

To: T10 Technical Committee
From: Heng Liao, PMC-Sierra (heng.liao@pmc-sierra.com),
Steve Gorshe, PMC Sierra (steve_gorshe@pmc.com) and
Tom Grieff, HP (thomas.grieff@hp.com)
Date: 19 April 2005
Subject: T10/05-144r0 SAS-2 zoning

Revision History

Revision 0 (19 April 2005) first revision

Related Documents

sas1r09 - Serial Attached SCSI 1.1 revision 9

Overview

SAS is starting to gain interest for use as a small storage area network (SAN). With this, there is a need to add a feature to segregate traffic among devices like Fibre Channel provides with zones or Ethernet TCP/IP provides with virtual LANs. The model requires no changes to end devices; initiators continue to perform normal SAS discovery, and initiators and targets perform OPENS as normal. The difference is they do not see the entire SAS domain; they see only the portions of the domain that they are allowed to see based on the permission table.

The model supports up to 128 **groups**. The group assignment is strictly based on expander ports of the zoning fabric. Each fabric port can only belong to one group. A permission table defines whether communication is allowed between two groups.

The zoning fabric boundary is defined by the system designer by marking phys in expanders as trusted or not trusted. Devices inside the boundary are considered trusted, devices outside the boundary are not trusted.

SMP functions provide the means to control and configure the expander phy Group Assignment, expander phy Trusted bit, expander phy Supervisor bit and expander permission table. If a device is attached to an expander phy marked as supervisor, the device is considered to have supervisor privilege to the zoning fabric. Multiple phys in the zoning fabric can be marked as supervisors allowing redundant supervisors to exist in one fabric.

An supervising expander is responsible for propagating zone permission table changes to all zoning expanders in the fabric. Any zoning expander in the fabric can act as the supervising expander if it receives a CONFIGURE PERMISSION TABLE command with SUPERVISE THIS bit from a supervisor. The purpose of having a supervising expander is to offload the supervisor and the associated traffic from the management port for permission table propagation.

The zoning configuration SMP functions are only allowed if the SMP function is originated from a supervisor device.

Suggested Changes

4.1 Zoning model

4.1.1 Zone Model Overview

A set of expander devices in a SAS domain may support zoning to restrict communication between SAS devices in that SAS domain. It is up to the system designer to decide if the SAS domain is going to be zoned or not. If so, expanders have to be chosen that implement special zoning features, and certain phys in those expanders have to be marked to indicate the boundaries of the secure part of the domain. Legacy expanders can be cascaded outside the boundary of the zoned fabric.

All devices attached to the legacy expanders will inherit the group assignment of the phy at the boundary of the zoned fabric.

The model supports up to 128 zones. It is vendor-specific how many zones each expander device supports; it is up to the system designer to select expander devices that meet the minimum requirements of the system. The SMP REPORT GENERAL command reports the number of zones supported by a zoning expander.

4.1.2 Zoning Configuration

The Expander zoning configuration consists of two parts: the per expander phy configuration, and per expander zoning permission table.

Every expander phy is associated with the following per-phy configuration parameters:

Table 1. Per-expander phy zoning configuration

Name	Description
Trusted	<p>If set to 0, this phy is on the boundary of the zoning fabric. All message (primitives and frames) that come across this phy shall be mapped to be backwards-compatible to SAS standard without zoning features, except for the new SMP commands defined by the zoning extension.</p> <p>If set to 1, this phy is inside the fabric boundary. The new primitives and frame formats that are defined by the zoning extension are allowed to pass through this phy.</p>
Group ID[6:0]	<p>The GID defines the zoning Group ID in the range from 0..127.</p> <p>GID=0: Group 0 is a special group that is not allowed to communicate with any other group except for group 127. Note that a device belonging to group 0 can still discover all the expanders and communicate with the SMP virtual target in the expanders (i.e. SMP virtual target within the zoning expanders are considered to have GID=127).</p> <p>GID=127: Group 127 is a special group that is allowed to communicate with all other groups. All trusted phys shall be automatically assigned to have GID =127 by the zoning expanders.</p> <p>GID=1..126: User defined groups. The communications amongst the user defined groups are restricted by the zoning permission table.</p>
Supervisor	<p>If Set to 1, the device attached to this phy is allowed to originate SMP commands to set up and change zoning configuration.</p> <p>If set to 0, the device attached to this phy is not allowed to originate SMP commands to change the zoning information.</p>

As shown in table 2, the zoning permission table has up to 128 entries with each entry corresponding to one group. Each entry is also referred to as a row of the permission table.

