Expander Election Priority	9
Table 4 - REPORT GENERAL Response.....	10
Table 5 - DISCOVER Request	11
Table 6 - DISCOVER Response	12
Table 7 - PHY_ZONE Reassignment after a phy Reset Event	13
Table 8 – SMP responses to requests with access control.....	15
Table 9 - OPEN Address Frame Format	16
Table 10 - ADDRESS FRAME TYPE Field	16
Table 11 - CONFIGURE PHY_ZONE Request.....	18
Table 12 - PHY_ZONE Configuration Entry Descriptor.....	19
Table 13 - CONFIGURE PHY_ZONE Response	19
Table 14 - CONFIGURE_ZONE_PERMISSION Request.....	20
Table 15 - ZONE Permission Entry Descriptor.....	21
Table 16 - CONFIGURE_ZONE_PERMISSION Response.....	22
Table 17 - REPORT_ZONE_PERMISSION Request	22
Table 18 - REPORT_ZONE_PERMISSION Response	23
Table 19 - REPORT_ZONE_ROUTE_TABLE Request.....	25
Table 20 - REPORT_ZONE_ROUTE_TABLE Response	26
Table 21 - ZONE_ROUTE_TABLE Entry Descriptor	27
Table 22 - ZONED_BROADCAST Address Frame Format	29
Table 23 - BROADCAST_TYPE Field.....	29

1 Scope

A SAS infrastructure can be used in small storage area networks (SANs). For this application, there is a need to segregate traffic among devices the way that Fibre Channel provides with zones or Ethernet with virtual LANs.

SAS zoning is a mechanism that allows initiators and targets to be restricted to see only the portions of the topology that they have been granted permission to see based on a defined zoning permission table. SAS zoning is managed zoning expander devices, and requires no enhanced features for end devices.

2 Normative References

SAS1r09, Serial Attached SCSI 1.1 Revision 9

3 Definitions

3.1. 1 Active Supervisor: A supervisor device that has the highest election priority among all the other supervisors within the zoned service delivery subsystem. The active supervisor is the only device that can change the settings of the zone. It is responsible for propagating zone permission table updates to all zoning expanders within the zoned service delivery subsystem consistently. A zoning service delivery subsystem shall have only one active supervisor.

3.1. 2 Group ID: A per phy attribute, this 7-bit field indicates the group that the phy belongs to in the range 0 thru 127.

3.1. 3 Participating Phy: A phy that interconnects between two zoning capable devices. The port is required to generate and act upon received zone configuration information such as OPEN address frames and BROADCAST address frames with the zone information that the SAS zoning service delivery subsystem relies on for zone management and access control.

3.1. 4 Self-Configuring Expander: (Reference SAS1r09) An expander device that shall contain a management application client and SMP initiator port to perform the discover process to configure its own expander route table.

3.1. 5 Supervising priority: A per phy attribute, this 4-bit field indicates the supervisor priority of the device connected to the phy. If the phy is not connected to a supervising device, then this field shall have a value of zero.

3.1. 6 Supervisor: A zoning expander or end device that is an SMP initiator and target. It supports all SMP zoning commands and has SUPERVISING PRIORITY attribute is not equal to zero. If the supervisor is an end device it shall be directly attached to a zoning expander phy, and the SUPERVISING PRIORITY attribute will be applied to the phy connected to the supervisor capable end device.

3.1. 7 Topology: The arrangement of all devices in a SAS domain and how they are connected to each other.

3.1. 8 Zone Route Table: The structure that provides an association between routed SAS addresses to expander phy identifiers using the PHY_ZONE configuration attributes PARTICIPATING, GROUP ID, SUPERVISING PRIORITY.

3.1. 9 Zoned Broadcast: To reduce the amount of traffic within the SAS infrastructure, a zoned broadcast is restricted to just the groups that are affected by the event.

3.1. 11 Zoning Expander: A SAS expander capable of being an SMP initiator, supporting the SMP zone commands. It has a zone permission table to support at least three zones and up to 128 zones, and has a zone route table.

3.1. 12 Zoning Permission Table: The reference table defining group affiliations that determine the zones for data access and communication. This may be capable of supporting up to 128 zones. A zoning service delivery subsystem supports the number of zones of the smallest zone table in the domain.

3.1. 13 Zoning Service Delivery Subsystem: A zoning service delivery subsystem is either part of a SAS domain, or is a complete SAS domain that provides the additional services of zone management and access control services.

4 General

SAS zoning supports up to 128 zoned groups. A set of zoning expander devices interoperates to define the permission table that specifies the membership of each phy to a group. This is referred to as the zoning service delivery subsystem. Each device is managed by the attributes of its phy connection. The zoning permission table controls whether communication is allowed between multiple groups, defining a zone.

The zoning delivery subsystem is composed of zoning capable expander devices that exchange SMP functions to configure and manage it. To denote the phys that are connected to other zone capable devices, the interconnecting phys have their PARTICIPATING attribute set to one. All other phys have the PARTICIPATING attribute set to zero. This defines the boundary of the zoning delivery subsystem.

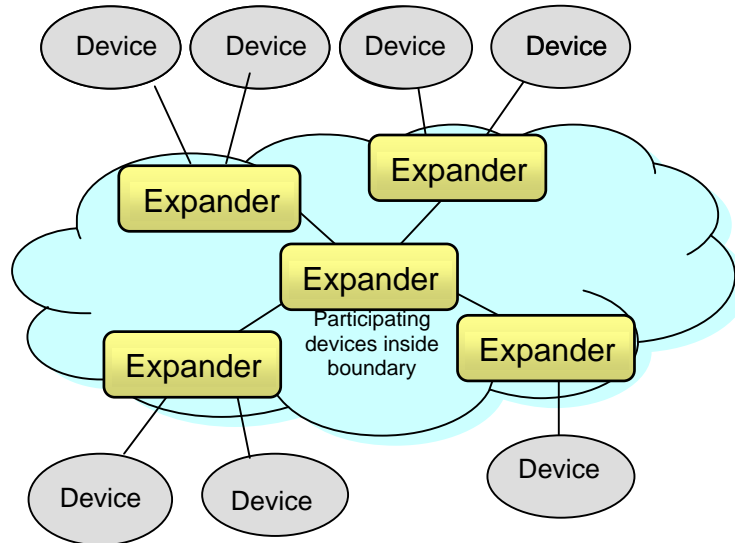


Figure 1 : Zoning Service Delivery Subsystem

SMP functions for zoning provide the means to control the zoning service delivery subsystem. This includes configuring the expander phy GROUP ID, PARTICIPATING bit, SUPERVISING PRIORITY and zoning permission table.

