

Date: 6 January 2006
To: T10 Technical Committee
From: Steve Johnson and Ralph Weber
Subject: SAS-2 Expander Configuration Supervisor and SMP CHANGE Request

Summary

Changes are proposed that:

- a) generalize the Supervising Device function proposed in 06-019, making it independent of SAS-2 Zoning and increasing the opportunities for self-configuring expander sets; and
- b) replace the ZONE BROADCAST address frame proposed in 06-019 with a new SMP Request, SMP CHANGE.

Referenced Documents

SAS-2 revision 1 (<http://www.t10.org/ftp/t10/drafts/sas2/sas2r01.pdf>)
SAS-2 Zoning [06-019r1] (<http://www.t10.org/ftp/t10/document.06/06-019r0.pdf>)

Revision History

r0 Initial revision

Overview

During the fall of 2005, issues were raised in the SAS Protocol Working Group regarding the general suitability of Expander Device implementations of the SAS-2 Discovery Process (see 4.7 in SAS-2 revision 1). While there was no question about Expander Devices discovering the portion of the SAS Domain tree below them, questions were raised about discovering the portion of the tree above the Expander Device, at least in the context of the normal operation for SAS Expander Devices.

It has since been realized that generalizing the Supervising Device function proposed as part of SAS-2 Zoning (see 06-019r1) would provide the perfect context for SAS Expander Devices whose product goals include performing the SAS Discovery Process. Based on this revelation, this proposal details the changes needed to promote the SAS-2 Zoning concept of Supervising Device from a zoning-specific feature to a general feature that any SAS-2 Device, including but not limited to Expander Devices, can implement.

Similarly, most of the recent SAS Protocol Working Group discussions about the 06-019r1 proposed ZONE BROADCAST address frame have been dominated by complaints about it.

- Flow Control
- Non-Deterministic Behavior
- Insufficient Information in Broadcast
 - resulting in extra message traffic to discover the rest of the story
- Frame Size
- Frame versus Primitive Signal usage

Every attempt to address these issues has produced greater complexity. This proposal eliminates the ZONE BROADCAST and replaces it with a SMP CHANGE request/response sent:

- from the device that detects a change in the SAS domain configuration,
- to the Supervising Expander

The SMP CHANGE request/response proposed in this document has many advantages.

- More Simple than any proposal to date
- Deterministic Behavior
- Confirmed Delivery
- Elimination of Superfluous Traffic
- Notification that is Coherent With Routing And Permission Tables

Changes Proposed in 06-019r1

3.1.4 active zone supervisor device: The supervisor device (see 3.1.264) that has been elected (see 4.7a.3) to configure the expander route tables of those devices with the CONFIGURE ROUTE TABLE bit set to one in their REPORT GENERAL response, all zoning expander zone route tables, and all zone permission tables in the ZSDS SAS service delivery subsystem.

...

3.1.263 zone-supervising supervisor device priority: The priority of a zone supervisor device for election to the roll of active zone supervisor device (see 3.1.4).

3.1.264 zone supervisor device: A zoning device An SMP target/initiator port that is capable of performing the duties of the active supervisor device (see 3.1.4) and so indicated by setting has its zone-supervising supervisor device priority (see 3.1.263) to a value that is greater than zero. See 4.7a.

...

4. General

...

4.9 Zoning

4.9.1 Zoning overview

...

The zone route table is an extended version a variation of the expander route table (See Section see 4.6.7.3) that is extended to include zone phy information. The zone permission table and zone route table are configured by the active zone supervisor device (see 3.1.4).

...

4.9.2 Zoning expander device model overview

A zoning expander device shall support the following attributes:

- a) contains an SMP initiator and SMP target (See SAS-2 revision 1 section 4.6.1);
- b) allows or denies connection requests following the zoning permissions;
- e) self-configuring;
- d) supports all zoning related SMP functions;
- e) contains a zone permission table that supports 128 zone groups; and
- f) contains a zone expander zone route table.

ROW Note 1: I believe the requirement for a supervisor device in the service delivery subsystem eliminates the need for the 'self-configuring' requirement, but I may need to be corrected on this point.

The proper configuration of zoning expander devices depends on the presence of at least one supervisor device (see 3.1.264) in the ZSDS. This dependency may be satisfied by including the supervisor device in one or more of the zoning expander devices in the ZSDS.

... {an unspecified number of 06-019r1 subclause(s) have been skipped} ...