Table 2. Zoning permission table

	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

Note that the grey area of table 2 is user-defined permissions among group 1..126.

Each entry contains a bitmask of up to 128 bits with each bit corresponding to the access permission of a target group. Each bit within the bitmask is also referred to as a column of the permission table.

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

P[X,Y] =1: means group X has permission to access Group Y

P[X,Y] =0: means group X has no permission to access group Y

Note that 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 equal to P[Y,X].

Note that group 0 is not allowed to access any other group except for 127. P[0, 0..126] are always set to all zeros. P[0, 127] is always set to 1. P[0..126, 0] are always zero.

Note that group 127 is allowed to access all other groups. Therefore, P[0..127, 127] is always set to all 1s, and P[127, 0..127] is also set to all 1s.

4.1.3 OPEN address frame handling

The OPEN address frame used in a zoned SAS environment includes a new ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID field (in the COMPATIBLE FEATURES area). The ACCESS ZONE MANAGEMENT bit and SOURCE GROUP ID field in these OPEN address frames are only valid among devices inside the zoning fabric on trusted expander phys.

When an untrusted expander phy receives an OPEN frame, it sets the ACCESS ZONE MANAGEMENT bit according to the value of SUPERVISOR bit of the expander phy, and sets the SOURCE GROUP ID according to the value of the GROUP ID of the expander phy. When an untrusted expander phy transmits or forwards an OPEN, it sets ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID to zero.

When a trusted expander phy transmits or forwards an OPEN, the value of the ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID are preserved and transmitted. This mechanism allows the use of the new OPEN frame format inside zoning fabric across trusted phys, and ensures legacy OPEN address frame format is used outside the zoned fabric boundary. This preserves backwards compatibility.

The ACCESS ZONE MANAGEMENT bit in the OPEN address frame determines whether an expander's SMP target shall execute the SMP zoning management SMP functions during the connection. If the SMP connection is opened with an OPEN address frame with ACCESS ZONE MANAGEMENT bit set high, the expander SMP target supports the complete SMP command set including the Zone management SMP functions. Otherwise, the expander SMP target shall not support the Zone management SMP functions and generate SMP response frame with FUNCTION RESULT equal to UNKNOWN SMP FUNCTION should a zone management SMP function be requested.

The SOURCE GROUP ID field in the OPEN frame allows the expander SMP target port to respond to SMP commands according to the source group. For instance, the SMP PHY_CONTROL command from the specified group can only access phys that the source group has permission to see according to the permission table. Similarly, SMP DISCOVER command shall only report the phys the source group can see, and the remaining phys are reported as disabled; SMP REPORT ROUTE INFO and SMP CONFIGURE ROUTE INFO commands can only access the routing table entries that the source group is allowed to see according to the permission table.

4.1.4 SMP functions

The SMP REPORT ZONE ROUTE TABLE function is defined to configure expander phy zone membership; read and write the zone table in an expander. The SMP target returns an error

(UNKNOWN SMP FUNCTION) if an attempt is made to access these functions in a connection that does not have the ACCESS ZONE MANAGEMENT bit set to 1.

The SMP CONFIGURE PHY ZONE function is added to configure expander phy zone membership.

The SMP DISCOVER command is extended to report zone configuration for specific expander phys when they are executed in a connection with ACCESS ZONE MANAGEMENT bit set to 1.

The SMP CONFIGURE ZONE PERMISSION function sets the permission table for the expander.

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

The SMP REPORT GENERAL function is extended to support a ZONING SUPPORTED Bit that indicates whether the expander supports zoning, the ZONE PERMISSION CHANGE ENTRY COUNT field and ZONE PERMISSION CHANGE ENTRY INDEX field that provides information about the permission table change status.