The zoning service delivery subsystem elects a single device to be the 'active supervisor' based on the value of its SUPERVISING PRIORITY value and SAS address value. The active supervisor is the zoning master. All configuration changes are issued by the active supervisor to the devices in the zoned service delivery subsystem. This ensures the consistency of permission tables in all the zoning expanders in the zoning service delivery subsystem. All expander devices within the zoning delivery subsystem have a complete copy of the zone permission table. Should the active supervisor fail or be removed from the system, then a replacement active supervisor shall be elected. SMP functions for zoning configuration shall only be accepted by the other zoning devices if the SMP function has originated from the active supervisor.

5 SAS Zoning

5.1 Zoned service delivery subsystem

A set of zoning capable expander devices in a SAS system may support zoning to restrict communication between SAS devices. The zone service delivery subsystem is defined by setting the participating phy attribute on phys marking the periphery of the group of zoning expander devices. Non-zoning expanders can be part of the system, but remain outside the boundary of the zoned service delivery subsystem. All phys in a wide port shall be assigned to the same group and have the same supervisor priority.

All devices outside the boundary inherit the group assignment of the boundary phy to which they are connected. A phy connecting to a non-zoning capable expander shall have its PARTICIPATING bit set to zero. Up to 128 zoning groups are supported, but the actual number of zones each expander device supports is vendor-specific. The number of zones supported by each device is reported in the REPORT GENERAL response (see Table 4). A zoned service delivery subsystem can only support the number of zones supported by the device with the smallest zone count.

5.2 Zoning Configuration

Zoning expander configuration consists of two parts:

- Per-expander phy configuration,
- Per-expander zoning permission table.

5.2.1 Per Expander Phy Configuration

Every expander phy is associated with the per-phy configuration parameters in Table 1. The information about any phy is reported in the DISCOVER response (see Table 6).

Table 1- Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)

Name	Description
PARTICIPATING	If set to zero, this phy is on the boundary of the zoning service delivery subsystem. All messages (primitives and frames) that come across this phy shall be mapped to be compatible with non-zoning devices. If set to one, this phy is inside the service delivery subsystem boundary. Zoning configuration frames and SMP commands issued by the active supervisor can pass through this phy.
SUPERVISING PRIORITY [3:0]	The election priority for the supervisor device attached to the phy. This field is used in the active supervisor selection process. If the device is not a supervisor device or is not to be elected then this value shall be zero. '1111b' is the highest priority.
GROUP ID [6:0]	The GROUP ID defines the zoning group ID in the range from 0 to 127.

5.2.2 Zoning permission table

The zoning permission table has up to 128 entries with each entry (a row of the permission table) corresponding to one group (one column of the permission table).

Table 2 - Zoning Permission Table

Group X Group Y	0	1	2	3	4	5	...	126	127
0	0	0	0	0	0	0	...	0	1
1	0	P[1,1]	P[2,1]	P[3,1]	P[126,1]	1
2	0	P[1,2]							1
3	0								1
4	0								1
5	0								1
...
126	0	P[1,126]					...	P[126,126]	1
127	1	1	1	1	1	1	...	1	1

The shaded area of Table 2 shows the user-defined permissions among groups 1 to 126.

P[X,Y] refers to permission bit Y of entry X:

- If P[X,Y] is set to one, then Group X has permission to access Group Y;
- If P[X,Y] is set to zero, then Group X has no permission to access Group Y

The access permission between groups are always symmetrical, therefore, setting access permission from Group X to Group Y automatically sets the permission from Group Y to Group X. In other words, P[X,Y] always equals P[Y,X].

Group 0 is not allowed to access any other group except Group 127:

- P[0, 0...126] shall be set to zero.
- P[0..126, 0] shall be set to zero.

Group 127 is allowed to access all other groups. Therefore:

- P[0..127, 127] shall be set to all ones
- P[127, 0..127] shall be set to one.

5.3 Supervisors

5.3.1 Supervisor Election Process

When a SAS end device is attached to an expander phy with the SUPERVISING PRIORITY value greater than zero, the device is eligible for active supervisor privilege to perform zoning changes. When the SUPERVISING PRIORITY field is set to zero, the expander is not allowed to participate in the election process for active supervising device selection.

Multiple phys and expanders in the zoning domain can have a SUPERVISING PRIORITY value greater than zero thus allowing redundant supervisors to coexist in a system. Only one shall be elected to be the active supervisor. During the active supervisor election process the SUPERVISING PRIORITY value gets pre-pended to a candidate device's SAS address. The device with the highest value in the 68-bit field defined by [SUPERVISING PRIORITY/SAS address], shall be elected as the active supervisor. The address of the elected device shall be reported as the SUPERVISING EXPANDER ADDRESS in the REPORT GENERAL response. If there are no devices available for election then the SUPERVISING EXPANDER ADDRESS shall be set to zero to indicate the election process has failed.

The election priority of a zoning expander is defined in Table 3.

Table 3 - Zoning Expander Election Priority

MSB	LSB
Bits [67:64]	Bits [63:0]
SUPERVISING PRIORITY	EXPANDER SAS ADDRESS

The supervising expander address is recalculated every time a zoning expander performs discovery. Since the address of the active supervisor is reported by the SMP REPORT GENERAL response frame, any device can use the REPORT GENERAL request to obtain the active supervisor address.

5.3.2 Topology Discovery

All zoning expanders shall be self-configuring expanders that are responsible for configuring the zone route table without host intervention. The zone route table is an extension of the routing table defined by SAS 1.1 (Section 4.7) with the additional information for phy zone assignment of the associated device. To discover the full SAS domain, a supervisor shall issue the DISCOVER request with the NO_ZONE_MASKING bit set to one (see Table 5), called PHYSICAL DISCOVERY. When the DISCOVERY request is made with the NO_ZONE_MASKING bit is set to zero, then only the devices accessible to the initiators group shall be reported, called LOGICAL DISCOVERY.

The zoning expander traverses the SAS topology and uses the REPORT GENERAL, DISCOVER, and REPORT ZONE ROUTE TABLE commands to gather topology information from adjacent expanders to populate the zone route table entries as defined in SAS1r09. When the topology discovery is complete, then the expanders will elect the active supervisor, and issue to BROADCAST (CHANGE) with group ID 127 to update all devices with the active supervisor address..

