

To: T10 Technical Committee
 From: Heng Liao, PMC-Sierra (liaoheng@pmc-sierra.com),
 Steve Gorshe, PMC-Sierra (steve_gorshe@pmc-sierra.com),
 Tim Symons, PMC-Sierra (Tim_Symons@pmc-sierra.com) and
 Tom Grieff, HP (thomas.grieff@hp.com)
 Date: August 15, 2005
 Subject: T10/05-144r5 SAS-2 zoning

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.

Related Documents

- sas1r09 - Serial Attached SCSI 1.1 Revision 9

Overview

SAS is starting to gain interest for use in small storage area networks (SANs). For this application, there is a need to add a feature to segregate traffic among devices the way that Fibre Channel provides with zones or Ethernet with virtual LANs.

The SAS zoning model requires no change to end devices; initiators continue to perform normal SAS discovery, and initiators and targets perform OPEN address frame functions as normal. However, SAS zoning means that initiators and targets can be restricted so that they are aware of only part of the entire SAS domain; they see only the portions of the domain that they have been granted permission to see based on a defined zoning permission table.

The SAS zoning model supports up to 128 zoning groups. A set of zoning expander devices interoperate to define a permission table which specifies the membership of each phy to a group. This is referred to as the zoning service delivery subsystem. An end device is managed by the attributes of the phy that it's connected to. Each phy belongs to one group. The permission table also defines whether communication is allowed between multiple groups, defining a zone.

The zoning delivery subsystem is controlled and configured by SMP functions. Attributes are assigned to each phy; group assignments, TRUSTED bit, SUPERVISOR bit, and the zone permission table. Expander devices participating in the zoning delivery subsystem are considered trusted, and can exchange SMP zone commands. Phys linked within the subsystem are marked

as trusted, with all other phy connections marked as untrusted. This defines the boundary of the zoning delivery subsystem..

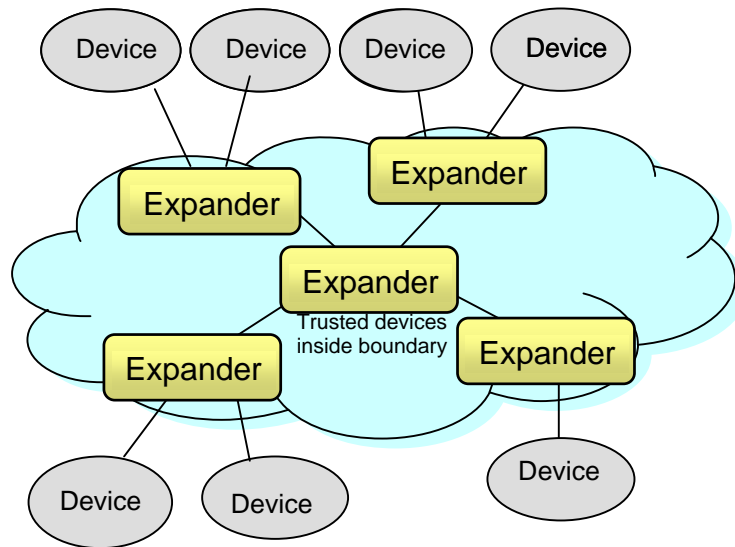


Figure 1 : Zoning Service Delivery Subsystem

SMP functions provide the means to control and configure the expander phy group assignments, the expander phy TRUSTED bit, the expander phy SUPERVISOR bit, and the expander permission table. If a device is attached to an expander phy marked as a supervisor, the device is given supervisor privilege to the zoning domain. Multiple phys in the zoning domain can be marked as supervisors allowing redundant supervisors to coexist in one domain. To avoid information conflicts a single supervising expander is assigned to manage the update of all expander devices inside the boundary.

The *supervising expander* is responsible for propagating zone permission table changes to all zoning expanders in the domain. Within a SAS topology, a single supervising expander is elected based on the value of the SUPERVISING PRIORITY attribute and the EXPANDER SAS ADDRESS value. The supervising expander ensures the consistency of all of the expander permission tables in the SAS topology when multiple supervisors concurrently update the permission table of the zoning domain.

All expanders have a complete copy of the zone permission table. Should the supervising expander fail or be removed from the system, then another of the expanders will assume the function of supervising expander by virtue of its own SUPERVISING PRIORITY attribute.

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

3.1 Definitions

3.1.260 Zoning Expander: A SAS expander capable of being an SMP initiator, and will support the additional 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.