4.9.3.4 Zoning Expander zone route table

~~Zoning expander devices shall be self-configuring. Zoning expander devices in the ZSDS cooperate to complete the zone route tables.~~ The active supervisor device (see 3.1.4) shall configure the zone route tables in each zoning expander device. The zone route table shall contain zone phy information in addition to the expander phy information.

Figure 2 shows a representation of a zoning expander zone route table.

{In figure 2, remove the ZONE SUPERVISING PRIORITY field from the box labeled Expander Zone Route Entry.}

... {an unspecified number of 06-019r1 subclause(s) have been skipped} ...

{The intention of this proposal is that the Supervisor Device functions proposed in 06-019r1 be moved out of the Zoning model subclause and into a new model subclause, 4.7a, that is inserted between 4.7 and 4.8. The notation in the following changes may not clearly show this intention.}

4.7a Supervisor devices

4.7a.1 Overview

Supervisor devices are an extension of the discovery process defined in 4.7.

Among the supervisor devices in a SAS domain, only one (i.e., the active supervisor device) performs the entire discovery process, configures each configurable expander device with the necessary expander route entries to allow routing of connection requests through the SAS domain, and configures each zoning expander device with zone expander route tables and zone permission tables.

Maintaining the limit of one active supervisor device completing the discovery process is accomplished via an election process (see 4.7a.3).

4.7a.2 Supervisor device requirements

Supervisor devices shall:

- a) set the SMP TARGET PORT bit and SMP INITIATOR PORT bit to one in their IDENTIFY address frame (see 7.8.2);
- b) support the SMP REPORT GENERAL function;
- c) set the ACCESS ZONE MANAGEMENT bit to one in the OPEN frame that precedes any SMP request;
- d) set the SUPERVISOR PRIORITY field in its SMP REPORT GENERAL response to a value that is greater than zero;
- e) be capable of configuring expander route tables, zone expander route tables and zone permission tables; and
- f) perform the discovery process defined in 4.7 with the additions described in 4.7a.3.2.

If a supervisor device determines that a change has occurred in its ability to communicate with a device other than the active supervisor device, it shall not originate a BROADCAST (CHANGE) primitive. Instead, it shall perform an

SMP CHANGE function (see 10.4.3.bb) with the active supervisor device indicating the nature of the configuration change that it has detected.

If a supervisor device determines that it is not able to communicate with the active supervisor device (see 3.1.4), it shall originate a BROADCAST (CHANGE SUPERVISOR) primitive (see 7.2.5.4).

Except during the election of an active supervisor, a supervisor device that receives a REPORT GENERAL in which the contents of the ACTIVE SUPERVISOR SAS ADDRESS field do not match its own ACTIVE SUPERVISOR SAS ADDRESS field contents shall originate a BROADCAST (CHANGE SUPERVISOR) primitive.

If a device that is not a supervisor device determines that a change has occurred in its ability to communicate with a device other than the active supervisor device, it should perform an SMP CHANGE function with the active supervisor device indicating the nature of the configuration change that it has detected, but it may originate a BROADCAST (CHANGE) primitive instead.

If a device that is not a supervisor device determines that is not able to communicate with the active supervisor device, it should originate a BROADCAST (CHANGE SUPERVISOR) primitive, but it may originate a BROADCAST (CHANGE) primitive instead.

Except when it is the active supervisor device, a supervisor device's response to a BROADCAST (CHANGE) primitive shall not include configuring expander route tables, configuring zone expander route tables or configuring zone permission tables.

ROW Note 2: The discussions of which I am aware (including apparently 06-019r1) have assumed that an active supervisor election is initiated by receipt of a BROADCAST (CHANGE). This assumption appears to result in an infinite loop.

The last step in an active supervisor election is changing the CONFIGURING bit value from one to zero. When this is done, SAS-2 10.4.3.3 requires the sending of a BROADCAST (CHANGE) primitive. The purpose of this action appears to be informing SAS initiator ports that it is time to begin discovering accessible devices in the SAS domain. For this reason, changing the SAS-2 behavior with respect to the CONFIGURING bit seems ill advised.

If the SAS-2-required BROADCAST (CHANGE) also initiates a new active supervisor election, then the process of electing an active supervisor will never end. From this arises the need to invent BROADCAST (CHANGE SUPERVISOR).

It is possible that the proposal to convert BROADCAST (RESERVED 1) to BROADCAST (CHANGE SUPERVISOR) will break existing silicon. If this proves to be the case, alternative approaches will need to be evaluated.