The SMP CONFIGURE PHY ZONE function sets the Group ID (GID), the SUPERVISOR Bit and the TRUSTED Bit for an expander phy.

The SMP CONFIGURE ZONE PERMISSION function sets the permission table for a zoning expander.

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

The SMP DISCOVER function returns the physical link configuration information for the specified phy including the Group ID (GID), the SUPERVISOR bit and the TRUSTED bit for the specified expander phy.

The SMP REPORT ZONE ROUTE TABLE function returns the route table contents. Since SMP frames are limited to 1024 bytes, it is not possible to access the entire zone/route table within one SMP request or response frame. The function can report multiple contiguous table entries at a time such that the full table can be retrieved in multiple steps (but does not require a separate SMP connection function to retrieve each individual entry, as does the REPORT/CONFIGURE ROUTE INFORMATION).

Note that the zone supervisor device is expected to initiate SMP CONFIGURE PHY ZONE and SMP CONFIGURE ZONE PERMISSION commands to set the per expander phy zone membership information and the zone permission table.

Zoning expanders shall be "self-configuring expanders". The extended discovery process in the zoning expanders are responsible for establishing SAS WWN to expander phy mapping relationship, the SAS WWN to Group ID relationship, the Supervisor and the Trusted bits.

The Supervisor or a Supervising expander shall be responsible for sending consistent zone permission table updates via SMP CONFIGURE ZONE PERMISSION command to all zoning expanders in the fabric. After the zone permission table updates to all zoning expanders are completed, the Supervisor or Supervising Expander is responsible for sending BROADCAST(CHANGE) message to force topology rediscovery in the expanders and the hosts. Note that the supervisor function may be implemented in any device that has SMP initiator capability including host end devices and expander devices. The expander self discovery process is not responsible for propagation of the zone permission table among each other.

The zoning expanders use SMP REPORT ZONE ROUTE TABLE and SMP DISCOVER commands to exchange topology and establish the mapping from SAS address to group ID assignment.

The zoning expanders still support the SMP REPORT ROUTE INFO command for host topology discovery. The zoning expander SMP virtual target uses the GID field in the SMP OPEN frame to determine which source port is doing topology discovery and presents only the route entries the source group is permitted to see based on the zoning permission table, all remaining entries are reported as disabled.

Every time a zoned expander completes a PHY ZONE configuration update with the UPDATE COMPLETE flag set high, the zoned expander is responsible for sending out a BROADCAST CHANGE message to all ports to trigger topology rediscovery.

In other words, the zoning supervisor is responsible for configuring expander phy to group assignment and for setting the zoning permission table for zoning expanders in the fabric. For changing the PHY ZONE configuration, the zoning supervisor sends CONFIGURE PHY ZONE command to the specific expander with the phy to be configured. The self configuring zoning expander extended topology discovery process is responsible for establishing the SAS address to expander port and group ID mapping and distributing the zone/route table information among expanders without the direct involvement of the zoning supervisor.

For redundancy purpose, multiple zoning supervisors are allowed to co-exist and be attached to any zoning expanders in the topology. The zoning supervisors shall cooperate and ensure the phy zone configuration and the zone permission table are maintained in a consistent manner.

4.1.5 Broadcasts

Broadcasts must be controlled to avoid crossing zones. The ZONED BROADCAST address frame (see 7.8.x) is used within zoning fabric to propagate BROADCAST primitives without crossing zone boundaries:

When a zoned expander detects a phy event or received a broadcast primitive that would normally spawn a BROADCAST (CHANGE), it instead sends a ZONED BROADCAST address frame with the SGID field set to the Group ID (GID) of the phy which caused the event, to all phys that are accessible from the source Group.

When an expander receives a ZONED BROADCAST address frame on a phy, it forwards it only to the phys that are accessible from the source Group specified by the SGID field received in the address frame. The rule for wide links still apply to this broadcast method that apply to the legacy BROADCAST. This if more than one phy are joined in a wide expander port, then the address frame should only be sent one time, on any or all available phy.