3.1.260 Zone 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.

3.1.260 Zone Route Table: A zoning expander device that supports the table routing method shall contain an 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 TRUSTED, SUPERVISOR, GROUP ID, SUPERVISING PRIORITY.

3.1.260 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. 3.1.261 Topology: A topology is the arrangement by which the devices of a SAS domain are connected to each other.

3.1.262 Trusted Phy: A zoning expander phy that is attached to a device that is required to generate and receive 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. A trusted device can either be a zoning expander that is part of the zoning domain, or be a zone capable end device that participates in the zoning function.

3.1.263 Untrusted phy: A zoning expander phy that is attached to a device that is not required to generate and receive secured zoning information. An untrusted phy is considered to be outside the zoning domain; therefore, the zoning domain does not communicate with zoning information to such a device.

3.1.264 Supervisor: A management entity that is either an expander device inside the zoning service delivery subsystem, or an end device that is attached to a zoning service delivery subsystem. The zoning supervisor is an SMP initiator that generates SMP commands for SAS zoning configuration and management. A SAS zoning domain can have one or multiple supervisors.

3.1.265 Supervising Expander: A SAS zoning expander that is designated by a SAS zoning supervisor expander election process. The supervising expander is the zoning expander device that has highest election priority among all the expanders within the SAS domain. It is responsible for propagating zone permission table updates to all zoning expanders within the SAS zoning domain consistently when the supervisor(s) update the zone permission table.

3.1.266 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 initially when the expander powers up, and subsequently when the expander detects a topology change event.

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 domain. It is up to the system administrator to decide if the SAS domain is going to be zoned or not. If so, the zone service delivery subsystem will be defined by setting the trusted phy attribute on phys marking the periphery of the zone capable group of expander devices.. Legacy expanders can be cascaded outside the boundary of the zoned service delivery subsystem to form a complete SAS domain. All devices attached to the legacy expanders will inherit the group assignment of the phy at the boundary of the zoned service delivery subsystem.

The model supports up to 128 zones. The number of zones each expander device supports is vendor-specific; 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

Zoning expander configuration consists of two parts:

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

4.1.2.1 Per Expander Phy Configuration

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

Table 1: Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)

Name	Description
TRUSTED	If set to 0, this phy is on the boundary of the zoning service delivery subsystem. 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. If set to 1, this phy is inside the service delivery subsystem boundary. The new primitives and frame formats that are defined by the zoning extension are allowed to pass through this phy.
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. 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 GROUP ID = 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.
SUPERVISOR	If set to 1, the device attached to this phy is allowed to originate SMP commands to set up and change zoning configuration. If set to 0, the device attached to this phy is not allowed to originate SMP commands to change the zoning information.
SOURCE CHECK	This 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.

4.1.2.2 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

Note that the shaded area of Table 2: Zoning Permission Table 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

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 equals P[Y,X].

Note that 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, 127] shall be set to 1.
- P[0..126, 0] shall be set to 0.

Note that 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.

4.1.2.3 Special Groups

Group 0 is defined as a group that can only communicate with group 127. It is considered that when a new end device is introduced to a zoned system that it should default to group 0 until the supervisor assigns it to a relevant group membership. This provides a controlled level of management for large topologies.

Group 127 is defined as a group that can communicate with every group in a system. This group is necessary for zone management, so that a supervisor can update every device in the domain without impact to the defined user zones.

4.1.3 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 zoning expanders according to the rules defined in Table 3: PHY ZONE Reassignment after a phy Reset Event.

Table 3: 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.

4.1.4 OPEN Address Frame Handling

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

When an untrusted expander phy receives an OPEN address frame, it sets the ACCESS ZONE MANAGEMENT bit according to the value of the 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 address frame, it sets the ACCESS ZONE MANAGEMENT bit and the SOURCE GROUP ID to 0.

When a trusted expander phy transmits or forwards an OPEN address frame, 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 address frame format inside the zoning service delivery subsystem, across trusted phys, and ensures that the legacy OPEN address frame format is used outside the zoned service delivery subsystem 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 functions during the connection. If the SMP connection is opened with an OPEN address frame with the ACCESS ZONE MANAGEMENT bit set to 1, the expander SMP target supports the complete SMP command set including zone management SMP functions. Otherwise, if it is set to 0, the expander SMP target shall not support zone management SMP functions and, instead, will generate an SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION when a zone management SMP function is requested.

