To: T10 Technical Committee
From: Heng Liao, PMC-Sierra (liaoheng@pmc-sierra.com),
Steve Gorshe, PMC-Sierra (steve_gorshe@pmc-sierra.com), and
Tom Grieff, HP (thomas.grieff@hp.com)
Date: ~~11 May~~ June 13, 2005 2005
Subject: T10/05-144r~~32~~ 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.

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. Within a SAS topology, the zoning expander with SUPERVISING ALLOWED attribute with largest SAS address value is elected to be the supervising expander~~.~~. The purpose of having a supervising expander is to offload the supervisor and the associated traffic from the management port for permission table propagation and to ensure the consistency of permission tables in all the expanders in the SAS topology when multiple supervisors are doing permission table updates simultaneously.

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

Suggested Changes

3.1 Definitions

3.1.259 Fabric: A switch fabric is an interconnected network of switching devices that provide device-to-device pathways for transferring information. A SAS fabric consists of expander devices in a topology that connects them together to allow end devices such as SAS initiators and SAS target devices and expander devices to communicate with each other.

3.1.260 Zoning Fabric: A zoning fabric can be part of a SAS fabric, or be a complete SAS fabric that provides not only the services of a normal SAS fabric, but also implements and enforces zoning access control policies and zone management functions.

3.1.261 Topology: The arrangement in which the devices of a SAS domain are connected to each other. Sometimes, the term fabric and topology is used interchangeably to refer to the SAS domain and how devices are connected together.

3.1.262 Trusted Phy: A trusted Phy is a zoning expander Phy that is attached to a device that can be trusted to generate and receive secured information about zoning that the zoning fabric relies on the enforce the access control policies, which includes the zoning OPEN frames and Zoning broadcast frames. A trusted device can either be a zoning expander that is part of the zoning fabric, or be a secured end device that participate in the zoning function in a SAS domain with full knowledge about receiving and generating zoning traffic including zoning OPEN address farmes and zoning broadcast frames.

3.1.263 Untrusted Phy: A unstrusted Phy is a zoning expander Phy that is attached to a device that can not be trusted to generate and receive secured information about zoning. An untrusted Phy is considered to be outside the zoning fabric, therefore the zoning fabric does not expect to communicate with such a device with zoning information.

3.1.264 Supervisor: A supervisor is a management entity that is either part of the zoning fabric, or is an end device that is attached to a zoning fabric. The zoning supervisor is a SMP initiator that is capable of generate SMP commands for SAS zoning configuration and management. A SAS zoning fabric can have one or multiple supervisors.

3.1.265 Supervising Expander: A supervising expander is an SAS zoning expander that is designated by the SAS zoning supervisor expander election process. It has the largest SAS address among the zoning expander with the SUPERVISNG ALLOWED attribute. It is responsible for propagating zone permission table update to all zoning expanders within the SAS zoning fabric consistently when one or multiple supervisor attempts to update the zone permission table.

3.1.266 Expander Self-Discovery: The expander self-discovery process is an expander function that exchange information about the SAS topology using SMP protocol with adjacent expanders in order to come to accurate knowledge about the SAS zoning topology, such information is then 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 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 phy zoning configuration (PHY_ZONE CONFIGURATION)

Name	Description
TRUSTED	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. 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.
GROUP ID[6:0]	The GID defines the zoning Group ID in the range from 0..127. 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). 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. GID=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.~~

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 Device Group Reassignment

As noted, the group assignment of a device is defined as PHY ZONE configuration that is associated with the attached expander Phy. The PHY ZONE configuration shall be maintained by the zoning expanders according to following rules:

1) If the zoning expander detects a new SAS device attached to an expander Phy (i.e. the attached SAS address in the IDENTIFY frame is different from the previous attached SAS address), the expander shall automatically assign the Phy to the default group ID of 0.

2) If the zoning expander detects the current SAS device attached to an expander Phy to have the same SAS address of the last device attached to the expander Phy before the link lose PHY_RDY, the expander shall automatically reassign the expander Phy to the PHY ZONE configuration of the previous device prior to losing link PHY_RDY.

3) If the zoning expander detects a SAS device (from the identified SAS address) that has moved from a different Expander Phy that is part of the same expander, or a different expander in the topology, the expander shall assign the expander Phy to default group ID of 0. ~~The supervisor may discover the device has moved from one location of the topology to a different location of the fabric. And the supervisor may reassign the device to the same PHY~~

ZONE configuration that it used to belong, or reassign the device to a different PHY ZONE configuration based on the intention of the administrator. The policy of the supervisor is beyond the scope of this specification. The supervisor may reassign the device to a group other than Group 0 at a later time. The policy of supervisor reassignment of group ID to a device is beyond the scope of this specification.