5.3.3 REPORT GENERAL Response

Table 4 defines the response frame to the REPORT GENERAL request.

Table 4 - REPORT GENERAL Response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (00h)							
2	FUNCTION RESULT							
3	Reserved							
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)
5								
6	(MSB)	EXPANDER ROUTE INDEXES						(LSB)
7								
8	NUMBER OF ZONES							
9	NUMBER OF PHYS							
10	Reserved				Reserved		CONFIGURIN G	CONFIGURABL E ROUTE TABLE
11	Reserved				SUPERVISING PRIORITY			
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20								
27	ACTIVE SUPERVISOR SAS ADDRESS							
28	(MSB)	CRC						(LSB)
31								

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

The FUNCTION field shall be set to 00h.

The NUMBER OF ZONES field indicates the number of zones supported. For expanders that do not support zoning, this field shall be set to zero. Group 0 and Group 127 shall be supported in all zoning expanders. The user-defined group ID shall range from one to the (NUMBER OF ZONES - 2).

The SUPERVISING PRIORITY field is a 4-bit address field. This value enables a system administrator to prioritize the device to be selected as the active supervisor.

The SUPERVISING EXPANDER ADDRESS field indicates the SAS address of the active supervisor that is elected during topology discovery. The election process has failed to elect an Active Supervisor if both the ACTIVE SUPERVISOR ADDRESS and CONFIGURING fields are zero.

Table 5 - DISCOVER Request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (10h)								
2	Reserved								
3									
4	Ignored								
7									
8	Reserved							NO_ZONE_MASKING	
9	PHY IDENTIFIER								
10	Ignored								
11	Reserved								
12	(MSB)	CRC							
15								(LSB)	

NO_ZONE_MASKING bit shall be set to one to request information to be reported on any phy in the topology without any zone permission restriction called PHYSICAL DISCOVERY. When the NO_ZONE_MASKING bit is set to zero, then only the phys permitted access by the zone permission table shall be reported, called LOGICAL DISCOVERY

5.3.4 DISCOVER Response

The DISCOVER function is defined in SAS1r09 : Section 10.4.3.5

The DISCOVER response function responds with details of the physical link configuration information for the specified phy. This SMP function provides information from the IDENTIFY address frame received by the phy, zone membership assignment information, and additional phy-specific information. This SMP function shall be implemented by all SMP target ports.

If the DISCOVER request is issued by a device that does not have group access privilege to all the phys, then the response frame for the phys that cannot be accessed shall respond with all the zone information set to zero, indicating that the specified phy is VACANT.

Table 6 defines the format of the DISCOVER response.

Table 6 - DISCOVER Response

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	Reserved							
4	Ignored							
7								
8	Reserved							
9	PHY IDENTIFIER							
10	Ignored							
11	Reserved							
12	Ignored	ATTACHED DEVICE TYPE			Ignored			
13	Reserved				NEGOTIATED PHYSICAL LINK RATE			
14	Reserved				ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15					ATTACHED SATA PORT SELECTOR	ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET
16	SAS ADDRESS							
23	ATTACHED SAS ADDRESS							
24	ATTACHED SAS ADDRESS							
31	ATTACHED SAS ADDRESS							
32	ATTACHED PHY IDENTIFIER							
33	Reserved							
39								
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	Reserved				ROUTING ATTRIBUTE			
45	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved		ZONE VIOLATION	PARTICIPATING	SUPERVISING PRIORITY			
49	Reserved	GROUP ID						
50	(MSB)	Vendor Specific						(LSB)
51								
52	(MSB)	CRC						(LSB)
55								

NO_ZONE_MASK field shall be to one if the function is responding to a DISCOVERY request that had the NO_ZONE_MASK bit set to one. This is a confirmation that the response is a domain wide report.

If the NO-ZONE_MASKING bit is set to zero, the response is subject to the GROUP ID restrictions defined in the zoning permission table (see Table 2).

The ZONE VIOLATION field shall be set to one if any zone violation has occurred causing the specified phy to send an OPEN_REJECT(ZONE VIOLATION) abandon primitive. The ZONE VIOLATION field shall be cleared if a PHY CONTROL function with an operation code of CLEAR ERROR LOG for the specified phy received from a supervisor.

A primitive type OPEN_REJECT(ZONING VIOLATION) is defined to replace OPEN_REJECT(RESERVED ABANDON 0) for this purpose. (See SAS1r09 Section 7.2.5.11 OPEN_REJECT).

5.4 Device Group Reassignment

The group assignment of a device is defined by the PHY ZONE configuration that is associated with the attached expander phy. After a phy reset event, such as hot swap of an end device, the PHY ZONE configuration shall be maintained by the zoned service delivery subsystem according to the rules defined in the PHY ZONE Reassignment table (See Table 7).

Table 7 - PHY ZONE Reassignment after a phy Reset Event

Initial Condition	Event	New PHY ZONE Configuration
SAS device attached, Expander phy in the SP15 state, Initial PHY ZONE configuration assigned	Phy exits the SP15 state, and later re-enters the SP15 state, receiving an IDENTIFY frame with the same SAS address as before.	Expander restores the PHY ZONE configuration.
	Phy exits the SP15 state, and later re-enters the SP15 state, receiving an IDENTIFY frame with a different SAS address as before.	Expander assigns the phy to the default Group 0.
SATA device attached, Expander phy in the SP22 state, Initial PHY ZONE configuration assigned	Phy exits the SP22 state, and later re-enters the SP22 state without having a Hot-Plug Timeout timer expiration event in-between.	Expander restores the PHY ZONE configuration.
	Phy exits the SP22 state, and later re-enters the SP22 state with a Hot-Plug Timeout timer expiration event in-between.	Expander assigns the phy to the default Group 0.

5.5 OPEN Address Frame

5.5.1 SOURCE GROUP ID

The SOURCE GROUP ID field in the OPEN address frame (see Table 9) allows the expander's SMP target port to respond to SMP commands according to the device group assignment. A DISCOVER request, with the NO_ZONE_MASKING bit set to zero (see Table 5), shall only report the phys that accessible by the SOURCE GROUP ID of the requestor. The remaining phys are reported as vacant.

When an OPEN address frame receives and transmits on phys whose PARTICIPATING bits are set to one, the SOURCE GROUP ID is preserved and transmitted. This mechanism allows the OPEN address frame to exchange zoning information inside the zoning service delivery subsystem.