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. For instance, the SMP PHY

CONTROL command from a specified group can only access phys that the source group has permission to see as defined by the permission table. Similarly, the SMP DISCOVER command shall only report the phys the source group can see. The remaining phys are reported as disabled. The 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 as defined by the zone permission table.

4.1.5 SMP Functions

In this proposal, new SMP commands and extensions to existing SMP commands are added to support zone management and enhanced topology discovery functions. The zoning expanders shall support these new commands. A zoning expander shall execute the SMP commands that change the zone configurations only if the command is coming from a supervisor or supervising expander as indicated by the ACCESS ZONE MANAGEMENT bit in the OPEN address frame that initiates an SMP connection.

The SMP REPORT GENERAL command is extended to report the NUMBER OF ZONES supported. A zoning expander shall return the number of zones supported in this field in an SMP response frame. The zoning expander also reports the SAS address of the supervising expander in the topology that is elected by the zoning expanders during the expander topology discovery process.

The SMP DISCOVER command is extended to report the ZONE VIOLATION, ~~SOURCE CHECK~~, TRUSTED, SUPERVISOR, and GROUP ID information that is part of the PHY ZONE configuration of a specific phy. 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.

The SMP CONFIGURE PHY ZONE command is added to 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. Otherwise, the zoning expander shall ignore the command and return a FUNCTION FAILED. The zoning expander that executes the CONFIGURE PHY ZONE command shall send out a BROADCAST(CHANGE) message to the affected groups.

The SMP CONFIGURE ZONE PERMISSION command is added to allow a supervisor permission to update the zone permission table. The supervisor shall only send this command to the supervising expander with the PROPAGATE UPDATE bit set to 1. The supervising expander is responsible for sending SMP CONFIGURE ZONE PERMISSION commands with the PROPAGATE UPDATE bit set to 0 to all zoning expanders in the topology to propagate the zone permission table changes consistently. If the supervising expander has not completed the propagation of the current permission table update, it shall ignore any new updates to the permission table and return a FUNCTION FAILED. If a non-supervising expander receives a SMP CONFIGURE ZONE PERMISSION command with the PROPAGATE UPDATE bit set to 1, the expander shall ignore the command and return a FUNCTION FAILED. The CONFIGURE ZONE PERMISSION commands support 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. A zoning expander shall send out BROADCAST messages to the affected groups after it executes a CONFIGURE ZONE PERMISSION command with the UPDATE COMPLETE bit set to 1. The SMP REPORT ZONE PERMISSION command reports the zone permission table entries.

The SMP REPORT ZONE ROUTE TABLE command 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: TRUSTED, SUPERVISOR, GROUP ID, and ATTACHED DEVICE TYPE.

4.1.6 ZONE Management Updates

As described in the previous section, 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 expander self-discovery process propagates the new PHY ZONE configuration associated with the attached SAS address. Before the topology rediscovery is completed in the SAS domain, there might be a transition period where devices in the domain might send OPEN address frames to the reassigned device. Such OPEN address frames could be rejected by expanders along the way, or by the last expander that is connected to the reassigned device if the new group assignment prohibits the OPEN address frame from being accepted. In both cases, the new group assignment takes effect after the CONFIGURE PHONE ZONE command is completed.

Through the single supervising expander, the permission table update consistency is ensured among multiple supervisors, such that only one permission table update can take place at any time. During the period where the supervising expander is propagating the permission table updates to other expanders, the zone permission tables among the expanders could be temporarily inconsistent with each other. For example, if the permission between Group X and Y is changed from 0 to 1, during the propagation period, some expanders may have changed $P(X,Y)$ to 1, while others still have the old $P(X,Y)$ value of 0. This means that any new OPEN address frames from Group X may still be rejected by an expander that has the old permission value until the propagation process is complete. The new permission does not take effect until the propagation is completed. Similarly, if permission $P(X,Y)$ is being changed from 1 to 0, an OPEN address frame from Group X may still be routed to Group Y during the propagation period. However, as soon as the propagation process is complete, no new OPEN address frames will be routed between Groups X and Y. Again, at the system level, this means the permission update does not fully take effect until the propagation is completed. In another example, if the permission table update uses the batch mode, there could be temporary inconsistency between $P(X,Y)$ and $P(Y,X)$, during the update period. OPEN address frames may be routed from Group X to Y, but not on the reverse direction (rejection). During the propagation period, the new permissions between the affected groups may or may not take full effect. However, after the propagation process is completed, the new permissions on the affected groups take effect.