4.7a.3 Electing an active supervisor device

4.7a.3.1 Initiating an active supervisor election

A supervisor device (see 3.1.264) that experiences a power on or receives a BROADCAST (CHANGE SUPERVISOR) primitive shall:

- 1) set the CONFIGURING bit in its REPORT GENERAL response data (see 10.4.3.3) to one;
- 2) set the ACTIVE SUPERVISOR SAS ADDRESS field in its REPORT GENERAL response data to zero;
- 3) if the active supervisor election is being initiated in response to a power on, originate a BROADCAST (CHANGE SUPERVISOR) primitive (see 7.2.5.4); and

- 4) begin performing the active supervisor election process defined in 4.7a.3.2.

4.7a.3.2 Performing an active supervisor election

To elect an active supervisor device, each supervisor device shall traverse the SAS domain using the algorithm defined in 4.7.1 and while traversing the SAS domain each supervisor device shall:

- 1) send a REPORT GENERAL request to any port that has the SMP TARGET PORT bit set to one in its IDENTIFY address frame;
- 2) note whether the CONFIGURE ROUTE TABLE bit is set to one in the REPORT GENERAL response;
- 3) note whether the ZONING bit is set to one in the REPORT GENERAL response; and
- 4) stop performing the SAS domain traversal if:
 - 1) the SUPERVISOR PRIORITY field in the REPORT GENERAL response contains a value that is greater than the value that the supervisor device places in its own REPORT GENERAL responses;
 - 2) the SUPERVISOR PRIORITY field in the REPORT GENERAL response contains a value that is equal to the value that the supervisor device places in its own REPORT GENERAL responses and the value in the VENDOR-SPECIFIC IDENTIFIER field in the SAS address (see 4.2.2) to which the REPORT GENERAL request was sent is greater than the value in the VENDOR-SPECIFIC IDENTIFIER field of the supervisor device; or
 - 3) the SUPERVISOR PRIORITY field in the REPORT GENERAL response contains a value that is equal to the value that the supervisor device places in its own REPORT GENERAL responses, the value in the VENDOR-SPECIFIC IDENTIFIER field in the SAS address (see 4.2.2) to which the REPORT GENERAL request was sent is equal the value in the VENDOR-SPECIFIC IDENTIFIER field of the supervisor device, and the value in the IEEE COMPANY ID field in the SAS address (see 4.2.2) to which the REPORT GENERAL request was sent is greater than the value in the IEEE COMPANY ID field of the supervisor device.

The supervisor device that completes the traversal of the entire SAS domain without stopping (i.e., without encountering a supervisor priority and SAS address combination requiring that the traversal be stopped as described in list item 4) in this subclause) is the active supervisor device (see 3.1.4). The active supervisor device shall complete the election process as defined in 4.7a.3.3.

During the election of an active supervisor device, any supervisor device that detects a change communications accessibility with another supervisor device shall originate a BROADCAST (CHANGE SUPERVISOR) primitive.

Whether the SAS domain traversal has been stopped or not, the receipt of a BROADCAST (CHANGE SUPERVISOR) or BROADCAST (CHANGE) primitive (see 7.2.5.4) shall cause the supervisor device to restart the SAS domain traversal from the beginning.

4.7a.3.3 Completing an active supervisor election

A supervisor device that completes the entire SAS domain traversal defined in 4.7a.3.2 shall:

- 1) if any REPORT GENERAL response had the CONFIGURE ROUTE TABLE bit set to one, then route table information shall be collected and sent to each expander device with a configurable route table using the SMP CONFIGURE ROUTE INFORMATION function (see 4.6.7.3);
- 2) if any REPORT GENERAL response had the ZONING bit set to one, then zone route expander table information and zone permission table information shall be collected and sent to each zoning expander device using the SMP CONFIGURE PHY ZONE function and SMP CONFIGURE ZONE PERMISSION function, respectively (see 4.9.6);
- 3) perform an SMP ACTIVE SUPERVISOR ELECTED function (see 10.4.3.aa) with each device that reported a non-zero SUPERVISOR PRIORITY field in its REPORT GENERAL response and each device that reported a ZONING bit set to one in its REPORT GENERAL response;

- 4) set the ACTIVE SUPERVISOR SAS ADDRESS field in its REPORT GENERAL response data (see 10.4.3.3) to its SAS address;
- 5) set the CONFIGURING bit in its REPORT GENERAL response data to zero; and
- 6) originate a BROADCAST (CHANGE) primitive.