When an OPEN address frame is transmitted on a phy with the PARTICIPATING bit is set to zero the SOURCE GROUP ID field shall be 000 0000b.

When an expander receives an OPEN address frame on a phy with the PARTICIPATING bit set to zero, it sets the SOURCE GROUP ID according to the value of the GROUP ID of the receivers expander phy.

5.5.2 ACCESS ZONE MANAGEMENT

The ACCESS ZONE MANAGEMENT bit identifies devices that are supervisor capable. These devices are allowed to make enquiries about the zoning configuration, but cannot change anything.

The ACCESS ZONE MANAGEMENT bit shall be set to one if the SUPERVISING PRIORITY attribute of the phy is greater than zero to enable any supervising device to make change requests to the zoning subsystem. The ACCESS ZONE MANAGEMENT bit is set to zero if the supervising priority is zero.

When an SMP OPEN Address frame is received, the expander shall verify that the SAS SOURCE ADDRESS matches the active supervisor address. If the initiating device is the active supervisor, then the expander SMP target supports the complete SMP command set including zone management SMP functions.

If the ACCESS ZONE MANAGEMENT bit is set to one and SAS SOURCE ADDRESS is not the active supervisor then the expander SMP target shall only support non-configuration zoning commands and shall respond with UNKNOWN SMP FUNCTION (Reference table 170 of SAS1r09) if a zoning configuration command is received.

Table 8 – SMP responses to requests with access control

Access Zone Management	0	0	1	1	1
SGID has access to DGID	0	1	0	0	1
No Zone Mask	X	X	0	1	X
Report General	Expander Change Count is SGID dependent				
Report Mfr Info	No Effect				
Configure Route Info	Function Failed				
Discover					
Report Phy Error Log		zoning fields not populated			
Report Phy SATA					
Report Route Info	Vacant	zoning fields not populated. Routes with GIDs that cannot access the SGID are reported as VACANT.	Vacant	Full	Full
Phy Control					
Phy Test		Full			
Report Zone Permission			Full		
Report Zone Route Table	Unknown SMP function	Unknown SMP function	Vacant		Routes with GIDs that cannot access the SGID are reported as VACANT.
Configure Phy Zone			Function Failed (see note)	Function Failed (see note)	Function Failed (see note)
Configure Zone Permission					

NOTE 1. The ACTIVE SUPERVISOR ADDRESS is only checked for configure phy zone and configure zone permission requests.

NOTE 2. If the source address is the ACTIVE SUPERVISOR and the ACCESS ZONE MANAGEMENT bit is set this is a FULL response. The SOURCE GROUP ID access field and NO ZONE MASKING bit are not checked.

The SMP REPORT ROUTE INFORMATION (Reference SAS1r09 section 10.4.3.8) and SMP CONFIGURE ROUTE INFORMATION (Reference SAS1r09 section 10.4.3.9) requests can only access the routing table entries that the SOURCE GROUP ID is allowed to access, as defined by the zone permission table.

Table 9 The OPEN Address Frame Format defines the OPEN address frame format used for connection requests.

Table 9 - 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	ACCESS ZONE MANAGEMENT	SOURCE GROUP ID						
21	PATHWAY BLOCK COUNT							
22	(MSB)	ARBITRATION WAIT TIME						(LSB)
23								
24	MORE COMPATIBLE FEATURES							
27								
28	(MSB)	CRC						(LSB)
31								

The ACCESS ZONE MANAGEMENT bit defines whether the OPEN address frame has originated from a supervisor device.

The SOURCE GROUP ID field defines the group assignment of the device that issued the OPEN address frame.

The ADDRESS FRAME TYPE field specifies the type of address frame. This field determines the definition of the frame type dependent bytes, defined in Table 10

Table 10 - ADDRESS FRAME TYPE Field

Code	Frame Type	Description
0h	Identify	Identification sequence
1h	Open	Connection request
2h	Broadcast	Broadcasts within zones
All others	Reserved	

5.6 SMP Functions for zone management

5.6.1 Zone initialization

SMP zoning commands support zone management and topology discovery functions. A zoning expander shall report the number of zones that it supports in the SMP REPORT GENERAL response (see Table 4). During the discover process the expander with the smallest number of zones supported determines the number of zones for the zoned service delivery subsystem. If a device is added to an active zoned service delivery subsystem, and it has a lesser number of supported zones than the established zoning delivery subsystem then the active supervisor shall set the device PARTICIPATING bit to zero, and the device shall be excluded from the zoned service delivery subsystem and treated as a non-zone aware device.

The CONFIGURE ZONE PERMISSION request and response allow the active supervisor to update the zone permission table. The active supervisor shall send the CONFIGURE ZONE PERMISSION request to all zoning expanders in the topology to propagate the zone permission table changes consistently. If the active supervisor has not completed the propagation of the current permission table update, it shall respond to any new updates with FUNCTION FAILED.

The CONFIGURE ZONE PERMISSION request supports both a batch mode and single mode operation. The single mode sets the value for the permission between a pair of group IDs and ensures the permission for the reverse direction is set symmetrically. The batch mode allows multiple rows of the permission table to be downloaded by one command. The supervisor shall ensure the symmetry of the permission table when it downloads the permission table in batch mode. The zoning expander is not required to check and enforce the symmetry of the zone permission table. The active supervisor shall issue a BROADCAST message to the affected groups when the zone change has been completed. The REPORT ZONE ROUTE TABLE response reports the zone route table details. The ZONE route table is logically organized in a similar way as the SAS 1.1 routing table. Each entry of the table is extended to contain the following fields in addition to the attached SAS address: PARTICIPATING, SUPERVISING PRIORITY, GROUP ID, and ATTACHED DEVICE TYPE.

5.6.2 CONFIGURE PHY ZONE Function

The CONFIGURE PHY ZONE request and response (See Table 11 and Table 13) allows the active supervisor to change the PHY ZONE configuration of a specific phy. The zoning expander shall only perform this command if the SAS address of the device that initiated the OPEN Request Frame matches that of the Active Supervisor. Otherwise, the zoning expander shall ignore the command and respond with UNKNOWN SMP FUNCTION. The zoning expander that performs the CONFIGURE PHY ZONE request shall send out a BROADCAST(CHANGE) message to the affected groups. (See Table 8)

The CONFIGURE PHY ZONE function sets the GROUP ID field, PARTICIPATION bit, and SUPERVISING PRIORITY field. The CONFIGURE PHY ZONE can set the attributes of a single phy or multiple phys. with contiguous phy indexes. The active supervisor shall send out a BROADCAST(CHANGE) message with the SOURCE GROUP ID of 127 to all zones after the CONFIGURE PHY ZONE function is completed.