4) For a SATA device that has achieved PHY_RDY, the expander shall assign it to group ID 0 if the device has not lost and regain link PHY_RDY due to the expander executing a SMP PHY_CONTROL LINK_RESET or HARD_RESET command to the affected PHY. If the SATA device has lost and regain link PHY_RDY due to the expander executing a SMP LINK_RESET and HARD_RESET command on the attached expander Phy, the Expander shall restore the previous PHY_ZONE configuration of the Expander Phy.

Editor's Note 1: A vendor may choose to implement an supervisor policy that always automatically reassigns the PHY_ZONE configuration to a device after the device has moved from one location of the topology to a different location. This would allow a persistent SAS WWN based group assignment to be implemented. A different vendor may choose to implement a supervisor policy that requires human administrator's intervention on assigning the newly moved device to a new group. Such design choices are implementation specific features of the vendors' management software running in a supervisor device and are beyond the scope of this specification.

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 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.

4.1.5 SMP functions

New SMP commands are added to support the 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 “ACCESS ZONE MANAGEMENT” bit in the OPEN frame that initiates the SMP connection.

The SMP REPORT GENERAL command is extended to report NUMBER OF ZONES SUPPORTED. A zoning expander shall return the number of zones supported in this field in the 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 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 the specific Phy. When the SMP DISCOVER command is executed from a source group (as indicated by the SGID in the OPEN 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 frame has ACCESS ZONE MANAGEMENT bit set to 1. Otherwise, the Zoning expander shall ignore the command and return FUNCTION FAILED. The zoning expander that executes the CONFIGURE PHY ZONE command shall send out a BROADCAST(CHANGE) with SGID of 127 and force PHY change count associated with the affected PHY and the expander change count to be incremented.

The SMP CONFIGURE ZONE PERMISSION command is added to allow a supervisor to update the ZONE PERMISSION TABLE. The supervisor shall only send this command to the supervising expander with PROPAGATE UPDATE bit set to 1. The supervising expander is responsible for sending SMP CONFIGURE ZONE PERMISSION commands with 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 permission table and return FUNCTION FAILED. If a non-supervising expander receives a SMP CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE bit set to 1, the expander shall ignore the command and return 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 pair of group ID and ensures the permission for the reverse direction is set symmetrically. The batch mode allow multiple rows of the permission table is downloaded at each 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 expected to check and enforce the symmetry of the zone permission table. A zoning expander shall send out BROADCAST (CHANGE) with SGID of 127 after it executes a CONFIGURE ZONE PERMISSION command with UPDATE COMPLETE bit set to 1 if it contains expander Phys belong to the groups whose permission has been changed. The zoning expander shall force the PHY change counter to be incremented on the Phys that are affected by the permission change and the expander change count shall be incremented if the BROADCAST(CHANGE) message is sent.

The SMP REPORT ZONE PERMISSION reports the zone permission table entries.

The SMP REPORT ZONE ROUTE TABLE reports the zone route table, 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 routing table. Each entry of the table is extended to contain fields in addition to attached SAS address including TRUSTED, SUPERVISOR, GROUP ID, ATTACHED DEVICE TYPE.

4.1.6 ZONE Management Updates

As described in the previous section, there are two types of zone configuration updates: 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 is being arbitrated.

The PHY ZONE update operation is an atomic operation send from a supervisor to a zoning expander with a single SMP command. The new PHY ZONE configuration associated with the specific PHY takes effect instantaneously after the CONFIGURE PHY ZONE command is executed. The expander self-discovery process propagates the new PHY ZONE configuration association with the attached SAS address. Prior to topology rediscovery is completed in the SAS domain, there could be a period of time when the devices in the domain would still send OPENs to the reassigned device, such OPEN 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 from being accepted. In both cases, the new group assignment takes effect.