ROW Note 3: If a SAS domain contains zoning devices but no supervisor devices, no active supervisor device is elected and no zoning configuration SMP requests are sent. As a result, the zones remain in whatever state they were in before the discovery process was initiated.

Upon receipt of an SMP ACTIVE SUPERVISOR ELECTED request, each supervisor device that stopped traversing the SAS domain during the election process (see 4.7a.3.2) shall:

- a) copy the ACTIVE SUPERVISOR SAS ADDRESS field the ACTIVE SUPERVISOR ELECTED request to the ACTIVE SUPERVISOR SAS ADDRESS field in its REPORT GENERAL response data (see 10.4.3.3);
- b) set the CONFIGURING bit in its REPORT GENERAL response data to zero; and
- c) not originate a BROADCAST (CHANGE) primitive.

ROW Note 4: One BROADCAST (CHANGE) sent by the active supervisor device should be enough, right?

4.7a.4 Active supervisor device duties between elections

After the single active supervisor device in a SAS domain has been determined as described in 4.7a.3, the active supervisor device shall:

- a) maintain current and complete information regarding the SAS domain configuration;
- b) update expander route tables, zone expander route tables and zone permission tables based on:
 - A) information received in SMP CHANGE requests (see 10.4.3.bb); and
 - B) information discovered by polling devices in the SAS domain;
- c) poll for configuration changes as necessary in response to BROADCAST (CHANGE) primitives received;
- d) use the SMP ACTIVE SUPERVISOR ELECTED function (see 10.4.3.aa) as necessary to update the active supervisor SAS address in devices;
- e) originate BROADCAST (CHANGE) primitives as needed to start new device polling by SAS initiators; and
- f) restart the active supervisor election process in response to receipt of a BROADCAST (CHANGE SUPERVISOR) primitive.

{remove all of 4.9.4 - including table 3 - from 06-019r1. Text shown here for reference}

4.9.4 Zone supervisor devices

~~... 4.9.4.1 Zone supervisor priority~~

~~When a SAS device is attached to a phy that reports a ZONE SUPERVISING PRIORITY field value greater than zero, the device is eligible for active zone supervisor device.~~

~~If a phy is attached to a device that is not a zone supervisor device then the ZONE SUPERVISING PRIORITY field shall be set to zero.~~

~~... 4.9.4.2 Active zone supervisor~~

~~There is one active zone supervisor device in a ZSDS. The zoning expander devices elect a device to be the active zone supervisor device based on the value of its ZONE SUPERVISING PRIORITY field. The highest value of all the zone~~

~~supervising priority entries shall be elected to be the active zone supervisor device. If there are two or more supervisor priority entries with the same zone supervising priority value then the device with the highest SAS address shall be elected to be the active zone supervisor device.~~

~~... 4.9.4.3 Active zone supervisor device election~~

~~The active zone supervisor device election occurs after all zoning expander devices in the ZSDS complete the self-configuration process. The completion of the self-configuration process is identified when the zoning expander device sets the CONFIGURING bit from one to zero in the SMP REPORT GENERAL response the sends BROADCAST (CHANGE), (see section 4.7.1).~~

~~A zoning device shall issue SMP DISCOVER request to each phy in the ZSDS and evaluate the highest ZONE SUPERVISING PRIORITY field value (see Table 3). If the same highest value is reported by two or more phys or devices, then the highest SAS address is elected.~~

~~A zoning expander device shall set the ACTIVE ZONE SUPERVISOR PRIORITY and ACTIVE ZONE SUPERVISOR SAS ADDRESS fields to zero until the active zone supervisor election process is complete.~~

~~Table 3 indicates the significance of priority values in the ZONE SUPERVISING PRIORITY field.~~

{remove table 3}

~~If all zone table entries have the zone supervising priority set to zero then the election process fails and the ACTIVE ZONE SUPERVISOR PRIORITY field and the ACTIVE ZONE SUPERVISOR SAS ADDRESS field is set to zero in the SMP REPORT GENERAL response frame.~~

... 4.9.5 Access zone management

A **zone** supervisor device may make SMP REPORT ZONE PERMISSION and SMP REPORT ZONE ROUTE TABLE requests

~~A zoning device that reports a ZONE SUPERVISING PRIORITY field value greater than zero, shall set the ACCESS ZONE MANAGEMENT bit set to one in the OPEN address frame. (See section 7.8.3).~~

A zoning device that reports a **ZONE SUPERVISING SUPERVISOR** PRIORITY field of zero, shall set the ACCESS ZONE MANAGEMENT bit set to zero in the OPEN address frame. (See section 7.8.3).