When transmitting a ZONED BROADCAST event on a trusted expander phy, the expander transmits a ZONED BROADCAST address frame and sets the outgoing SGID field to the SGID field received (or the phy GID of the phy causing the broadcast).

A BROADCAST primitive received on a trusted phy shall be treated the same way as a BROADCAST address frame with SGID of 127 (unrestricted broadcast).

A BROADCAST primitive received on a untrusted phy shall be treated the same way as a BROADCAST address frame with the SGID assigned to the Group ID of the receiving phy.

When transmitting a ZONED BROADCAST event on a untrusted phy, the expander transmits a BROADCAST primitive that represents the type of broadcast event represented by the ZONED BROADCAST event, but the SGID information is discarded.

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

4.1.6 Source SAS address checking

For phys with SOURCE CHECK bit set to 1 or phys attached to end devices (as indicated by the IDENTIFY address frame), the expander device shall check that every OPEN address frame received by the phy has the same source SAS address as originally indicated by the IDENTIFY address frame. If it receives an OPEN address frame with a different source SAS address, the expander shall reject the connection request with OPEN_REJECT (ZONING VIOLATION) and shall set a TBD bit retrievable via the DISCOVER function indicating that the phy has received an invalid OPEN address frame.

A primitive type OPEN_REJECT(ZONING VIOLATION) is defined to replace OPEN_REJECT(RESERVED ABANDON 0) for backwards compatibility.

Editor's Note 2: maybe define OPEN_REJECT (ZONING VIOLATION) to replace OPEN_REJECT (RESERVED ABANDON 0) and use for this situation (and others)

Editor's Note 3: The security violation bit could be a single bit or a counter. It must be clearable by the management agent (unless it's a big counter that just rolls over). Provide an optional mechanism to record the last (few) attempted rogue address(es).

Editor's Note 4: Define optional BROADCAST to report zoning violations?

4.1.7 Zone checks

When an expander phy receives an OPEN, the expander normally compares the SAS address to the SAS addresses that are directly attached to the expander and then looks up the destination SAS address in the routing table.

With zoning, if the OPEN frame is received from an untrusted phy, the expander device retrieves the zone membership (GROUP ID) of the phy that receives the OPEN address frame and sets the SOURCE GROUP ID field to this value. When the OPEN frame is received on a trusted phy, the SOURCE GROUP ID field in the OPEN frame carries the source group information (SGID).

The SAS direct and table routing mechanism is extended to map the destination address of the OPEN frame to the destination port, as well as the Destination Group ID (DGID). If the destination device is directly attached to the expander device, the direct routing lookup should find a match. In this case, the group ID assignment of the matching destination phy is used as the destination group ID (DGID). If the direct lookup fails to find a match, the expander performs table routing lookup. The table routing lookup maps the destination address to a destination group ID (DGID) and a destination port.

If direct lookup and zone table lookup maps the destination address to a Destination Group ID (DGID), the expander checks the zoning permission between the Source Group ID (SGID) and the destination group (DGID). If the zoning permission table does not have a permission between SGID and DGID, the OPEN is rejected with OPEN_REJECT(ZONING VIOLATION). If the zoning permission table allows routing between SGID and DGID, the expander shall proceed with the normal ECM arbitration procedure.

If neither table routing nor direct routing finds a match for the destination and the expander supports subtractive routing, the OPEN is routed according to subtractive routing rules without zoning permission check. If the subtractive port is used to connect to an adjacent expander, it is a requirement that the subtractive port must be trusted and set to group 127. This allows the OPEN to be forwarded to the next expander which has the topology knowledge to do table or direct routing and the permission check.

Editor's Note 5: The security violation bit could be a single bit or a counter per phy. It must be clearable by the management agent (unless it's a big counter that just rolls over). Provide an optional mechanism to record the last (few) attempted rogue address(es).

Editor's Note 6: Define optional BROADCAST to report zoning violations?

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

This should be – Although it is compatible to perform route to destination phy before the zoning permission decision is made, it is preferred that zoning expanders perform the zoning permissions check at the boundary of the zone. This gives the additional benefit of stopping an illegal route at the ingress and avoid tying up intermediate links in arbitration before being backed off. The rest of this is implementation specific and vendor value add.

