

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: T10/05-144r6 SAS-2 zoning

1 Revision History

- 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.
-

2 Related Documents

- Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9

3 Overview

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.

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 which phys are connected to other zone capable devices, the interconnecting phys have their PARTICIPATING attribute set to 1. All other phys have the PARTICIPATING attribute set to 0. This defines the boundary of the zoning delivery subsystem.

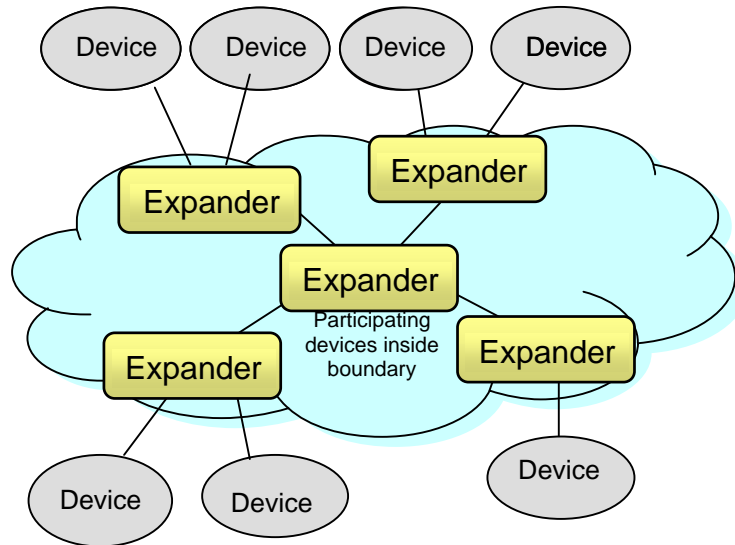


Figure 1 : Zoning Service Delivery Subsystem

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 will elect a single device to be the 'active supervisor' based on the value of its SUPERVISING PRIORITY bit field 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 will be elected. SMP functions for zoning configuration will only be accepted by the other zoning devices if the SMP function has originated from the active supervisor.

4 Definitions

4.1. 1: 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 3 zones and up to 128 zones, and has a zone route table.

4.1. 2: Zoning Permission Table: The reference table defining group affiliations which determine the zones for data access and communication. This may be capable of supporting up to 128 zones. A zoning service delivery subsystem will support the number of zones of the smallest zone table in the domain.

4.1. 3: 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.

4.1. 4: 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.

4.1. 5: Topology: The arrangement in which the devices of a SAS domain are connected to each other.

4.1. 6: 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.

4.1. 7: Supervisor: A zoning expander or end device that is an SMP initiator. It can generate SMP commands for reporting information about the SAS zoning configuration. A supervisor may be nominated to be an Active supervisor through the setting of its SUPERVISING PRIORITY bits.

4.1. 8: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 will have only one active supervisor.

4.1. 9:Expander Self-Discovery: The expander self-discovery process is an expander function for exchanging information about the SAS topology using SMP protocol with adjacent expanders in order to derive accurate knowledge about the SAS zoning topology. Such information is stored in the expander zone-route table. This process is completed automatically without intervention from any host when the expander powers up, and subsequently when the expander detects a topology change event.

4.1. 10: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.

4.1. 11: Group ID: A per phy attribute, this 7-bit field indicates the group that the phy belongs to in the range 0 thru 127.

4.1. 12: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 will have a value of zero.

5 SAS Zoning

5.1 Zoned service delivery subsystem

A set of zoning capable expander devices in a SAS topology may support zoning to restrict communication between SAS devices in that topology. 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 topology, but will remain outside the boundary of the zoned service delivery subsystem. All devices outside the boundary will inherit the group assignment of the boundary phy that they are connected to. The phy will have its PARTICIPATING bit set to 0. 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 (Table 4). A zoned service delivery subsystem can only support the number of zones supported by the device with the lowest 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 (Table 6).