The SMP target shall only process the CONFIGURE PHY ZONE request if it originates from the active supervisor.

If the ACCESS ZONE MANAGEMENT bit set to one but the initiators SAS address is not the active supervisor then the SMP target shall respond with FUNCTION FAILED.

If the ACCESS ZONE MANAGEMENT bit set to zero the SMP target shall respond with UNKNOWN SMP FUNCTION.

Table 11 defines the CONFIGURE PHY ZONE request format.

Table 11 - CONFIGURE PHY ZONE Request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (93h)								
2	Reserved								
3									
4	Reserved								
5									
6	UPDATE COMPLET E	START PHY INDEX							
7	NUMBER OF ZONE PHY ENTRIES								
	PHY ZONE configuration entry list								
8	First PHY ZONE configuration entry descriptor								
9									
...	...								
m-1	Last PHY ZONE configuration entry descriptor								
m									
	Fill bytes, if needed								
n-3	(MSB)	CRC							
n								(LSB)	

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

The FUNCTION field shall be set to 93h.SAS1r09

The CRC field is defined in SAS1r09

The START PHY INDEX field defines the first phy index to be configured by the CONFIGURE PHY ZONE command. Note that this command configures one or multiple contiguous expander phys starting from the START PHY INDEX.

The NUMBER OF ZONE PHY ENTRIES field defines how many phy zone entries the CONFIGURE PHY ZONE request intends to configure from 1 to 255 phys.

The PHY ZONE entry descriptor list contains one or more PHY ZONE entry descriptors.

The UPDATE COMPLETE bit indicates whether the current CONFIGURE PHY ZONE request is the last command of a sequence of CONFIGURE PHY ZONE requests. The zoning expander shall send out BROADCAST(CHANGE) messages to the affected zones after a CONFIGURE PHY ZONE request with UPDATE COMPLETE bit is set to one.

Table 12 defines the PHY ZONE configuration entry descriptor.

Table 12 - PHY ZONE Configuration Entry Descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved			PARTICIPATING	SUPERVISING PRIORITY			
1	Reserved	GROUP ID						

The GROUP ID field specifies the group membership to be assigned to the specified phy (see Table 1).

The SUPERVISING PRIORITY field specifies whether the specified phy is a candidate for active supervisor and if so what priority it has (see Table 1).

The PARTICIPATING bit shall be set to one when the phy is connected to another phy within the zoned service delivery subsystem to indicate that the device is participating in zone management (see Table 1).

Table 13 defines the CONFIGURE PHY ZONE response format.

Table 13 - CONFIGURE PHY ZONE Response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (93h)							
2	FUNCTION RESULT							
3	Reserved							
4	(MSB)	CRC						(LSB)
7								

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

The FUNCTION field shall be set to 93h.

The FUNCTION RESULT field is defined in SAS1r09 Section 10.4.3.2.

The CRC field is defined in SAS1r09 Section 10.4.3.1.

5.6.3 CONFIGURE ZONE PERMISSION Function

The CONFIGURE ZONE PERMISSION function provides two modes for setting the values in the zoning permission table controlled by the value of the BATCH bit (see Table 11): single set and batch set mode.

Single set operation is simplest to setup and maintains a small request frame size. The value for a single permission table entry is set to the value provided by the NEW VALUE field of the CONFIGURE ZONE PERMISSION request frame. The single set operation shall apply the NEW VALUE symmetrically to both the [X,Y] and the [Y,X] entries of the permission table.

In batch set operation, the CONFIGURE ZONE PERMISSION assigns the bit value for one or multiple contiguous entries of the permission table. When the zone permission table is set in a batch operation, the active supervisor shall ensure the permission table is symmetrical.

The SMP target shall only process the CONFIGURE ZONE PERMISSION request if it originates from the active supervisor.

If the ACCESS ZONE MANAGEMENT bit set to one but the initiating device SAS address is not the active supervisor then the SMP target shall respond with FUNCTION FAILED

If the ACCESS ZONE MANAGEMENT bit set to zero the SMP target shall respond with UNKNOWN SMP FUNCTION.

Table 14 defines the CONFIGURE ZONE PERMISSION request format.

Table 14 - CONFIGURE ZONE PERMISSION Request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (80h)							
2	Reserved							
3								
4	Reserved							
5								
6	Reserved	SOURCE GROUP ID						
7	NEW VALUE	TARGET GROUP ID						
8	Reserved							
9	Reserved							UPDATE COMPLETE
10	BATCH	START ZONE ENTRY INDEX						
11	NUMBER OF ZONE PERMISSION ENTRIES							
	ZONE PERMISSION entry list							
12	First ZONE PERMISSION entry descriptor							
27								
...	...							
n-20	Last ZONE PERMISSION entry descriptor							
n-4								
n-3	(MSB)	CRC						(LSB)
n								

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

The FUNCTION field shall be set to 80h. The BATCH bit when set to one indicates that more than one contiguous permission table entry is to be updated by this function. The SOURCE GROUP ID and the TARGET GROUP ID shall be set to zero in this case. When the BATCH bit is set to zero this function sets the access permission between the specified SOURCE GROUP ID and the TARGET GROUP ID.

The NEW VALUE field is only used in single set mode. It provides the value for the permission table entry. For batch set mode, this bit shall be set to zero. The SOURCE GROUP ID and TARGET GROUP ID fields define the PERMISSION ZONE TABLE entry to be modified. It is only used in single set mode. For batch set mode, both the SOURCE GROUP ID and the TARGET GROUP ID fields shall be set to 000 0000b.

If a non-supervising expander receives CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE is set to zero, the non-supervising expander shall updates its own permission table without propagation.

If a non-supervising expander receives CONFIGURE ZONE PERMISISON command with PROPAGATE UPDATE is set to one, the expander shall ignore this command and return FUNCTION FAILED.

If a supervising expander receives CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE is set to one, the supervising expander shall update its permission table accordingly and propagate the same permission table updates to all other zoning expanders within the domain by sending CONFIGURE ZONE PERMISSION commands to individual expanders with PROPGATE UPDATE is set to zero. If the supervising expander receives another CONFIGURE ZONE PERMISSION command before the previous command has been completed, the supervising expander shall ignore the new CONFIGURE ZONE PERMISSION command and return FUNCTION FAILED.