Note that All SMP target in the zoning expanders are considered to be part of group 127. Therefore, all OPENS to the SMP targets are permitted.

If an end device is attached to an untrusted phy, it is by default assigned to group 0, which is not permitted to OPEN any connections except for the target phys assigned to group 127 until it has been added to a zone; all other connection requests from these phys result in OPEN_REJECT (ZONING VIOLATION). If an end device is attached to a trusted phy, that expander shall set the phy GID to 127 for that phy and allow that phy to OPEN the expander device itself or any other device in the fabric.

7.8 Address frames

7.8.1 Address frames overview

...

The ADDRESS FRAME TYPE field specifies the type of address frame and is defined in table 78. This field determines the definition of the frame type dependent bytes.

Table 3 — ADDRESS FRAME TYPE field

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

...

7.8.3 OPEN address frame

Table 81 defines the OPEN address frame format used for connection requests.

Table 81 — 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						
3	(LSB)							
4	DESTINATION SAS ADDRESS							
11	DESTINATION SAS ADDRESS							
12	SOURCE SAS ADDRESS							
19	SOURCE SAS ADDRESS							
20	ACCESS ZONE MANAGEMENT	SOURCE GROUP ID (SGID)						
21	PATHWAY BLOCK COUNT							
22	(MSB)	ARBITRATION WAIT TIME						
23	(LSB)							
24	MORE COMPATIBLE FEATURES							
27	MORE COMPATIBLE FEATURES							
28	(MSB)	CRC						
31	(LSB)							

...

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

The SOURCE GROUP ID field defines which source group the OPEN is coming from.

...

7.8.x ZONED BROADCAST address frame

Table 4 defines the ZONED BROADCAST address frame.

Table 4 — 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						
31	(LSB)							

....

....

Table 5 defines the BROADCAST TYPE field.

Table 5 — 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 shall forward the broadcast event only to other phys that is accessible from SOURCE GROUP ID according to the expander zoning permission table.

10.4.3 SMP functions

10.4.3.1 SMP function request frame format

...

10.4.3.2 SMP function response frame format

10.4.3.3 REPORT GENERAL function

...

Table 167 — 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	Reserved	NUMBER OF ZONES						
9	NUMBER OF PHYS							
10	Reserved				ZONING SUPPORTED	CONFIGURING	CONFIGURABLE ROUTE TABLE	
11	Reserved	ZONE PERMISSION CHANGE ENTRY COUNT						
12	Reserved	ZONE PERMISSION CHANGE ENTRY INDEX						
13								
27	Reserved							
28	(MSB)	CRC						(LSB)
31								

The NUMBER OF ZONES field indicates the number of zones supported when the ZONING SUPPORTED bit is 1. When the ZONING SUPPORTED bit is 0, this field should be set to 0. Note that group 0 and group 127 must be supported in all zoning expanders. The remaining zone indexes should range from 1 to (NUMBER OF ZONES –2).

The ZONING SUPPORTED bit indicates whether the expander device supports the zoning feature.

The ZONING PERMISSION CHANGE ENTRY COUNT field indicates the number of contiguous zone permission table entries that have been changed.

The ZONING PERMISSION CHANGE ENTRY INDEX field indicates the first entry of the zone permission table that has been changed.

...

10.4.3.5 DISCOVER function

The DISCOVER function returns 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 SMP target shall reply to the DISCOVER command according to the SOURCE GROUP ID fielding the OPEN frame for the SMP connection. If the specified phy is accessible from the SOURCE GROUP ID according to expander permission table, the SMP target shall provide accurate information about the specified expander phy. Otherwise, the SMP target shall generate a response frame that indicates the specified phy is disabled with zone information set to all zero.

....

Table 171 defines the response format.

Table 171 — 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	Reserved		ATTACHED SSP TARGET
16	SAS ADDRESS							
23								
24	ATTACHED SAS ADDRESS							
31								
32	ATTACHED PHY IDENTIFIER							
33	Reserved							
39								
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMAED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	TRUSTED	ZONE VIOLATION	SOURCE CHECK	Reserved		ROUTING ATTRIBUTE		
45	SUPERVISOR	GROUP ID						
46	Reserved							
49								
50	Vendor Specific							
51								
52	(MSB)	CRC						(LSB)
55								

...