Through the single supervising expander, the PERMISSION table update consistency is ensure among multiple supervisors, such that only one permission table update can take place at any time. During the time period the supervising expander is propagating the permission table updates to other expanders, there could be temporary inconsistency of the zone permission tables among the expanders. It can illustrated that the SAS domain can still work properly with the temporary inconsistency using a simple example. If the permission between group X and Y is changed from 0 to 1, during the propagation period, some expander may have changed P(X,Y) to 1, while others still has the P(X,Y)=0. This only means any new OPEN from group X may still be rejected by the expander that has the old permission value, until the propagation is completed. The net effect is the new permission does not take effect until the propagation is completed. Similarly, if the permission P(X,Y) is being changed from 1 to 0, the OPEN from group X may still be routed to group Y during the propagation period, but as soon as the propagation is completed, no new OPEN can be routed between group X and Y. Again, at the system level, this means the permission update does not fully take effect until the propagation is completed. In yet 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 this update period, OPENs could be routed from X to Y, but not on the reverse direction (rejection). This is may result in higher layer time out on the devices in the affect group, but the normal error recovery mechanism should be able to handle such a condition because the permission table change is expected to cause such disruption on the devices in the affected groups (either X can now access a new set of devices in group Y in the case of 0->1 change, or X can no longer access devices in group Y in the case of 1->0 change).

The supervisor is required to redownload the permission table whenever it detects a topology change in a SAS domain that involves expanders, i.e. when one or more existing expander is removed from the SAS domain, or one or more expanders are added to a SAS domain, the supervisor must send the entire permission table to the latest supervising expander. This prevents addition or removal of expander in a SAS domain creates inconsistency in permission table among the expanders.

In conclusion, the zoning expanders shall not affect existing connections, or new OPEN processing during the zone management update periods. During the 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. But after the update and propagation period, the new zone management configuration shall be in full effect. At the system level, the normal data traffic should still be able to be 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, the PHY ZONE configuration attributed such as TRUSTED, SUPERVISOR, GROUP ID, Supervising Allowed, attached device type. Each association represents an expander route entry.

An expander device reports the size of its expander route table and indicates if the expander route table is configurable in the SMP REPORT GENERAL function (see 10.4.3.3). Each expander route entry shall be disabled after power on.

A management application client may reference a specific expander route entry within an expander route table with the SMP REPORT ZONE ROUT TABLE function. A zoning expander is responsible for configuring its zone route table automatically (self-configuring expanders) through the topology discover process described in 4.1.8.

Figure 44 shows a representation of an expander route table.