If a supervising expander receives CONFIGURE ZONE PERMISISON command with PROPAGATE UPDATE is set to zero, the supervising expander shall ignore the command and return FUNCTION FAILED.

The UPDATE COMPLETE bit indicates whether the current CONFIGURE ZONE PERMISSION command is the last command of a sequence of CONFIGURE ZONE PERMISSION commands. This bit is valid in both batch set mode and single set mode. The zoning expander shall send out BROADCAST(CHANGE) messages to the affected zones after a CONFIGURE ZONE PERMISSION command with UPDATE COMPLETE bit is set to one.

The START ZONE ENTRY INDEX field specifies the first zone permission table entry index to be configured in batch set mode. For single set mode, the START ZONE ENTRY INDEX field shall be set to 000 0000b.

The NUMBER OF ZONE PERMISSION ENTRIES field defines how many zone permission entries the CONFIGURE ZONE PERMISSION request intends to configure. This command configures contiguous permission table entries starting from the START ZONE ENTRY INDEX field in batch set mode.

The PHY ZONE entry descriptor list contains one or more ZONE PERMISSION entry descriptors in batch set mode. For single set mode, the PHY ZONE entry descriptor list shall not be present. Note that *n* (total number of bytes) is required to be equal to or less than 1032. This limits the number of CONFIGURE ZONE PERMISSION changes to 63.

The CRC field is defined in SAS1r09Section 10.4.3.1. Table 15 defines the ZONE PERMISSION entry descriptor.

Table 15 - ZONE Permission Entry Descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	ZONE PERMISSION							P[0,Y]
15	P[127,Y]							P[120,Y]

The ZONE PERMISSION field is the zoning permission entry as defined in Table 2. The bits corresponding to unused zone indexes shall be set to zero.

Table 16 defines the CONFIGURE ZONE PERMISSION response format.

Table 16 - CONFIGURE ZONE PERMISSION Response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (03h)							
2	FUNCTION RESULT							
3	Reserved							
4	(MSB) _____							
7	CRC _____ (LSB)							

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

The FUNCTION field shall be set to 03h.

The FUNCTION RESULT field is defined in SAS1r09 Section 10.4.3.2.

The CRC field is defined in SAS1r09 Section 10.4.3.1

5.6.4 REPORT ZONE PERMISSION Function

The REPORT ZONE PERMISSION function reports the zoning permission table of the expander.

Table 17 defines the REPORT ZONE PERMISSION request format.

Table 17 - REPORT ZONE PERMISSION Request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (14h)							
2	Reserved							
3								
4								
5	Reserved							
6	START ZONE ENTRY INDEX							
7	NUMBER OF ZONE PERMISSION ENTRIES							
8	(MSB) _____ CRC _____ (LSB)							
11								

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

The FUNCTION field shall be set to 14h.

The START ZONE ENTRY INDEX field specifies the first zone permission table entry index to be reported.

The NUMBER OF ZONE PERMISSION ENTRIES field defines how many zone permission entries the REPORT ZONE PERMISSION request intends to report. This command reports contiguous permission table entries starting from the START ZONE ENTRY INDEX field.

The CRC field is defined in SAS1r09 Section 10.4.3.1

Table 18 defines the REPORT ZONE PERMISSION response format.

Table 18 - REPORT ZONE PERMISSION Response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (14h)								
2	FUNCTION RESULT								
3	Reserved						CONFIGU RING	Reserved	
4	Reserved								
5									
6	START ZONE ENTRY INDEX								
7	NUMBER OF ZONE PERMISSION ENTRIES								
	ZONE PERMISSION entry list								
8	First ZONE PERMISSION entry descriptor								
23									
...	...								
n-19	Last ZONE PERMISSION entry descriptor								
n-4									
n-3	(MSB)	CRC							
n							(LSB)		

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

The FUNCTION field shall be set to 14h. The FUNCTION RESULT field is defined in Section 10.4.3.2.

The CONFIGURING field indicates the expander is in the process of zone permission table update. The expander shall issue a BROADCAST message when the update is completed.

The START ZONE ENTRY INDEX field specifies the first zone permission table entry index of the first zone permission entry contained in this response frame.

The NUMBER OF ZONE PERMISSION ENTRIES field defines the number of zone permission entries in the response frame. The response frame contains contiguous permission table entries starting from the START ZONE ENTRY INDEX field.

The PHY ZONE entry descriptor list contains zero or more ZONE PERMISSION entry descriptors. Note that n (total number of bytes) is required to be equal to or less than 1032. This limits the number of REPORT ZONE changes to 63.

The CRC field is defined in SAS1r09 Section 10.4.3.1

5.7 ZONE Management Updates

There are two types of zone configuration updates: the phy zone update and the zone permission table update. The zoning expanders shall implement both types of updates without affecting the existing connection or the OPEN requests that are being arbitrated.

The phy zone update operation is an atomic operation sent from a supervisor to a zoning expander with a single SMP command. The new phy zone configuration associated with the specific phy takes effect immediately after the CONFIGURE PHY ZONE command is performed. The active supervisor is required to re-download the permission table whenever it detects a topology change in a zoned service delivery subsystem. That is, when an expander(s) is removed from or added to the SAS domain, the active supervisor shall refresh the permission table to the whole zoned service delivery subsystem. This prevents the addition or removal of expanders in a SAS domain from creating inconsistency in the permission tables. Zone configuration updates shall not affect existing connections.

5.7.1 Expander Zone Route Table

The expander zone route table is a structure that provides an association between routed SAS addresses to expander phy identifiers using the phy zone configuration attributes PARTICIPATING, SUPERVISOR, GROUP ID, SUPERVISING PRIORITY, and the attached device type. Each association is represented by one expander route entry.

An expander device reports the size of its expander route table and indicates if the table is configurable in the SMP REPORT GENERAL response (see Table 4). Each expander route entry shall be disabled after power on.

A management application client may report a specific expander route entry within an expander route table in response to the SMP REPORT ZONE ROUTE TABLE request (see Table 19). All zoning expanders are responsible for configuring the zone route table automatically through the topology discovery process described in the SAS1r09 Topology Discovery Section 4.3.1.

Figure 2 shows a representation of an expander route table.