The supervisor is required to re-download the permission table whenever it detects a topology change in a SAS domain that involves expanders. That is, when an existing expander(s) is removed from or added to the SAS domain, the supervisor must send the entire permission table to the supervising expander. This prevents the addition or removal of expanders in a SAS domain from creating inconsistency in the permission tables.

Therefore, the zoning expanders shall not affect existing connections, or new OPEN address frame processing during zone management update periods. During a zone management update period, the updated PHY ZONE configuration and the permission table update is either fully in effect, partially in effect, or not in effect at all. However, after the update and completion of the propagation period, the new zone management configuration shall be in full effect. At the system level, the normal data traffic is delivered by the SAS domain without service interruption.

4.1.7 Expander Zone Route Table

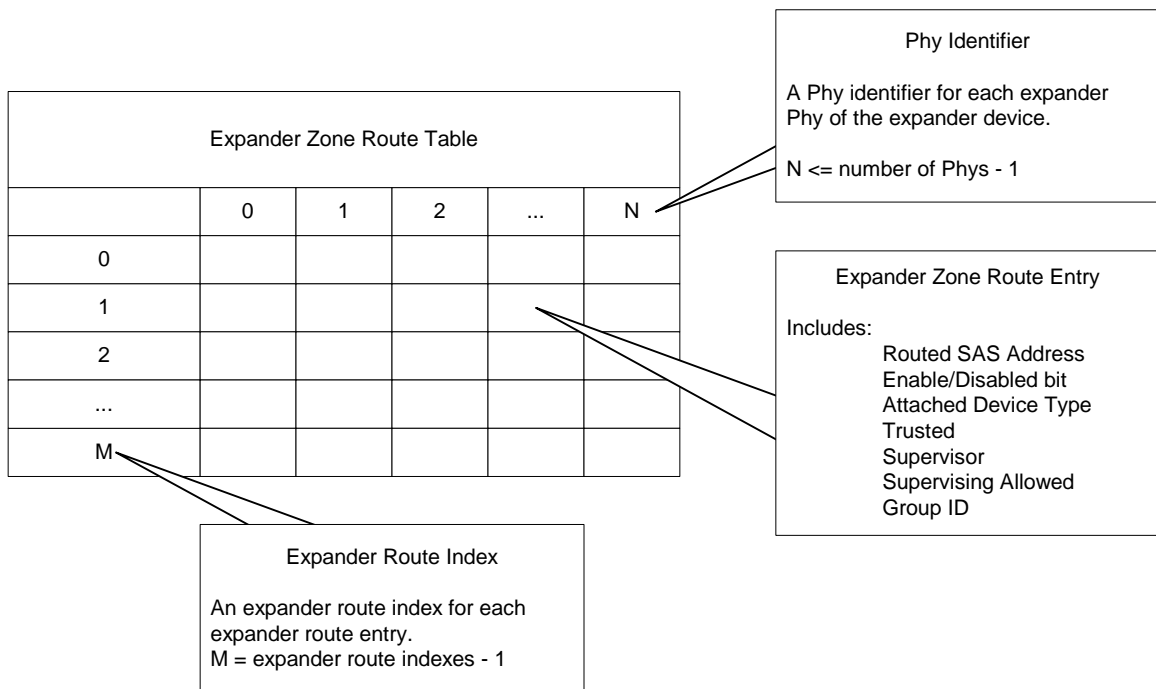
A zoning expander device that supports the table routing method shall contain an expander zone route table. The expander zone route table is a structure that provides an association between routed SAS addresses to expander phy identifiers using the PHY ZONE configuration attributes TRUSTED, SUPERVISOR, 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 using the SMP REPORT GENERAL function (see Section 10.4.3.3). 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 with the SMP REPORT ZONE ROUTE TABLE function. All zoning expanders are responsible for configuring the zone route table automatically through the topology discovery process described in Section 4.1.8.

Figure 2: Expander Zone Route Table Example shows a representation of an expander route table.

Figure 2: Expander Zone Route Table Example



4.1.8 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 in SAS 1.1 with the additional information about the 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 as a supervising expander as well its priority for election. By default, all zoning expanders shall have this attribute set to 0 to not allow the device to be a supervising expander. The SAS domain supervisor can change the SUPERVISING PRIORITY attribute of the zoning expanders within a SAS domain, thereby controlling which expanders can be used as the supervising expander and their priorities for election. All zoning expanders with a non-zero SUPERVISING PRIORITY attribute become candidates for the supervising expander election process.