When an OPEN address frame is received on a phy that has the **ZONE SUPERVISING SUPERVISOR** PRIORITY field set to a value greater than zero and the ZONE PARTICIPATING bit set to zero, then the zoning expander device shall respond **for as if** the ACCESS ZONE MANAGEMENT bit **is** set to one in the OPEN address frame **and ignore the value of the ACCESS ZONE MANAGEMENT bit in the OPEN address frame.** (See section 7.8.3 Table 10). If the OPEN address frame is forwarded to another device then the zoning expander device shall set the ACCESS ZONE MANAGEMENT bit to one.

... 4.9.6 Zone configuration

SMP CONFIGURE ZONE PERMISSION and SMP CONFIGURE PHY ZONE requests are used to manage the ZSDS. These requests **if shall be** originated **from by** the active **zone** supervisor device. If **an SMP CONFIGURE ZONE PERMISSION or SMP CONFIGURE PHY ZONE request the request is made by is received from** a device that is not the active **zone** supervisor device, then the response shall be as defined in section 7.8.3 Table 10.

The SMP CONFIGURE ZONE PERMISSION function requests change to entries in the zone permission table. (See section 10.4.3.14). The active **zone** supervisor device shall send an SMP CONFIGURE ZONE PERMISSION request to **all each** zoning expanders in the ZSDS.

The SMP CONFIGURE PHY ZONE function requests change to the zoning attributes for phys in a zoning expander device. (See section 10.4.3.13). **The active supervisor device shall send an SMP CONFIGURE PHY ZONE request to each zoning expander in the ZSDS.**

... 4.9.6 Zoning expander phy reset event {no changes to the 06-019r1 proposal in this subclause}

~~... 4.9.7 ZONE BROADCAST address frames ...~~ {remove all of this subclause from the 06-019r1 proposal}

... {two 06-019r1 subclause(s) have been skipped} ...

... 7.8 Address frames

~~... 7.8.1 Address frames overview~~ {remove/undo all the changes shown in the 06-019r1 proposal}

... 7.8.2 IDENTIFY address frame

Table 92 — ~~Table 7~~ IDENTIFY address frame format

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
21	Reserved				ZONE DEVICE	ZONE BROADCAST METHOD Reserved		
	...							

...

A ZONE DEVICE bit set to one specifies that the device is a zoning device. A ZONE DEVICE bit set to zero specifies that the device does not support zone functions.

~~The ZONE BROADCAST METHOD field specifies the type of ZONE BROADCAST address frames supported by the device. (See Table 8)~~

...

{per 06-019r1 [Revision of existing text shown in red]}

See 4.1.3 for additional requirements concerning the DEVICE TYPE field, SSP INITIATOR PORT bit, STP INITIATOR PORT bit, SMP INITIATOR PORT bit, SSP TARGET PORT bit, STP TARGET PORT bit, SMP TARGET PORT bit, **and** SAS ADDRESS field, **and** ZONE DEVICE bit **and** ~~ZONE BROADCAST METHOD~~ field.

{per 06-019r1 [End of revision of existing text shown in red]}

{remove table 8 from the 06-019r1 proposal}

... 7.8.3 OPEN address frame

...

{no changes to the 06-019r1 proposal vis a vis table 9 aka table 94 in SAS-2 revision 1}

...

An ACCESS ZONE MANAGEMENT bit set to one specifies that the OPEN request is from a **zone supervisor device** and shall respond as defined in Table 10. If the ACCESS ZONE MANAGEMENT bit is set to zero then all zone type SMP functions shall return a FUNCTION RESULT of UNKNOWN SMP FUNCTION.

ROW Note 5: In most cases, 'zone supervisor device' has been changed to 'supervisor device'. The appropriateness of such a change is unclear in this case.

{no other changes or issues in 7.8.3 of the 06-019r1 proposal}

~~... 7.8.4 ZONE BROADCAST type address frame ...~~ {remove all of this subclause from the 06-019r1 proposal}

~~... 7.8.5 Short ZONE BROADCAST type address frame ...~~
 {remove all of this subclause from the 06-019r1 proposal}

10 Application layer...

... 10.4.3 SMP functions

... 10.4.3.2 SMP function response format {no changes to the 06-019r1 proposal in this subclause}

10.4.3.3 REPORT GENERAL function

...