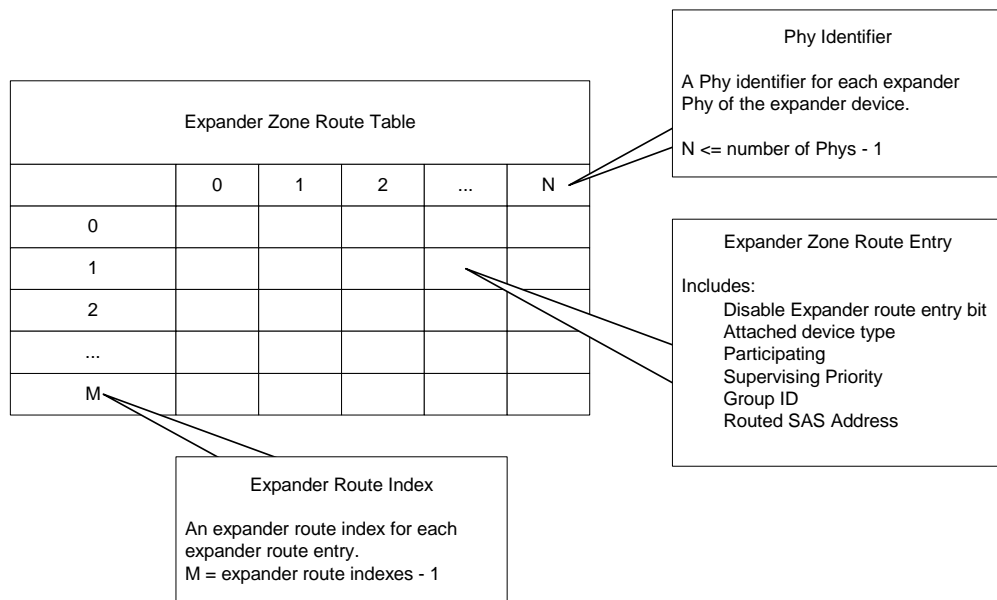


Figure 2 Expander Zone Route Table Example

5.7.2 REPORT ZONE ROUTE TABLE Function

The REPORT ZONE ROUTE TABLE is an SMP function that requires an expander device to respond with zone information about each phy of the device. If the NUMBER OF ZONES supported field is set to a value greater than zero and the EXPANDER ROUTE INDEXES field is set to a value greater than zero in the REPORT GENERAL response function then the expander device shall provide the SMP REPORT ZONE ROUTE TABLE Response. The SMP target shall only perform the REPORT ZONE ROUTE TABLE function if the OPEN address frame that established the SMP connection had its ACCESS ZONE MANAGEMENT bit set to one. If the ACCESS ZONE MANAGEMENT bit is set to zero, the SMP target shall generate an SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION. If the Phy Identifier field in the Report Zone Route Table Request has a Group ID that is not accessible to the device that initiated the Open Request Frame then this function shall respond with phy vacant.

Table 19 defines the REPORT ZONE ROUTE TABLE request format.

Table 19 - REPORT ZONE ROUTE TABLE Request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (14h)								
2	Reserved								
3									
4	NUMBER OF ZONE ROUTE TABLE ENTRIES								
5	PHY IDENTIFIER								
6	(MSB)	STARTING EXPANDER ROUTE INDEX						(LSB)	
7									
8	Reserved								
11									
12	(MSB)	CRC						(LSB)	
15									

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

The FUNCTION field shall be set to 14h.

The NUMBER OF ZONE ROUTE ENTRIES defines how many zone route table entries the REPORT ZONE ROUTE TABLE request intends to read. This command reads the zone route table entries with the contiguous expander route index starting from the value specified by the STARTING EXPANDER ROUTE INDEX field for the phy specified by the PHY IDENTIFIER field

The PHY IDENTIFIER field specifies the phy for that the expander route entry is reporting. (See Section 4.6.7.3).

The STARTING EXPANDER ROUTE INDEX field specifies the first expander route index for the expander route entry being reported The CRC field is defined in Section 10.4.3.1.

Table 20 defines the REPORT ZONE ROUTE TABLE response format.

Table 20 - REPORT ZONE ROUTE TABLE Response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (14h)								
2	FUNCTION RESULT								
3	Reserved								
4	NUMBER OF ZONE ROUTE TABLE ENTRIES								
5	PHY IDENTIFIER								
6	(MSB)	STARTING EXPANDER ROUTE INDEX							
7								(LSB)	
8	RESERVED								
9	RESERVED								
10	Reserved								
11	Reserved						CONFIGU RING	END OF ENTRIES	
ZONE ROUTE TABLE entry list									
12	First ZONE ROUTE TABLE entry descriptor								
23	Last ZONE ROUTE TABLE entry descriptor								
...	...								
n-15	Last ZONE ROUTE TABLE entry descriptor								
n-4	Last ZONE ROUTE TABLE entry descriptor								
n-3	(MSB)	CRC							
n								(LSB)	

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

The FUNCTION field shall be set to 14h.

The FUNCTION RESULT field is defined in SAS1r09 Section 10.4.3.2.

The NUMBER OF ZONE ROUTE ENTRIES field indicates how many zone route table entries the REPORT ZONE ROUTE TABLE response frame contains.

The PHY IDENTIFIER field indicates the phy that the expander route entry is reporting (see Section 4.6.7.3).

The END OF ENTRIES field indicates whether the response frame contains the last enabled zoning route table entry of the request phy.

The ZONE ROUTE TABLE entry descriptor list contains zero or more Zone Route Table Entry descriptors.

The CRC field is defined in SAS1r09 Section 10.4.3.1

The Zone Route Table Entry descriptor format is defined in Table 21

Table 21 - ZONE ROUTE TABLE Entry Descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	DISABLE EXPANDER ROUTE ENTRY	ATTACHED DEVICE TYPE			Reserved			
1	Reserved			PARTICIPATING	SUPERVISING PRIORITY			
2	Reserved	GROUP ID						
3	Reserved							
7								
8	ROUTED SAS ADDRESS							
15								

The DISABLE EXPANDER ROUTE ENTRY bit specifies whether the ECM shall use the expander route entry to route connection requests (see Section 4.6.7.3). If the DISABLE EXPANDER ROUTE ENTRY bit is set to zero, then the ECM shall use the expander route entry to route connection requests. If the DISABLE EXPANDER ROUTE ENTRY bit is set to one, the ECM shall not use the expander route entry to route connection requests.