Table 1: Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)

Name	Description
PARTICIPATING	If set to 0, 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 1, 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 will be zero. '1111b' is the highest priority.
GROUP ID [6:0]	The GROUP ID defines the zoning group ID in the range from 0..127. GROUP ID = 0: Group 0 is a special group that is not allowed to communicate with any other group except for group 127. GROUP ID = 127: Group 127 is a special group that is allowed to communicate with all other groups. GROUP ID = 1..126: User defined groups. The communications amongst the user-defined groups are restricted by the zoning permission table.

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

The shaded area of the Zoning Permission Table (Table 2) shows the user-defined permissions among groups 1..126.

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

- If P[X,Y] is set to 1, then Group X has permission to access Group Y;
- If P[X,Y] is set to 0, 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 all zeros.
- P[0..126, 0] shall be set to 0.

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 all ones.

5.3 Supervisors

A supervisor can be an expander or a SAS end device that is an SMP initiator and target. It supports all SMP zoning commands and has SUPERVISING PRIORITY attribute is not equal to 0. 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.

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. A supervising capable zoning expander may also have a SUPERVISING PRIORITY value greater than zero. 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 domain. 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. When the SUPERVISING PRIORITY field is set to 0, the expander is not allowed to participate in the election process for active supervising device selection.

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. 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. The exact topology discovery algorithm is vendor-specific.

The SUPERVISING PRIORITY attribute of a zoning expander determines whether the zoning expander can be a candidate for active supervisor and its priority for election. All zoning expanders and end devices that have their SUPERVISING PRIORITY bits set to a value greater than zero are candidates for the active supervisor election process. If a device has the SUPERVISING PRIORITY attribute set to zero, then it will not take part in the election process and cannot be chosen as the active supervisor.

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 repeats topology 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. If there are no devices available for election to be the active supervisor then the REPORT GENERAL response frame shall return value of zero in the SUPERVISING EXPANDER ADDRESS to indicate the election process has failed.

During normal operation, if a supervisor device status changes, a BROADCAST (CHANGE) is transmitted and the election process begins again and a replacement supervisor will be elected per the election process. Whenever there is a change in the topology each expander shall repeat the Active Supervisor Election process.

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 (10h)								
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									
19	ENCLOSURE LOGICAL IDENTIFIER								
20									
27	ACTIVE SUPERVISOR SAS ADDRESS								
28	(MSB)	CRC						(LSB)	
31									

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 must be supported in all zoning expanders. The user-defined group ID shall range from 1 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.

5.3.4 Supervisor Topology reporting

To discover the full SAS domain, a supervisor shall issue the DISCOVER request with the NO_ZONE_MASKING bit set to 1. (Table 5), called PHYSICAL DISCOVERY. When the DISCOVERY request is made with the NO_ZONE_MASKING bit is set to 0, then only the devices accessible to the initiators group will be reported, called LOGICAL DISCOVERY.

This mechanism permits the any supervisor device to gain full topology information for control, whilst belonging to any group for data traffic purposes.

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	Reserved								
4	Ignored								
7									
8	Reserved							NO_ZONE_MASKING	
9	PHY IDENTIFIER								
10	Ignored								
11	Reserved								
12	(MSB)	CRC						(LSB)	
15									

NO_ZONE_MASKING bit shall be set to 1 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 0, then only the phys permitted access by the zone permission table will be reported, called LOGICAL DISCOVERY

5.3.5 DISCOVER Response

The DISCOVER function is defined in Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9: 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 source that does not have group access privilege to all the phys, then the response frame for the phys that cannot be accessed will respond with all the zone information set to 0, 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	Reserved							
7								
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11	Reserved							
12	Reserved	ATTACHED DEVICE TYPE			Reserved			
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	(MSB)	CRC						(LSB)
52	(MSB)	CRC						(LSB)
55	(MSB)	CRC						(LSB)

NO_ZONE_MASK field is set to 1 if the function is responding to a DSICOVERY request that had the NO_ZONE_MASK bit set to 1. This is a confirmation that the response is a domain wide report. If the NO_ZONE_MASKING bit is set to 0, the response is subject to the group id restrictions defined in the zoning permission table (Table 2).

The ZONE VIOLATION field is set to 1 if any ZONE violation has occurred causing the specified phy to send an OPEN_REJECT(ZONE VIOLATION) message. 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.