The ZONE VIOLATION field is set to 1 if any ZONE violation has occurred causing the specified phy to send OPEN_REJECT(ZONE VIOLATION). The ZONE VIOLATION shall be cleared if a PHY CONTROL function with operation code of CLEAR ERROR LOG for the specified phy is received from a supervisor.

The TRUSTED bit reports whether the specified phy is currently configured as trusted phy or untrusted phy by the supervisor.

The SUPERVISOR bit reports whether the specified phy is currently configured as a zone supervisor phy.

The SOURCE CHECK bit reports whether the specified phy is doing the source SAS address checking on the specific phy.

The GROUP ID fields reports the source group ID assignment of the specified phy.

...

10.4.3.12 CONFIGURE PHY ZONE function

The CONFIGURE PHY ZONE function sets the expander phy zone membership. This command sets the zone membership assignment for one or multiple expander phys with contiguous phy indexes.

This SMP function shall be supported by SMP target ports in expander devices if the ZONING SUPPORTED bit is set to one in the REPORT GENERAL function. The SMP target should only execute the CONFIGURE PHY ZONE function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

Table 6 defines the CONFIGURE PHY ZONE request.

Table 6 — CONFIGURE PHY ZONE request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (xxh)							
2	Reserved							
3								
4								
5	Ignored							
6	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						(LSB)
n								

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

The FUNCTION field shall be set to xxh. (TBD)

The CRC field is defined in 10.4.3.1.

The START PHY INDEX field defines the first phy index to be configured by the CONFIGURE PHY ZONE command.

The NUMBER OF ZONE PHY ENTRIES field defines how many phy zone entries the CONFIGURE PHY ZONE request intends to configure. This field has a range of 0 to 255.

Note that this command configures one or multiple contiguous expander phys starting from START PHY INDEX.

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

Table 7 defines the PHY ZONE configuration entry descriptor.

Table 7 — PHY ZONE configuration entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0						SOURCE CHECK	TRUSTED	SUPERVISOR
1	Reserved	GROUP ID						

The GROUP ID field specifies the group ID to be assigned to the specified phy.

The SUPERVISOR field specifies whether the specified phy is a supervisor.

The TRUSTED field specifies whether the specified phy is trusted or untrusted.

The SOURCE CHECK field specifies whether the specified phy shall check the SOURCE SAS address against the SAS address in the IDENTIFY address frame received on the specific phy.

Table 8 defines the response format.

Table 8 — CONFIGURE PHY ZONE response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (xxh)								
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 xxh. (TBD)

The FUNCTION RESULT field is defined in 10.4.3.2.

The CRC field is defined in 10.4.3.1.

10.4.3.13 CONFIGURE ZONE PERMISSION function

The CONFIGURE ZONE PERMISSION function provides two modes of setting the values in the zoning permission table controlled by the value of BATCH_SET bit: single set and batch set mode.

In the 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.

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

Note that the zone permission table shall always be symmetrical. ZONE_PERMISSION [X,Y] (Entry X, bit Y) shall always have the same value as ZONE_PERMISSION [Y,X] (Entry Y, bit X).

The single set operation causes the same NEW VALUE to be set to both the [SOURCE GROUP ID bit, TARGET GROUP ID] of the permission table, as well as the [TARGET GROUP bit, SOURCE GROUP ID]. When the zone permission table is set in a batch operation, the supervisor that originates the CONFIGURE ZONE PERMISSION function shall ensure the permission table to be symmetrical.

This SMP function shall be supported by SMP target ports in expander devices if the ZONING SUPPORTED bit is set to one in the REPORT GENERAL function. The SMP target should only execute the CONFIGURE ZONE PERMISSION function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to \ UNKNOWN SMP FUNCTION.

Table 9 defines the CONFIGURE ZONE PERMISSION request.