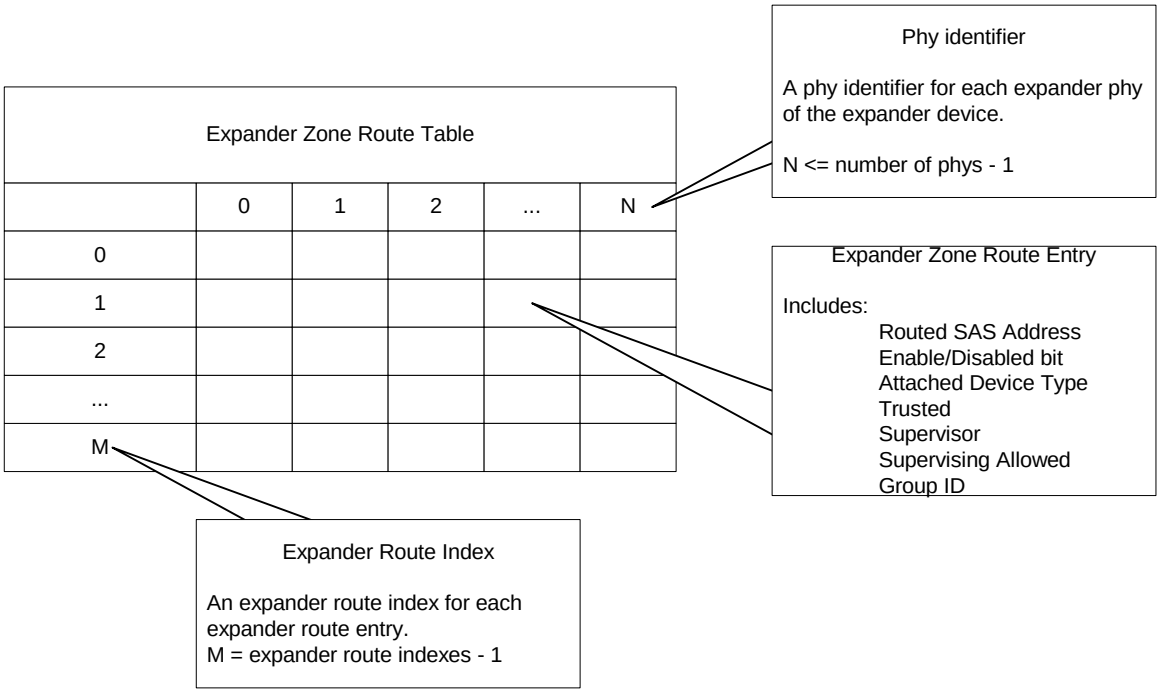

Figure 44 — Expander zone route table example

4.1.~~8~~7 Topology Discovery

All Zoning expanders shall be self-configuring expanders that are responsible for configuring 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 of about the PHY ZONE assignment of the associated device. The zoning expander traverse the SAS topology and use 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 ALLOWED attribute of a zoning expander determines whether the zoning expander can be a candidate as a supervising expander. By default, all zoning expanders shall have this attribute disabled. The SAS fabric supervisor can enable/disable the SUPERVISING ALLOWED attribute of zoning expanders within a SAS fabric, thereby controlling which

expanders can be used as the supervising expander. All zoning expanders with the SUPERVISING ALLOWED attribute become candidates for supervising expander election process.

The zoning expander topology process also accomplishes the election of supervising expander implicitly. As a zoning expander traverse the adjacent expanders in the SAS domain, the zoning expander can compare the SAS addresses of all zoning expanders with SUPERVISING ALLOWED attribute enabled in the SAS domain, and find the zoning expander with SUPERVISING ALLOWED attribute within the fabric that has the largest SAS address. This expander address is elected to be the supervising expander address. The supervising expander address is recalculated every time a zoning expander redo topology discovery. Since the supervising expander address is reported by 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.

The zoning expander shall support the host topology discovery from SAS 1.1 compliant host devices through REPORT GENERAL, DISCOVER, and REPORT ROUTE INFO commands. Note that the zoning expanders shall generate response frame to these commands based on the group assignment of the host device that is carried in the SGID of the OPEN frame that sets up the SMP connection such that the host devices can only discover the host devices that the host device is permitted to access based on the PERMISSION table.

To support topology discovery, all zoning expander ~~shall~~ should report PHY CHANGE count according to the group assignment of the host. The PHY change count maintains the number of BROADCAST(CHANGE) 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 SGID of OPEN frame is not permitted to access the requested Phy, other wise, the actual Phy change count value is reported.

To support topology discovery, all zoning expander ~~shall~~ may report EXPANDER CHANGE count according to the group assignment of the host. The Expander Change count maintains the number of BROADCAST(CHANGE) originated by the expander that is propagated to the group that the host is a member of. Logically, this requires the zoning expander to maintain a expander change count on per source group basis that counts the number of BROADCAST(CHANGE) originated from each source group.

~~Editor's note: The actual Expander CHANGE count may be calculated from the Phy Change counts of the expander. The expander does so by going through each Phy of the expander and sum up the Phy change counts of the Phys the SGID is permitted to access.~~

4.1.~~9~~6 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 Goup 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 the 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 1 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~~8~~ 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 VIOLIATION). 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								
12		SOURCE SAS ADDRESS						
19								
20	ACCESS ZONE MANAGE MENT	SOURCE GROUP ID (SGID)						
21	PATHWAY BLOCK COUNT							
22	(MSB)	ARBITRATION WAIT TIME						
23								(LSB)
24		MORE COMPATIBLE FEATURES						
27								
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 forwards 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						
5								(LSB)
6	(MSB)	EXPANDER ROUTE INDEXES						
7								(LSB)
8	Reserved	NUMBER OF ZONES						
9	NUMBER OF PHYS							
10	Reserved				SUPERVISING STATUS	SUPERVISING ALLOWED	CONFIGURING	CONFIGURABLE ROUTE TABLE
11	Reserved							
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20	Reserved							
27								
28	(MSB)	CRC						
31								(LSB)

The NUMBER OF ZONES field indicates the number of zones supported when the ZONING SUPPORTED bit is 1. For expanders that do not support zoning, 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 SUPERVISING ALLOWED indicates whether this expander is allowed to be a supervising expander. This is a zoning expander attribute configured by the supervisor through SMP CONFIGURE ZONE PERMISSION command.

The SUPERVISING STATUS returns the status of the supervising expander zone permission update propagation.

00: The current expander is not the supervising expander

01: 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 CONFIGURE ZONE PERMISSION command with FUNCTION RESULT code of FUNCTION ACCEPTED.