The PARTICIPATING bit reports whether the specified phy is inside the zoned boundary and attached to another device inside the zoned boundary. If set to 1, this indicates that the device attached to this phy is participating in zoning operations and that the ACCESS_ZONE MANAGEMENT and SOURCE GROUP ID fields of the OPEN address frames received on this phy are not replaced with those assigned to this phy. If set to 0, the attached device is outside of the zoned service delivery subsystem.

If the ACCESS_ZONE MANAGEMENT bit is 0 then all zoning fields shall be set to zero.

The SUPERVISING PRIORITY bits indicate the priority for election of active supervisor for the device attached to this phy. This value will be zero if the device is not a supervisor capable device. The GROUP ID fields reports the source group ID assignment of the specified phy. If the attached device is a zoning-aware expander, then the Supervising Priority is returned by REPORT GENERAL and the PHY Supervising Priority is ignored for election purposes.

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 (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 Handling

The OPEN address frame used in a zoned SAS environment includes ACCESS ZONE MANAGEMENT bit and SOURCE GROUP ID (Table 8). These fields are only valid when the OPEN address frames are passed between devices inside the zoning service delivery subsystem on participating expander phys.

When an expander receives an OPEN address frame on a phy with the participating bit set to 0, it sets the SOURCE GROUP ID according to the value of the GROUP ID of the expander phy and sets the ACCESS ZONE MANAGEMENT bit according to the SUPERVISING PRIORITY attribute of the phy. The AZM bit is set to zero if the supervising priority is zero, otherwise it is set to one. When the non-participating expander phy transmits or forwards an OPEN address frame, it sets the SOURCE GROUP ID to 0.

When a participating expander phy transmits or forwards an OPEN address frame, 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 SMP Open Address frame is received, the expander shall check whether the SAS SOURCE ADDRESS matches the active supervisor address. If the source is confirmed to be the active supervisor, then the expander SMP target supports the complete SMP command set including zone management SMP functions. If the source is not the active supervisor then the expander SMP target shall only support non-configuration zoning commands and shall respond with UNKNOWN SMP FUNCTION if a zoning configuration command is received. If the AZM bit is not set then all zoning SMP Functions are responded to with UNKNOWN SMP FUNCTION and the zoning-specific fields of the SMP DISCOVER response are set to 0.

The SOURCE GROUP ID field in the OPEN address frame allows the expander's SMP target port to respond to SMP commands according to the source group. An SMP DISCOVER request, with the NO_ZONE_MASKING bit set to 0 (Table 5), shall only report the phys that have group affiliation with the SOURCE GROUP ID of the requestor. The remaining phys are reported as vacant. The SMP REPORT ROUTE INFO (reference), SMP CONFIGURE ROUTE INFO (reference) and SMP REPORT ZONE ROUTE TABLE (reference) commands can only access the routing table entries that the SOURCE GROUP ID is allowed to see, as defined by the zone permission table.

5.5.1 OPEN Address Frame

Table 8 The OPEN Address Frame Format defines the OPEN address frame format used for connection requests.

Table 8: 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 which source group the OPEN address frame is coming from.

The ADDRESS FRAME TYPE field specifies the type of address frame. This field is defined in Table 9 This field determines the definition of the frame type dependent bytes.

Table 9: 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

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 (Table 4). During the discover process the minimum zones supported value determines the zones for a 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 0, and the device will be excluded from the zoned service delivery subsystem and treated as a non-zone aware device.

The SMP DISCOVER response indicates ZONE VIOLATION,-PARTICIPATING, SUPERVISING PRIORITY, and GROUP ID information for each phy (Table 6). When the SMP DISCOVER command is executed from a source group (as indicated by the SOURCE GROUP ID in the OPEN address frame that set up the SMP connection), the zoning expander shall report the accurate information for the phys that the source group is allowed to access according to the zone permission table. The phys that are inaccessible from the source group shall be reported as VACANT.