The ROUTED SAS ADDRESS field contains the routed SAS address for the expander route entry being configured (see 4.6.7.3).

The GROUP ID field contains the GROUP ID for the expander route entry being configured (see Section 4.6.7.3).

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence. DEVICE TYPE values are defined in SAS1r09 Table 85

The SUPERVISING PRIORITY is a 4-bit field that prioritizes each supervisor expander device in a system for election as the supervising expander device. This is a system configuration feature.

6 Zone Operation

6.1 PHY and EXPANDER CHANGE COUNT reporting.

To support topology change notification, zoning expanders shall increment the EXPANDER CHANGE COUNT as defined in SAS1r09 Section 10.4.3.5.

The PHY CHANGE COUNT field counts the number of BROADCAST (CHANGE)s originated by an expander phy as defined in SAS1r09 Section 10.4.3.5.

6.2 No Zone Masking

If the entire SAS topology information is required then the NO_ZONE_MASKING bit shall be set to one in the DISCOVER request (see Table 5) to indicate full topology data is required. Any phy that is not accessible to the SOURCE GROUP ID shall be reported as vacant.

6.3 Broadcasts

To minimize the volume of traffic within the SAS infrastructure, broadcast events can be restricted to just the groups that are affected by the event using a zoned broadcast frame (see Table 22).

Zoning expanders shall support reception of broadcast primitives and the BROADCAST address frame. Zoning expanders shall transmit broadcast primitives to devices outside the zone and they should transmit BROADCAST address frames to devices within the zoned area.

When a zoned expander detects a phy event that would spawn a BROADCAST message, it instead sends a ZONED BROADCAST request with the SOURCE GROUP ID field set to the source group ID of the phy that caused the event to the expander Broadcast Primitive Processor.

When a participating expander phy receives a BROADCAST address frame with SOURCE GROUP ID information, the BROADCAST TYPE and the SOURCE GROUP ID value from the received frame shall be sent to the expander Broadcast Primitive Processor as a ZONED BROADCAST request.

When an expander phy receives a BROADCAST primitive, a ZONED BROADCAST request shall be sent to the Broadcast Primitive Processor with the BROADCAST TYPE from the primitive and the SOURCE GROUP ID from the receiving phy.

If a BROADCAST address frame is received on a phy whose participation bit is set to zero, the zoning expander shall send the BROADCAST TYPE from the BROADCAST address frame and the SOURCE GROUP ID of the receiving phy to the Broadcast Primitive Processor as a ZONED BROADCAST request.

When the expander Broadcast Primitive Processor receives a ZONED BROADCAST request, it forwards the broadcast indication to all of the phys that are accessible by the source group of the broadcast request except the phys on the initiating device port. When transmitting a ZONED BROADCAST event on a participating phy, the expander may transmit a ZONED BROADCAST address frame with the outgoing SOURCE GROUP ID field set to the SOURCE GROUP ID field received in the Zoned Broadcast Request.

The Broadcast Primitive Processor may contain a Broadcast Queue to hold Zoned Broadcast requests. If the Broadcast Queue overflows, the Broadcast Primitive Processor may remove

entries from the queue and send Broadcast Primitives to all ports scheduled to receive the Zoned Broadcast Address Frame.

When the Broadcast Primitive Processor receives a ZONED BROADCAST request with SOURCE GROUP ID of 127, Broadcast Primitives may be sent in place of a Zoned Broadcast Address Frame.

6.3.1 ZONED BROADCAST Address Frame

Table 22 defines the address frame format for zoned broadcasts. The ZONED BROADCAST address frame is used within a zoning service delivery subsystem to propagate BROADCAST messages within zone boundaries. A ZONED BROADCAST address frame shall be transmitted three times to provide redundancy against broadcast messages that have been lost due to bit errors. The reception of a single broadcast address frame is sufficient for 'detection' of the broadcast message. A SAS device that receives repeated broadcast frames shall treat the repeated sequence of identical broadcast address frames as a single broadcast event.

Table 22 - ZONED BROADCAST Address Frame Format

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved	BROADCAST TYPE			ADDRESS FRAME TYPE (2h)			
1	Reserved							
2								
3	Reserved	SOURCE GROUP ID						
4	Reserved							
27								
28	(MSB)	CRC						(LSB)
31								

Table 23 defines the valid values of the BROADCAST TYPE field.

Table 23 - BROADCAST TYPE Field

Code	BROADCAST Primitive Represented
0h	BROADCAST (CHANGE)
1h	BROADCAST (SES)
2h	BROADCAST (RESERVED 1)
3h	BROADCAST (RESERVED 2)
4h	BROADCAST (RESERVED 3)
5h	BROADCAST (RESERVED 4)
6h	BROADCAST (RESERVED CHANGE 0)
7h	BROADCAST (RESERVED CHANGE 1)

The SOURCE GROUP ID field indicates the group ID of the broadcast request. The expander device forwards the broadcast event only to the other phys that have access permission of the SOURCE GROUP ID field defined by the expander zoning permission table. (See Table 2).

6.4 Zone Checks

During SAS zoning operation, if an OPEN address frame is received on a phy with the participation bit is set to zero, the expander sets the SOURCE GROUP ID field of the OPEN address frame to the GROUP ID value of the ingress phy. When the OPEN address frame is received on a participating phy, the SOURCE GROUP ID field in the OPEN address frame is preserved.

The SAS direct and table routing mechanism is extended to map the destination address of the OPEN address frame to the destination port, as well as the Destination group ID.

If the destination device is directly attached to the expander device, the Destination group ID is set to the group ID of the destination phy.

If the destination device is indirectly attached through a table routed port the Destination group ID is set to the group ID of the mapped destination phy.

If the destination device is indirectly attached through a subtractive routing port, the Destination group ID is set to the group ID of the subtractive port.

The expander checks the zoning permission between the Source group ID and the Destination group ID. If the zoning permission table does not have permission between Source group ID and Destination group ID, the OPEN address frame is responded to with OPEN_REJECT(ZONING VIOLATION) message. If the zoning permission table allows routing between the Source group ID and Destination group ID, the expander shall proceed with the normal ECM arbitration procedure.

This zoning check mechanism allows zoning permission checking (and routing) to be done in either a single hop or a multi-hop fashion. It is up to the system administrator to choose the approach that best satisfies the application requirement and configure the expander devices accordingly.

Note that all SMP ports in the zoning expanders shall be part of Group 127. Therefore, all OPEN address frames to the SMP port are permitted.