10: The current expander is the supervising expander and it has successfully propagated the zone permission update that was successfully received from a supervisor through CONFIGURE ZONE PERMISSION command with FUNCTION RESULT code of FUNCTION ACCEPTED.

11: 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 CONFIGURE_ZONE_PERMISSION command with 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.

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

…

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 7	Ignored							
8	Reserved							
9	PHY IDENTIFIER							
10	Ignored							
11	Reserved							
12	Ignored	ATTACHED DEVICE TYPE			Ignored			
13	Reserved				NEGOTIATED PHYSICAL LINK RATE			
14					ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15	ATTACHED SATA PORT SELECTOR	Reserved			ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA DEVICE
16 23	SAS ADDRESS							
24 31	ATTACHED SAS ADDRESS							
32		ATTACHED PHY IDENTIFIER						
33 39	Reserved							
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	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved				ZONE VIOLATION	~~SOURCE CHECK~~	TRUSTED	SUPERVISOR
49	Reserved	GROUP ID						
50 51	Vendor Specific							
52 55	(MSB) CRC (LSB)							

…

The ZONE VIOLATION field is set to 1 if any ZONE violation has occurred causing the specified phy to send OPEN_REJECT(ZONE VIOLIATION). 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. The Zoning expander shall send out BROADCAST(CHANGE) with SGID=127 to all zones after a CONFIGURE PHY ZONE function is executed.

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 requests

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	Ignored							
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 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 requests

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	Reserved							
5								
6	~~SET BATCH~~ Reserved	SOURCE GROUP ID						
7	NEW VALUE	TARGET GROUP ID						
8	Ignored							
9	OPERATION						PROPAGATE UPDATE	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 OPERATION field defines which of the following operations shall be executed

00: 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.

01: 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.

10: SET SUPERVISING ALLOWED: This sets the SUPERVISING ALLOWED attribute of the zoning expander, which allows the current expander to be a supervising expander.

11: CLEAR SUPERVISING ALLOWED: This clears the SUPERVISING ALLOWED attribute of the zoning expander, which disallow the current expander to be a supervising expander.

~~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 PROPAGATE UPDATE bit is set to ~~indicate that the Supervisor is handing the command off to this expander to become the Supervising expander.~~ 1 to indicate the UPDATE ZONE PERMISSION command is sent from a supervisor to the supervising expander of the zoning fabric. 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 fabric.

If a non-supervising expander receive CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE =0, the non-supervising expander shall updates its own permission table without propagation.

If a non-supervising expander receives CONFIGURE ZONE PERMISISON command with PROPAGATE UPDATE =1, the non-supervising expander shall ignore this command and return FUNCTION FAILED.

If a supervising expander receives CONFIGURE ZONE PERMISSION command with PROPAGATE UPDATE =1, the supervising expander shall update its permission table accordingly and propagate the same permission table updates to all other zoning expanders within the domain by sending SMP CONFIGURE ZONE PERMISSION commands to individual expanders with PROPGATE UPDATE =0. If the supervising expander receives another CONFIGURE ZONE PERMISSION command before the propagate of the zone permission update from the previous command has not been completed, the supervising expander shall ignore the new CONFIGURE ZONE PERMISSION command and return FUNCTION FAILED.

If a supervising expander receives CONFIGURE ZONE PERMISISON command with PROPAGATE UPDATE =0, the supervising expander shall ignore the command and return FUNCITON FAILED.

The UPDATE COMPLETE bit indicates whether the current CONFIGURE ZONE PERMISSION command is the last command of a sequence of CONFIGURE ZONE PERMISSION commands. This bit is valid in both batch set mode and single set mode. The zoning expander shall send out BROADCAST(CHANGE) messages after a CONFIGURE ZONE PERMISSION command with UPDATE COMPLETE bit is set to 1.

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				
11								(LSB)

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							CONFIG URING
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 10.4.3.2.

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						
7								(LSB)
8	~~GENERATION CODE~~RESERVED							
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						
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 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	Reserved						
1	~~Ignored~~SUPERVISING ALLOWED	ATTACHED DEVICE TYPE			Ignored		TRUSTED	SUPERVISOR
2	Ignored	GROUP ID						
3	Ignored							
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).

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and is defined in table 178.

The SUPERVISING ALLOWED indicates whether this expander designated by the routed SAS address is allowed to be a supervising expander.