5.6.1 CONFIGURE PHY ZONE Function

The SMP CONFIGURE PHY ZONE request and response allow a supervisor to change the PHY ZONE configuration of a specific phy. The zoning expander shall only execute this command if the OPEN address frame has its ACCESS ZONE MANAGEMENT bit set to 1 and the source SAS address from 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 executes the CONFIGURE PHY ZONE request shall send out a BROADCAST(CHANGE) message to the affected groups.

The CONFIGURE PHY ZONE function sets the expander phy zone group membership, whether the device attached to this phy will participate in active supervisor election and if it is outside the zoned service delivery subsystem or a zoning participant. This command sets the zone membership assignment for one or multiple expander phys with contiguous phy indexes. The active supervisor shall send out a BROADCAST(CHANGE) message with a Source group ID of 127 to all zones after the CONFIGURE PHY ZONE function is executed.

This SMP function shall be supported by a SMP target port in an expander device if the NUMBER OF ZONES supported field is greater than zero in the REPORT GENERAL function. The SMP target will only execute the CONFIGURE PHY ZONE function if the OPEN address frame that set up the SMP connect has its ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN address frame has its ACCESS ZONE MANAGEMENT bit set to 0 the SMP target shall generate a SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION. This command shall return FUNCTION FAILED if the Source SAS Address from the Open SMP Request Frame is not equal to the Active Supervisor SAS Address. If the Active Supervisor's SMP Target Port Receives this command it shall respond with FUNCTION FAILED.

Table 10 defines the CONFIGURE PHY ZONE request format.

Table 10: 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	Reserved	START PHY INDEX						
7	NUMBER OF ZONE PHY ENTRIES							
	PHY ZONE configuration entry list							
8	First PHY ZONE configuration entry descriptor							
9								
...	...							
n-5	Last PHY ZONE configuration entry descriptor							
n-4								
n-3	(MSB)	CRC						
n								(LSB)

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

The FUNCTION field shall be set to 93h.

The CRC field is defined in Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.1.

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 0 to 255.

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

Table 11 defines the PHY ZONE configuration entry descriptor.

Table 11: 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 ID to be assigned to the specified phy.

The SUPERVISING PRIORITY field specifies whether the specified phy is a candidate for active supervisor and if so what priority it has. The PARTICIPATING field indicates whether the device

attached to this specified phy is participating or not participating in zone management. It will be set to 1 if the specified phy is connected to another phy within the zoned service delivery subsystem.

Table 12 defines the CONFIGURE PHY ZONE response format.

Table 12: 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							
7							(LSB)		

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 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.2.

The CRC field is defined in Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.1.

5.6.2 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 (Table 10): single set and batch set mode.

In single set operation, the value for a single permission table between one source group and a target group is set to the value provided by the NEW VALUE field of the CONFIGURE ZONE PERMISSION request frame (Table 10).

In batch set operation, the CONFIGURE ZONE PERMISSION sets one or multiple contiguous entries of the permission table.

The single set operation causes the same NEW VALUE to be set to both the [Source group ID, Target group ID] bits as well as the [Target group ID, Source group ID] bits 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.

This SMP function shall be supported by a SMP target port in an expander device if the NUMBER OF ZONES SUPPORTED field is greater than zero in the REPORT GENERAL function. The SMP target shall only execute the CONFIGURE ZONE PERMISSION function if the OPEN address frame that set up the SMP connection has its ACCESS ZONE MANAGEMENT bit set to 1. If the ACCESS ZONE MANAGEMENT bit is set to 0, the SMP target shall generate an SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION. If the Source SAS Address from the OPEN Request Frame does not match the Active Supervisor SAS Address then the response shall be FUNCTION FAILED.

Table 10 defines the CONFIGURE ZONE PERMISSION request format.

Table 13: 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						PROPAGATE UPDATE	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 field when set to 1 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 set to 0 this function will set 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 between the Source group ID and the Target group ID. For batch set mode, this field shall be set to 0. Note that this value is set for both the permission table entry [Source group ID, Target group ID] bits, and the permission table entry [Target group ID bit, Source group ID] bits.

The SOURCE GROUP ID field is only used in single set mode. It provides the source group ID to be modified by the single set operation. For batch set mode, this field shall be set to 0.