Table 193 ~~Table 14~~ defines the response format.

Table 193 — ~~Table 14~~ — REPORT GENERAL response

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
20	Reserved					ZONE DEVICE	CONFIGUR- ING	CONFIG- URABLE ROUTE TABLE
	...							

...

The **ZONE DEVICE** bit shall be set to one if the device supports zoning (see 4.9). If the device does not support zoning, the **ZONE DEVICE** bit shall be set to zero.

{The following paragraph does not appear in 06-019r1, but rather than listing changes to 10.4.3.3 in two places in this proposal, the paragraph has been copied from SAS-2 revision 1 to here.}

A CONFIGURING bit set to one indicates that a self-configuring expander device has not completed configuring its expander route table **or that an election is in progress to determine the active supervisor device (see 4.7a.3)**. A CONFIGURING bit set to zero indicates that configuration is complete and the expander device is ready for connection requests. Changes in this bit from one to zero result in a BROADCAST (CHANGE) being originated, **unless the change occurs in a supervisor device that has not been elected as the active supervisor device (see**

4.7a.3.3). SMP target ports in self-configuring expander devices shall support this bit. SMP target ports in configurable expander devices and in other device types shall set the CONFIGURING bit to zero.

...

~~The ACTIVE ZONE SUPERVISOR PRIORITY field indicates the zone supervising priority of the elected active zone supervisor device. This field shall be set to zero if there is no active zone supervisor device.~~

The ZONE SUPERVISING SUPERVISOR PRIORITY field indicates the supervising supervisor priority of the responding zoning expander supervisor device. If the responding device is not a supervisor device, the SUPERVISOR PRIORITY field shall be set to zero.

...

The ACTIVE ZONE SUPERVISOR SAS ADDRESS field indicates the SAS address of the elected active zone supervisor device (see 3.1.4). ~~This~~ The ACTIVE SUPERVISOR SAS ADDRESS field shall be set to zero if: ~~there is no active zone supervisor device.~~

- a) The SAS address of the active supervisor device is not known; or
- b) An election is in progress to determine the active supervisor device (see 4.7a.3).

... 10.4.3.5 DISCOVER function

...{no changes to the 06-019r1 proposed changes in the DISCOVER request definition}...

Table 197 ~~Table 16~~ defines the response format.

Table 197 — ~~Table 16~~ — DISCOVER response

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
33	Reserved				ATTACHED ZONE DEVICE	ATTACHED ZONE BROADCAST METHOD Reserved		
	...							
48	Reserved		ZONE VIOLATION	ZONE PAR- TICIPATING	ZONE SUPERVISING PRIORITY Reserved			
49	Reserved	ZONE GROUP						
	...							

...

The ATTACHED ZONE DEVICE bit indicates the value of the ZONE DEVICE bit received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

~~The ATTACHED ZONE BROADCAST METHOD field indicates the value of the ZONE BROADCAST METHOD field (See Table 8) received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.~~

A ZONE VIOLATION bit set to one indicates that the phy sent an OPEN_REJECT (ZONE VIOLATION) response to a prior DISCOVER command. (See section 7.2.5.11)

A ZONE PARTICIPATING bit set to one indicates that the phy is attached to another zoning device.

~~The ZONE SUPERVISING PRIORITY field indicates the zone supervisor device election priority of the phy.~~

The ZONE GROUP field indicates the zone group that contains the phy.

The CRC field is defined in 7.8.1.

...

10.4.3.13 CONFIGURE PHY ZONE function

The CONFIGURE PHY ZONE function sets the zoning attributes for phys in the zoning expander device. This function shall be supported by all zoning devices and shall only be issued by the active ~~zone~~ supervisor device. The ACCESS ZONE MANAGEMENT bit in the OPEN request shall be used to determine the function result as shown in section 7.8.3 Table 10.

... {several paragraphs and tables in 06-019r1 skipped because no changes are needed} ...

Table 18 — Phy zone configuration entry descriptor

Byte/ Bit	7	6	5	4	3	2	1	0
0	Reserved			ZONE PAR- TICIPATING	ZONE SUPERVISING PRIORITY Reserved			
1	Reserved	ZONE GROUP						

A ZONE PARTICIPATING bit set to one indicates that the phy is attached to another zoning device. A ZONE PARTICIPATING bit set to zero indicates that the phy is not attached to a zoning device.

~~The ZONE SUPERVISING PRIORITY field is defined in section 10.4.3.5.~~

The ZONE GROUP field is defined in section 10.4.3.5.