Note that the zoning expander topology process also accomplishes the election of supervising expander implicitly. As a zoning expander traverses the adjacent expanders in the SAS domain, the supervising expander is elected to be the zoning expander with highest election priority value within the SAS domain.

The election priority of a zoning expander is defined in Table 4: Zoning Expander Election Priority

Table 4: Zoning Expander Election Priority

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 supervising expander address is reported by the SMP REPORT GENERAL response frame, any supervisor end device can use the REPORT GENERAL command to obtain the supervising expander address without having to implement the supervising expander election function on the host device.

If the SUPERVISING PRIORITY attribute is 0, the expander shall not be chosen as the supervising expander. If all expanders within a SAS domain have supervising priority attribute of 0, then the election process will not elect a supervising expander, and REPORT GENERAL response frame shall return value of 0 in the SUPERVISING EXPANDER ADDRESS to indicate the election process has failed.

The zoning expander shall support the host topology discovery process defined by SAS 1.1 through the REPORT GENERAL, DISCOVER, and REPORT ROUTE INFO commands. Note that the zoning expanders shall generate a response frame to these commands based on the group assignment of the host device that is carried in the SOURCE GROUP ID of the OPEN address frame that sets up the SMP connection such that the host devices can only discover the other host devices as is permitted by the permission table.

To support topology discovery, all zoning expanders should report the PHY CHANGE COUNT according to the group assignment of the host. The PHY CHANGE COUNT maintains the number of BROADCAST (CHANGE) messages originated by the phy and is reported by the DISCOVER command. When a zoning expander receives a DISCOVER command, the expander

shall report the PHY CHANGE COUNT as 0 if the SOURCE GROUP ID field of the OPEN address frame is not permitted to access the requested phy. Otherwise, the actual PHY CHANGE COUNT value is reported.

To support topology discovery, all zoning expander may report the EXPANDER CHANGE COUNT according to the group assignment of the host using the REPORT GENERAL function. The EXPANDER CHANGE COUNT maintains the number of BROADCAST (CHANGE) messages originated by the expander that are propagated to the group where the host is a member.

4.1.9 Broadcasts

The ZONED BROADCAST address frame (see Section 7.8.x) is used within a zoning service delivery subsystem to propagate BROADCAST messages within zone boundaries. A ZONED BROADCAST address frame may be transmitted three times to provide redundancy against broadcast messages that have been lost due to bit errors. 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 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.

When a trusted 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.

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

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

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 legacy BROADCAST messages still applies to this broadcast method. That is, if more than one phy is joined in a wide expander port, then the address frame should only be sent at least one time, on any or all-available phys.

When transmitting a ZONED BROADCAST event on a trusted 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).

When transmitting a ZONED BROADCAST event on an untrusted phy, the expander shall transmit a BROADCAST primitive that represents the type of broadcast event represented by the ZONED BROADCAST event, but the SOURCE GROUP ID 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 they shall transmit BROADCAST address frames to devices within the zoned area.

4.1.7 ~~Source SAS address checking~~

For phys with SOURCE CHECK bit set to one and 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?~~

4.1.10 Zone Checks

In typical SAS operation, when an expander phy receives an OPEN address frame, an expander 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.

During SAS zoning operation, if the OPEN address frame is received from an untrusted phy, 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 trusted 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 direct routing table will find a match. In this case, the group ID assignment of the matching destination phy is used as the Destination group ID. 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 and a destination port. If the subtractive routing method is used to route an OPEN address frame, 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 rejected with OPEN_REJECT(ZONING VIOLATION) response. 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.

7.8 Address Frames

7.8.1 Address Frames Overview

...

The ADDRESS FRAME TYPE field specifies the type of address frame. This field is defined in Table 5: ADDRESS FRAME TYPE Field. This field determines the definition of the frame type dependent bytes.

Table 5: 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 6: OPEN Address Frame Format defines the OPEN address frame format used for connection requests.

Table 6: 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.

...

7.8.4 ZONED BROADCAST Address Frame

Table 7: ZONED BROADCAST Address Frame Format defines the ZONED BROADCAST address frame format.

Table 7: 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								

....

....

The **BROADCAST TYPE** field defines the type of broadcast that is being made.

Table 8: BROADCAST TYPE Field defines the values of the BROADCAST TYPE field.

Table 8: 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 forward 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. Table 2: 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 9: REPORT GENERAL Response defines the REPORT GENERAL response frame format.