The TARGET GROUP ID field is only used in single set mode. It provides the target group ID to be modified by the single set operation. For batch set mode, this field shall be set to 0.

The PROPAGATE UPDATE bit is set to one to indicate the UPDATE ZONE PERMISSION command is sent to the active supervisor. The PROPAGATE UPDATE bit is set to zero to indicate the UPDATE ZONE PERMISSION command is sent from the active supervisor to other zoning expanders of the zoning domain.

If a non-supervising expander receive 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 non-supervising 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 SMP 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 FUNCITON 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, this field shall be set to 0.

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. For single set mode, this field shall be set to 0.

The PHY ZONE entry descriptor list contains zero 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 Section 10.4.3.1.

Table 14 defines the ZONE PERMISSION entry descriptor.

Table 14: ZONE Permission Entry Descriptor

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

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 0.

Table 15 defines the CONFIGURE ZONE PERMISSION response format.

Table 15: 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)	CRC							
7								(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 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.2.

The CRC field is defined in Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.1.

5.6.3 REPORT ZONE PERMISSION Function

The SMP CONFIGURE ZONE PERMISSION request and response allow the active supervisor to update the zone permission table. The Active supervisor shall send the SMP CONFIGURE ZONE PERMISSION commands 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 SMP REPORT ZONE ROUTE TABLE response reports the zone route table details, which is an extension of the routing table defined by SAS 1.1. 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.

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

This SMP function shall be supported by a SMP target ports in an expander device if the NUMBER OF ZONES SUPPORTED field is greater than zero in the REPORT GENERAL function.

The SMP target will only respond to the CONFIGURE ZONE PERMISSION request if the OPEN address frame that set up the SMP request has its ACCESS ZONE MANAGEMENT bit set to 1. If the ACCESS ZONE MANAGEMENT bit set to 0, the SMP target will generate the FUNCTION RESULT set to UNKNOWN SMP FUNCTION response.

Table 16 defines the REPORT ZONE PERMISSION request format.

Table 16: 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							
11								(LSB)	

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 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.1.

Table 17 defines the REPORT ZONE PERMISSION response format.

Table 17: 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						(LSB)
n								

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 will 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 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 executed. 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

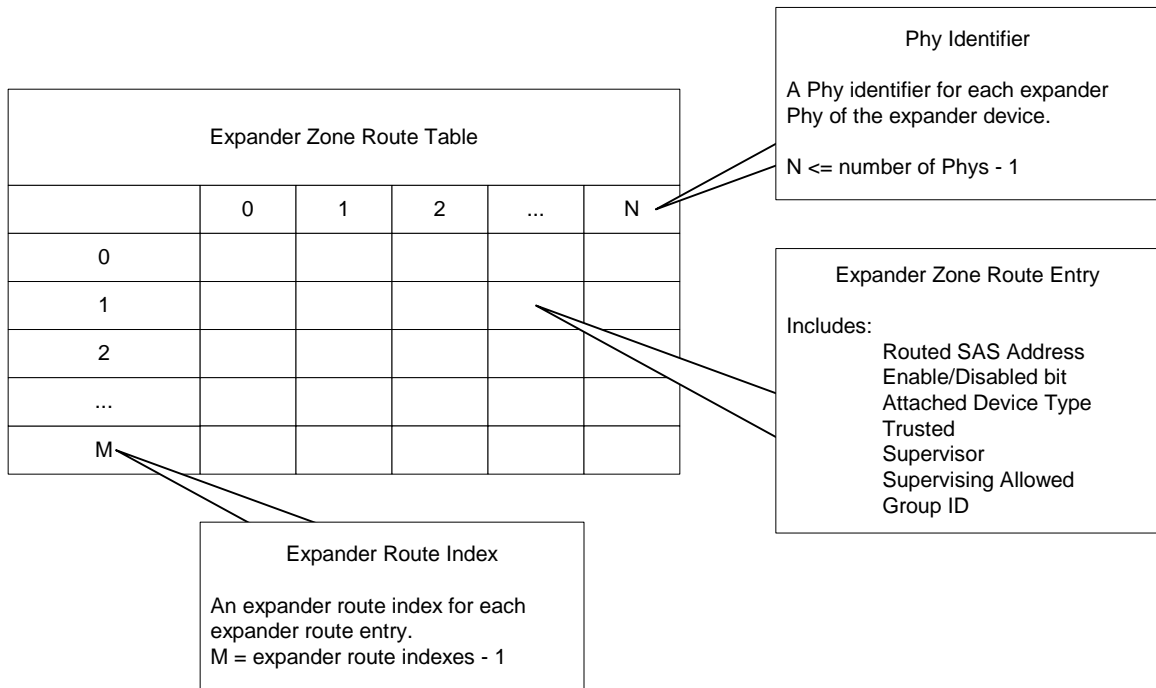
The expander zone route table is a structure that provides an association between routed SAS addresses to expander phy identifiers using the phy zone configuration attributes PARTICIPATING, SUPERVISOR, GROUP ID, SUPERVISING PRIORITY, and the attached device type. Each association is represented by one 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 (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 (Table 18). All zoning expanders are responsible for configuring the zone route table automatically through the topology discovery process described in the Topology Discovery Section 5.3.2