Table 9 — CONFIGURE ZONE PERMISSION request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (xxh)								
2	Reserved								
3									
4									
5	GENERATION CODE								
6	SET BATCH	SOURCE GROUP ID							
7	NEW VALUE	TARGET GROUP ID							
8	Ignored								
9	Reserved						SUPERVISE THIS	UPDATE COMPLETE	
10	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							
n	(LSB)								

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

The FUNCTION field shall be set to xxh. (TBD)

The SET BATCH field chooses between the single set mode or batch set mode. If the SET BATCH bit is set to 1, the batch set mode is chosen. In the batch mode, the NEW VALUE, SOURCE GROUP ID, and the TARGET GROUP ID fields are ignored. If the SET BATCH bit is set to 0, the single set mode is chosen. The START ZONE ENTRY INDEX and NUMBER OF ZONE PERMISSION ENTRIES shall be set to zero, and the ZONE PERMISSION entry list shall be empty.

The NEW_VALUE field is only used in single set mode. It provides the value for permission table between NEW VALUE, SOURCE GROUP ID and the TARGET GROUP ID. For batch set mode, this field should be set to zero. Note that this value is set to both permission table entry [SOURCE GROUP ID bit, TARGET GROUP ID], and permission table entry [TARGET GROUP ID bit, SOURCE GROUP ID].

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 should be set to zero.

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 should be set to zero.

The SUPERVISE THIS bit is set to indicate that the Supervisor is handing the command off to this expander to become the Supervising expander. The Supervising expander is now responsible for propagating the same zone permission table update to other zoning expanders in the fabric.

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. This may be used as a flag to a supervising expander that a BROADCAST can be generated to the appropriate groups for the changes to zoning.

The GENERATION CODE field specifies the generation code that must be in effect for the function to be accepted. If the GENERATION CODE field is not set to the current generation code, the SMP target shall return a response of (TBD). The generation code shall be incremented by one each time a CONFIGURE ZONE PERMISSION function with UPDATE COMPLETE set high is completed.

The START ZONE ENTRY INDEX field species the first Zone Permission table entry index to be configure in batch set mode. For single set mode, this field should be set to zero.

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 START ZONE ENTRY INDEX in batch set mode. For single set mode, this field should be set to zero.

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 contain zero.

Note that n (total number of bytes), is required to be equal to or less than 1032. This limits the number of CONFIGURE ZONE changes to 63.

The CRC field is defined in 10.4.3.1.

Table 10 defines the ZONE Permission entry descriptor.

Table 10 — ZONE permission entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0	
0	(MSB)	ZONE PERMISSION							
15								(LSB)	

The ZONE PERMISSION field is the zoning permission entry defined in Table 2. The bits corresponding to unused zone indexes should be set to 0.

Table 11 defines the response format.

Table 11 — CONFIGURE ZONE PERMISSION response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (xxh)								
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 xxh. (TBD)

The FUNCTION RESULT field is defined in 10.4.3.2.

The CRC field is defined in 10.4.3.1.

10.4.3.14 REPORT ZONE PERMISSION function

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

This SMP function shall be supported by SMP target ports in expander devices if the ZONING SUPPORTED bit is set to one in the REPORT GENERAL function. The SMP target should only

execute the CONFIGURE ZONE PERMISSION function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

Table 12 defines the REPORT ZONE PERMISSION requests.

Table 12 — REPORT ZONE PERMISSION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (xxh)							
2	Reserved							
3								
4	Ignored							
5	Ignored							
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 xxh. (TBD)

The START ZONE ENTRY INDEX field species 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 START ZONE ENTRY INDEX.

The CRC field is defined in 10.4.3.1.

Table 13 defines the response format.

Table 13 — REPORT ZONE PERMISSION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (xxh)							
2	FUNCTION RESULT							
3	Reserved						CONFIGURING	
4	GENERATION CODE							
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 xxh. (TBD)

The FUNCTION RESULT field is defined in 10.4.3.2.

The GENERATION CODE field indicates the generation of the data returned in the response frame. Each time the zone table changes, the generation code field is incremented. If the management application client detects a different value in the GENERATION CODE field while retrieving one page than it had while retrieving the previous page, it should go back and retrieve all the pages again to obtain a consistent set of information.