{no other changes in 10.4.3.13}

10.4.3.14 CONFIGURE ZONE PERMISSION function

The CONFIGURE ZONE PERMISSION function requests a change to the zone permissions table entries ~~and the zoning expander device zone supervising priority~~. This function shall be supported by all zoning devices and shall only be issued by the active ~~zone~~ supervisor device. The value of the ACCESS ZONE MANAGEMENT bit in the OPEN request shall be used to determine the function result as shown in section 7.8.3 Table 10.

- a) A single entry update that uses the SOURCE ZONE GROUP and the TARGET ZONE GROUP fields to identify the table entry bit to be written. The BATCH bit shall be set to zero; and
- b) A batch mode that allows multiple table row entries of a zone group field to be updated. In this mode the BATCH bit shall be set to one.

A request to change zone permission to zone group 0 or zone group 1 shall return a function result of FUNCTION FAILED in the response frame.

Table 20 defines the CONFIGURE ZONE PERMISSION request format.

Table 20 — CONFIGURE ZONE PERMISSION request

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
9	ZONE SUPERVISING PRIORITY ^b Reserved			Reserved		UPDATE PRIORITY Reserved	BATCH	UPDATE COM- PLETE ^c
	...							
<p>^a These fields are ignored if the BATCH bit is set to one</p> <p>^b This field is ignored if the UPDATE PRIORITY bit is set to zero</p> <p>^c These fields are ignored if the BATCH bit is set to zero</p>								

... {several paragraphs in 06-019r1 skipped because no changes are needed} ...

The GROUP PERMISSION bit specifies the new access permission to be updated between the two zone groups specified in the SOURCE ZONE GROUP and TARGET ZONE GROUP fields. The TARGET ZONE GROUP field specifies the target zone group of the zone permission entry.

~~The ZONE SUPERVISING PRIORITY field specifies the new zone supervisor device election priority to be written to the zoning device.~~

~~When the UPDATE PRIORITY bit is set to one the device shall update its zone supervising priority to the value contained in the ZONE SUPERVISING PRIORITY field.~~

A BATCH bit set to zero specifies that the CONFIGURE ZONE PERMISSION request is being used in single entry update mode. A BATCH bit set to one specifies that the CONFIGURE ZONE PERMISSION request is being used in batch mode, allowing updates to multiple access permission entries.

{no other changes in 10.4.3.14}

10.4.3.15 REPORT ZONE PERMISSION function

The REPORT ZONE PERMISSION function returns a set of zone permission table entries. This function shall be supported by all zoning devices and shall only be issued by ~~zone~~ supervisor devices. The value of the ACCESS ZONE MANAGEMENT bit in the OPEN request shall be used to determine the function result as shown in section 7.8.3 Table 10.

{no other changes in 10.4.3.15}

10.4.3.16 REPORT ZONE ROUTE TABLE function

The REPORT ZONE ROUTE TABLE request requires a zoning expander device to respond with zone information about each phy. This function shall be supported by all zoning devices and shall only be issued by ~~zone~~ supervisor devices. The value of the ACCESS ZONE MANAGEMENT bit in the OPEN request shall be used to determine the function result as shown in section 7.8.3 Table 10.

... {several paragraphs and tables in 06-019r1 skipped because no changes are needed} ...

Table 27 — Zone route table entry descriptor

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
1	Reserved			ZONE PAR- TICIPATING	ZONE SUPERVISING PRIORITY Reserved			
	...							

... {several paragraphs in 06-019r1 skipped because no changes are needed} ...

The ZONE PARTICIPATING bit indicates that the phy is attached to another zoning device.

~~The ZONE SUPERVISING PRIORITY field is defined in section 10.4.3.5.~~

The ZONE GROUP field is defined in section 10.4.3.5.

{no other changes in 10.4.3.15}

Changes Proposed in SAS-2 revision 1

4.6.6.5 BPP interface

Table 15 describes the requests from an expander phy to the BPP. See 7.11 for more information on broadcasts.

Table 15 — Expander phy to BPP requests

Message	Description
...	
Broadcast Event Notify (RESERVED CHANGE SUPERVISOR Received)	Request to transmit a BROADCAST (RESERVED CHANGE SUPERVISOR) on all other ports because a BROADCAST (RESERVED CHANGE SUPERVISOR) was received.

{no other changes in 4.6.6.5}

7.2.2 Primitive summary

Table 70 defines the primitives not specific to the type of connection.