Table 9: 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				Reserved		CONFIGU RING	CONFIGU RABLE ROUTE TABLE
11	SUPERVISING PRIORITY				SUPERVISING STATUS			
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20								
27	SUPERVISING EXPANDER ADDRESS							
28	(MSB)	CRC						(LSB)
31								

The NUMBER OF ZONES field indicates the number of zones supported (Group 0, Group 127, and the user-defined groups [126:1]). For expanders that do not support zoning, this field should be set to 0. Note that the values 1 and 2 are illegal for this field. Note that 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 prepended to the expander device SAS address. The 68-bit field is compared with other, and the expander devices and the highest value address will be elected as the supervisor. This provides a mechanism for system management to prioritise the device to be selected as a supervising device.. It is expected that two of these bits should be configured by hardware. When the SUPERVISING PRIORITY field is set to 0, the expander is not allowed to be a candidate for supervising expander election.

The SUPERVISING STATUS field returns the status of the supervising expander zone permission update propagation. Table 10: SUPERVISING STATUS Field defines the values of SUPERVISING STATUS field.

Table 10: SUPERVISING STATUS Field

Code	SUPERVISING STATUS Description
0000	The current expander is not the supervising expander.
0001	The current expander is the supervising expander and it is in the process of propagating the zone permission update that was successfully received from a supervisor through the CONFIGURE ZONE PERMISSION command with the FUNCTION RESULT code of FUNCTION ACCEPTED.
0010	The current expander is the supervising expander and it has successfully propagated the zone permission update that was successfully received from a supervisor through the CONFIGURE ZONE PERMISSION command with the FUNCTION RESULT code of FUNCTION ACCEPTED, or it has not received any CONFIGURE ZONE PERMISSION command.
0011	The current expander is the supervising expander and it has failed to propagate the zone permission update that was successfully received from a supervisor through the CONFIGURE_ZONE_PERMISSION command with the FUNCTION RESULT code of FUNCTION ACCEPTED. The propagation process has been aborted and the supervisor should re-download the zone permission table to the latest supervising expander.
0100~1111	Reserved

The SUPERVISING EXPANDER ADDRESS field indicates the SAS address of the supervising expander that is elected during the topology discovery process. The expander returns a 0 in the SUPERVISING EXPANDER ADDRESS field if the supervising expander election process has failed to elect a supervising expander.

10.4.3.5 DISCOVER Function

The DISCOVER command is responded to 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 disabled. Table 11: DISCOVER Response defines the DISCOVER response format.

Table 11: 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	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	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	SOURCE CHECK	TRUSTED	SUPERVISOR
49	Reserved	GROUP ID						
50	Vendor Specific							
51								
52	(MSB)	CRC						
55							(LSB)	

...

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 TRUSTED bit reports whether the supervisor has designed the specified phy as a trusted or untrusted phy.

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. The Zoning expander 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 0 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 will generate a SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

Table 12:CONFIGURE PHY ZONE Request defines the CONFIGURE PHY ZONE request format.

Table 12: 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	Reserved								
5									
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							
n								(LSB)	

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 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.

The PHY ZONE configuration entry descriptor is defined in Table 13: PHY ZONE Configuration Entry Descriptor.

Table 13: PHY ZONE Configuration Entry Descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	Reserved				Reserved	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 14: CONFIGURE PHY ZONE Response defines the CONFIGURE PHY ZONE response format.

Table 14: 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 Section 10.4.3.2.

The CRC field is defined in Section 10.4.3.1.

10.4.3.13 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_SET bit: 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.

In 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, 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 supervisor that originates the CONFIGURE ZONE PERMISSION function 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 0 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 should generate an SMP response frame with the FUNCTION RESULT set to UNKNOWN SMP FUNCTION.

Table 15: CONFIGURE ZONE PERMISSION Request show the request format.

Table 15: 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	Reserved							
5								
6	Reserved	SOURCE GROUP ID						
7	NEW VALUE	TARGET GROUP ID						
8	Reserved							
9	OPERATION				Reserved	Reserved	PROPAGATE UPDATE	UPDATE COMPLETE
10	START ZONE ENTRY INDEX							
11	NUMBER OF ZONE PERMISSION ENTRIES							
	ZONE PERMISSION entry list							
12	_____							
27	First ZONE PERMISSION entry descriptor _____							
...	_____							
n-20	_____							
n-4	Last ZONE PERMISSION entry descriptor _____							
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 OPERATION field defines operation codes to be executed for this command as shown in Table 16: OPERATION Field.