The CONFIGURING field indicates the expander is in the process of zone permission table update and 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 START ZONE ENTRY INDEX.

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 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	(MSB)							
	ZONE PERMISSION							
15	(LSB)							

The ZONE PERMISSION field is the zoning permission entry defined in Table 2. The ZONE PERMISSION field is the zoning permission entry defined in Table 2. The bits corresponding to unused zone indexes should be set to 0.

10.4.3.15 REPORT ZONE ROUTE TABLE function

The REPORT ZONE ROUTE TABLE function returns expander route entries from the expander route table within an expander device.

This SMP function shall be supported by SMP target ports in expander devices if the ZONING SUPPORTED bit is set to one and the EXPANDER ROUTE INDEXES field is non-zero in the REPORT GENERAL function. The SMP target should only execute the REPORT ZONE ROUTE TABLE function if the OPEN frame that set up the SMP connect has the ACCESS ZONE MANAGEMENT bit set to 1. If the SMP OPEN frame has ACCESS ZONE MANAGEMENT bit set to 0, the SMP target should generate a SMP response frame with FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

This SMP function may be used as a diagnostic tool to resolve topology issues.

Table 18 defines the REPORT ZONE ROUTE TABLE requests.

Table 18 — REPORT ZONE ROUTE TABLE request

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

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

The FUNCTION field shall be set to xxh.

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 contiguous expander route index starting from START EXPANDER ROUTE INDEX for PHY IDENTIFIER.

The PHY IDENTIFIER field specifies the phy for which the expander route entry is being read (see 4.6.7.3).

The START EXPANDER ROUTE INDEX field specifies the first expander route index for the expander route entry being reported (see 4.6.7.3).

The CRC field is defined in 10.4.3.1.

Table 19 defines the 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 (xxh)							
2	FUNCTION RESULT							
3	Ignored							
4	NUMBER OF ZONE ROUTE TABLE ENTRIES							
5	PHY IDENTIFIER							
6	(MSB)	START EXPANDER ROUTE INDEX						(LSB)
7								
8	GENERATION CODE							
9								
10	Ignored							
11	Reserved					CONFIGURING	END OF ENTRIES	
ZONE ROUTE TABLE entry list								
12	First ZONE ROUTE TABLE entry descriptor							
23								
...	...							
n-15	Last ZONE ROUTE TABLE 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 xxh. (TBD)

The FUNCTION RESULT field is defined in 10.4.3.2.

The NUMBER OF ZONE ROUTE ENTRIES defines how many zone route table entries the REPORT ZONE ROUTE TABLE response frame contains (this value should be smaller than or equal to the NUMBER OF ZONE ROUTE ENTRIES in the REPORT ZONE ROUTE TABLE request frame).

The PHY IDENTIFIER field specifies the phy for which the expander route entry is being read (see 4.6.7.3).

The CONFIGURING field indicates the expander is in the process of updating its ZONE ROUTE TABLE and the expander will issue a BROADCAST MESSAGE when the ZONE ROUTE TABLE update is completed.

The START EXPANDER ROUTE INDEX field specifies the first expander route index for the expander route entry being reported (see 4.6.7.3).

The GENERATION CODE field indicates the generation of the data returned in the response frame. Each time the zone table changes, the generation code field is incremented. If the management application client detects a different value in the GENERATION CODE field while retrieving one page than it had while retrieving the previous page, it should go back and retrieve all the pages again to obtain a consistent set of information.

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 10.4.3.1.

Table 20 defines the ZONE route table entry descriptor.

Table 20 —ZONE ROUTE TABLE entry descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	DISABLE EXPANDER ROUTE ENTRY	GROUP ID						
1	Ignored						TRUSTED	SUPERVISOR
2	Ignored							
3								
4	ROUTED SAS ADDRESS							
11								

The DISABLE EXPANDER ROUTE ENTRY bit specifies whether the ECM shall use the expander route entry to route connection requests (see 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 SUPERVISOR field specifies whether the specified SAS address corresponds to a supervisor.

The TRUSTED field specifies whether the specified SAS address is trusted or untrusted.

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 4.6.7.3).