Figure 2: Expander Zone Route Table Example 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 0 and the EXPANDER ROUTE INDEXES field is set to a non-zero value in the REPORT GENERAL response function then the expander device shall provide the SMP REPORT ZONE ROUTE TABLE Response. The SMP target shall only execute the REPORT ZONE ROUTE TABLE function if the OPEN address frame that established the SMP connection had its ACCESS ZONE MANAGEMENT bit set to 1. If the ACCESS ZONE MANAGEMENT bit is set to 0, the SMP target will 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 which is not accessible to the source Group ID from the Open Request Frame then this function shall respond with PHY VACANT.

Table 18 defines the REPORT ZONE ROUTE TABLE request format.

Table 18: 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 which the expander route entry is being read (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 19 defines the REPORT ZONE ROUTE TABLE response format.

Table 19: 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 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 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 for which the expander route entry is being read (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 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9 Section 10.4.3.1.

The Zone Route Table Entry descriptor format is defined in Table 20

Table 20: 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 0, then the ECM shall use the expander route entry to route connection requests. If the DISABLE EXPANDER ROUTE ENTRY bit is set to 1, 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 Table 85 Reference 1 : SAS1r09 - Serial Attached SCSI 1.1 Revision 9

The SUPERVISING PRIORITY is a 4-bit field that is set to prioritize which expander device in a system should be elected 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, all zoning expanders shall report the PHY CHANGE COUNT (Reference DISCOVER response Table 6) and EXPANDER CHANGE COUNT (Reference REPORT GENERAL response Table 4) according to the group assignment of the source ID that issued the request. If the entire SAS topology information is required then the NO_ZONE_MASKING bit shall be set to 1 in the DISCOVER request (Table 5) to indicate full topology data is required. The expander zone route table (Figure 2) maintains the number of BROADCAST (CHANGE) messages originated by each phy.

When a zoning expander receives a DISCOVER request and the NO_ZONE_MASKING bit is set to 1 then the expander shall report the PHY CHANGE COUNT. If the NO_ZONE_MASKING bit is set to 0 then the PHY CHANGE COUNT will be reported for phys that have access permission granted to the SOURCE GROUP ID field of the OPEN address frame. Any phy that is not accessible to the SOURCE GROUP ID shall be reported as vacant.

When a zoning expander receives a REPORT GENERAL request, the expander shall report the sum of the PHY CHANGE COUNT values that are accessible defined by the SOURCE GROUP ID. If no phys have access permission then the value of the EXPANDER CHANGE COUNT will be zero.

6.2 *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 (Table 21).

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 shall transmit BROADCAST address frames to devices within the zoned area.

When a zoned expander detects a phy event that would normally spawn a BROADCAST (CHANGE) 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 for the expander Broadcast Primitive Processor.

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.

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 preserved and be presented to the expander Broadcast Primitive Processor as a ZONED BROADCAST request.