Table 16: OPERATION Field

Code	SUPERVISING PRIORITY
0000	SET SUPERVISING PRIORITY attribute to 00
0001	SET SUPERVISING PRIORITY attribute to 01
0010	SET SUPERVISING PRIORITY attribute to 10
0011	SET SUPERVISING PRIORITY attribute to 11
0100~0111	Reserved
1000	SET SINGLE: This chooses the permission table single set function. This mode sets the permission between a pair of device groups as indicated by SOURCE GROUP ID and the TARGET GROUP ID fields to the value of NEW VALUE field.
1001	SET BATCH: This chooses the batch set mode operation, which sets a number of contiguous zone permission table entries at a time as specified by START ZONE ENTRY INDEX and NUMBER OF ZONE PERMISSION ENTRIES.
1010~1111	Reserved

The NEW VALUE field is only used in single set mode. It provides the value for the permission table between the new value, Source group ID and the Target group ID. For batch set mode, this field should 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 should 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 should be set to 0.

The PROPAGATE UPDATE bit is set to 1 to indicate the UPDATE ZONE PERMISSION command is sent from a supervisor to the supervising expander of the zoning domain. The PROPAGATE UPDATE bit is set to 0 to indicate the UPDATE ZONE PERMISSION command is sent from a supervising expander to other zoning expanders of the zoning domain.

If a non-supervising expander receives a CONFIGURE ZONE PERMISSION command with the PROPAGATE UPDATE bit set to 0, the non-supervising expander shall update its own permission table without propagation. If a non-supervising expander receives a CONFIGURE ZONE PERMISSION command with the PROPAGATE UPDATE bit set to 1, the non-supervising expander shall ignore this command and return a FUNCTION FAILED.

If a supervising expander receives a CONFIGURE ZONE PERMISSION command when the PROPAGATE UPDATE is set to 1, the supervising expander shall update its permission table accordingly and propagate the same updates to all other zoning expanders within the domain by sending SMP CONFIGURE ZONE PERMISSION commands to the individual expanders with the PROPAGATE UPDATE bit set to 0. 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 a FUNCTION FAILED.

If a supervising expander receives a CONFIGURE ZONE PERMISSION command with the PROPAGATE UPDATE bit set to 0, the supervising expander shall ignore the command and return a FUNCTION FAILED.

The UPDATE COMPLETE bit indicates whether the current CONFIGURE ZONE PERMISSION command is the last 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 receiving a CONFIGURE ZONE PERMISSION command with the UPDATE COMPLETE bit set to 1.

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 should 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 should 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 contain 0. 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 Section 10.4.3.1.

The ZONE PERMISSION entry descriptor is defined in Table 17: ZONE Permission Entry Descriptor

Table 17: 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 18: CONFIGURE ZONE PERMISSION Response defines the CONFIGURE ZONE PERMISSION response format.

Table 18: 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 Section 10.4.3.2.

The CRC field is defined in Section 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 a SMP target ports in an expander device if the NUMBER OF ZONES SUPPORTED field is greater than 0 in the REPORT GENERAL function.

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

Table 19: REPORT ZONE PERMISSION Request defines the REPORT ZONE PERMISSION request format.

Table 19: 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								
5	Reserved							
6	START ZONE ENTRY INDEX							
7	NUMBER OF ZONE PERMISSION ENTRIES							
8	(MSB)	CRC						(LSB)
11								

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

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

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 Section 10.4.3.1.

Table 20: REPORT ZONE PERMISSION Response defines the REPORT ZONE PERMISSION response format.

Table 20: 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							CONFIGU RING
4	Reserved							
5								
6	START ZONE ENTRY INDEX							
7	NUMBER OF ZONE PERMISSION ENTRIES							
	ZONE PERMISSION entry list							
8	First ZONE PERMISSION entry descriptor							
23								
...	...							
n-19	Last ZONE PERMISSION entry descriptor							
n-4								
n-3	(MSB)	CRC						
n								(LSB)

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

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

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.

The ZONE PERMISSION entry descriptor is defined in Table 21: ZONE PERMISSION Entry Descriptor.

Table 21: 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 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.

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

Table 22: REPORT ZONE ROUTE TABLE Request defines the REPORT ZONE ROUTE TABLE request format.

Table 22: 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)	STARTING EXPANDER ROUTE INDEX							
7								(LSB)	
8	Reserved								
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 the contiguous expander route index starting from the 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 (see Section 4.6.7.3).

The CRC field is defined in Section 10.4.3.1.

Table 23: REPORT ZONE ROUTE TABLE Response defines the REPORT ZONE ROUTE TABLE response format.

Table 23: 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	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 xxh. (TBD)

The FUNCTION RESULT field is defined in 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 CONFIGURING bit indicates that the expander device is in the process of updating its zone route table. The STARTING EXPANDER ROUTE INDEX field indicates the first expander route index for the expander route entry being reported (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 Section 10.4.3.1.

The Zone Route Table Entry descriptor format is defined in Table 24: ZONE ROUTE TABLE Entry Descriptor.

Table 24: ZONE ROUTE TABLE Entry Descriptor

Byte\Bit	7	6	5	4	3	2	1	0
0	DISABLE EXPANDER ROUTE ENTRY	Reserved						
1		ATTACHED DEVICE TYPE			SUPERVISING PRIORITY		TRUSTED	SUPERVISOR
2	Reserved	GROUP ID						
3	Reserved							
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 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.

A SUPERVISOR bit set to one indicates the specified SAS address corresponds to a supervisor. A SUPERVISOR bit set to zero indicates the specified SAS address does not correspond to a supervisor. A TRUSTED bit set to one indicates that the the specified SAS address is trusted. A TRUSTED bit set to zero indicates that the the specified SAS address is not trusted.

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.

The SUPERVISING PRIORITY is a 4 bit field that is set to prioritise which expander device in a system should be elected as the supervising device. This is a system configuration feature.

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

Table of contents

Revision History	1
Related Documents	1
Overview	1
3.1 Definitions	3
4.1 Zoning Model	4
4.1.1 Zone Model Overview	4
4.1.2 Zoning Configuration	4
4.1.2.1 Per Expander Phy Configuration	4
4.1.2.2 Zoning permission table	6
4.1.2.3 Special Groups	6
4.1.3 Device Group Reassignment.....	7
4.1.4 OPEN Address Frame Handling.....	7
4.1.5 SMP Functions	8
4.1.6 ZONE Management Updates	9
4.1.7 Expander Zone Route Table	10
4.1.8 Topology Discovery	11
4.1.9 Broadcasts.....	12
4.1.7 Source SAS address checking	13
4.1.10 Zone Checks.....	13
7.8 Address Frames	14
7.8.1 Address Frames Overview	14
7.8.3 OPEN Address Frame	15
7.8.4 ZONED BROADCAST Address Frame	16
10.4.3 SMP Functions	17
10.4.3.1 SMP Function Request Frame Format.....	17
10.4.3.2 SMP Function Response Frame Format.....	17
10.4.3.3 REPORT GENERAL Function.....	17
10.4.3.5 DISCOVER Function	18
10.4.3.12 CONFIGURE PHY ZONE Function.....	20
10.4.3.13 CONFIGURE ZONE PERMISSION Function	22
10.4.3.14 REPORT ZONE PERMISSION Function	26
10.4.3.15 REPORT ZONE ROUTE TABLE Function.....	28

Table of Figures

Figure 1 : Zoning Service Delivery Subsystem	2
Figure 2: Expander Zone Route Table Example	10

Table of Tables

Table 1: Per phy Zoning Configuration (PHY_ZONE CONFIGURATION)	4
Table 2: Zoning Permission Table.....	6
Table 3: PHY ZONE Reassignment after a phy Reset Event	7
Table 4: Zoning Expander Election Priority	11
Table 5: ADDRESS FRAME TYPE Field	14
Table 6: OPEN Address Frame Format	15
Table 7: ZONED BROADCAST Address Frame Format	16
Table 8: BROADCAST TYPE Field	16
Table 9: REPORT GENERAL Response	17
Table 10: SUPERVISING STATUS Field.....	18
Table 11: DISCOVER Response.....	19
Table 12:CONFIGURE PHY ZONE Request	21

Table 13: PHY ZONE Configuration Entry Descriptor.....	21
Table 14: CONFIGURE PHY ZONE Response	22
Table 15: CONFIGURE ZONE PERMISSION Request.....	23
Table 16: OPERATION Field.....	24
Table 17: ZONE Permission Entry Descriptor	25
Table 18: CONFIGURE ZONE PERMISSION Response	25
Table 19: REPORT ZONE PERMISSION Request	26
Table 20: REPORT ZONE PERMISSION Response.....	27
Table 21: ZONE PERMISSION Entry Descriptor	28
Table 22: REPORT ZONE ROUTE TABLE Request	28
Table 23: REPORT ZONE ROUTE TABLE Response	29
Table 24: ZONE ROUTE TABLE Entry Descriptor.....	30