When the expander Broadcast Primitive Processor receives a ZONED BROADCAST request from a phy, it forwards the broadcast indication to all the expanders phys that are accessible by

the source group of the broadcast request except the phys on the source port. The rule for wide links that applies to BROADCAST messages applies to this broadcast method. When transmitting a ZONED BROADCAST event on a participating expander phy, the expander shall transmit a ZONED BROADCAST address frame and set the outgoing SOURCE GROUP ID field to the SOURCE GROUP ID field received (or the phy group ID of the phy causing the broadcast).

If a BROADCAST address frame is received on a phy whose participation bit is set to 0, the zoning expander will replace the zone information with that of the receiving phy.

6.2.1 ZONED BROADCAST Address Frame

Table 21 defines the address frame format for zoned broadcasts.

Table 21: 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 22 defines the valid values of the BROADCAST TYPE field.

Table 22: 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 are accessible from the group ID specified by the SOURCE GROUP ID field defined by the expander zoning permission table. Reference (Table 2)

6.3 Zone Checks

During SAS zoning operation, if an OPEN address frame is received on a phy whose participation bit is set to 0, 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 which approach 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.

Editors Note: These reference tables are to assist update tracking and do not belong in the final document.

7 Index

7.1 Table of contents

1	Revision History	1
2	Related Documents	1
3	Overview	1
4	Definitions	3
5	SAS Zoning	4
5.1	Zoned service delivery subsystem	4
5.2	Zoning Configuration	4
5.2.1	Per Expander Phy Configuration	4
5.2.2	Zoning permission table	5
5.3	Supervisors	6
5.3.1	Supervisor Election Process	6
5.3.2	Topology Discovery	6
5.3.3	REPORT GENERAL Response	7
5.3.4	Supervisor Topology reporting	7
5.3.5	DISCOVER Response	8
5.4	Device Group Reassignment	10
5.5	OPEN Address Frame Handling	11
5.5.1	OPEN Address Frame	11
5.6	SMP Functions for zone management	13
5.6.1	CONFIGURE PHY ZONE Function	13
5.6.2	CONFIGURE ZONE PERMISSION Function	15
5.6.3	REPORT ZONE PERMISSION Function	18
5.7	ZONE Management Updates	21
5.7.1	Expander Zone Route Table	21
5.7.2	REPORT ZONE ROUTE TABLE Function	22
6	Zone Operation	26
6.1	PHY and EXPANDER CHANGE COUNT reporting	26
6.2	Broadcasts	26
6.2.1	ZONED BROADCAST Address Frame	27
6.3	Zone Checks	28
7	Index	30
7.1	Table of contents	30
7.2	Table of Figures	30
7.3	Table of Tables	30

7.2 Table of Figures

Figure 1	: Zoning Service Delivery Subsystem	2
Figure 2	: Expander Zone Route Table Example	22

7.3 Table of Tables

Table 1	: Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)	4
Table 2	: Zoning Permission Table	5
Table 3	: Zoning Expander Election Priority	6
Table 4	: REPORT GENERAL Response	7

Table 5: DISCOVER Request	8
Table 6: DISCOVER Response.....	9
Table 7: PHY ZONE Reassignment after a phy Reset Event	10
Table 8: OPEN Address Frame Format	12
Table 9: ADDRESS FRAME TYPE Field	12
Table 10: CONFIGURE PHY ZONE Request	14
Table 11: PHY ZONE Configuration Entry Descriptor.....	14
Table 12: CONFIGURE PHY ZONE Response	15
Table 13: CONFIGURE ZONE PERMISSION Request.....	16
Table 14: ZONE Permission Entry Descriptor	17
Table 15: CONFIGURE ZONE PERMISSION Response	18
Table 16: REPORT ZONE PERMISSION Request	19
Table 17: REPORT ZONE PERMISSION Response.....	20
Table 18: REPORT ZONE ROUTE TABLE Request.....	23
Table 19: REPORT ZONE ROUTE TABLE Response	24
Table 20: ZONE ROUTE TABLE Entry Descriptor.....	25
Table 21: ZONED BROADCAST Address Frame Format	27
Table 22: BROADCAST TYPE Field	27