Table 70 — Primitives not specific to type of connection ...

Primitive	Use ^a	From ^b			To ^b			Primitive sequence type ^c
		I	E	T	I	E	T	
...								
BROADCAST (CHANGE)	NoConn	I	E		I	E	T	Redundant
BROADCAST (SES)	NoConn			T	I	E	T	Redundant
BROADCAST (RESERVED CHANGE SUPERVISOR)	NoConn	I	E		I	E	T	Redundant
...								

{no other changes in 7.2.2}

7.2.3 Primitive encodings

Table 73 defines the primitive encoding for primitives not specific to type of connection.

Table 73 — Primitive encoding for primitives not specific to type of connection ...

Primitive	Character			
	1 st	2 nd	3 rd	4 th (last)
...				
BROADCAST (RESERVED CHANGE SUPERVISOR)	K28.5	D04.7	D01.4	D24.0
...				

{no other changes in 7.2.3}

7.2.5.4 BROADCAST

BROADCASTs are used to notify all SAS ports in a domain of an event.

The versions of BROADCAST representing different reasons are defined in table 79.

Table 79 — BROADCAST primitives

Primitive	Description
	...
BROADCAST (RESERVED 1 CHANGE SUPERVISOR)	Reserved. Notification of need to start active supervisor device election.
	...

{Several paragraphs skipped}

BROADCAST (SES) is sent by a SAS target port to notify SAS initiator ports that an asynchronous event has occurred in an enclosure, and SSP initiator ports should poll all the logical units with peripheral device types set to 0Dh (i.e., enclosure services devices) (see SPC-3 and SES-2) in the SAS domain. BROADCAST (SES) shall be ignored by SAS target ports.

BROADCAST (CHANGE SUPERVISOR) is sent by an expander device or a SAS target/initiator port to notify other expander devices and SAS target/initiator ports that a supervisor device election (see 4.7a.3) is needed to elect a new active supervisor device (see 3.1.4) (e.g., because the last known active supervisor device is not reachable).

BROADCAST (RESERVED CHANGE 0) and BROADCAST (RESERVED CHANGE 1) shall be processed the same as BROADCAST (CHANGE) by SAS ports. ~~BROADCAST (RESERVED 1)~~, BROADCAST (RESERVED 2), BROADCAST (RESERVED 3), and BROADCAST (RESERVED 4) shall be ignored by SAS ports.

{no other changes in 7.2.5.4}

10.4.3 SMP functions

10.4.3.1 SMP function request frame format

...

The FUNCTION field specifies which SMP function is being requested and is defined in table 189. If the value in the FUNCTION field is not supported by the SMP target port, it shall return a function result of UNKNOWN SMP FUNCTION as described in table 191.

{some of the additions in this table should appear in 06-019r1 but do not}

Table 189 — SMP functions (FUNCTION field)

Code	SMP Function	Description	Reference
...			
02h	READ GPIO REGISTER	See SFF-8485	
03h	REPORT ZONE PERMISSION	Return zone permission table entries	10.4.3.15
034h - 0Fh	Reserved for general SMP input functions		
...			
13h	REPORT ROUTE INFORMATION	Return route table information	10.4.3.8
14h	REPORT ZONE ROUTE TABLE	Return zone information about each phy	10.4.3.16
145h - 1Fh	Reserved for phy-based SMP input functions		
...			
81h	CHANGE	Reserved for a general SMP output function Report specific configuration changes	10.4.3.bb
82h	WRITE GPIO REGISTER	See SFF-8485	
83h	CONFIGURE ZONE PERMISSION	Change zone permissions table	10.4.3.14
84h	ACTIVE SUPERVISOR ELECTED	Report successful election of active supervisor device	10.4.3.aa
835h - 8Fh	Reserved for general SMP output functions		
...			
92h	PHY TEST FUNCTION	Request a test function by the specified phy	10.4.3.12
93h	CONFIGURE ZONE PHY	Change phy entries within a zone route table	10.4.3.13
923h - 9Fh	Reserved for phy-based SMP output functions		
...			

10.4.3.9 CONFIGURE GENERAL function

...

Table 210 — CONFIGURE GENERAL request

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
9	SUPERVISOR PRIORITY				UPDATE SUPERVI- SOR PRI- ORITY	UPDATE STP SMP L_T NEXUS LOSS TIME	UPDATE STP MAXI- MUM CON- NECT TIME LIMIT	UPDATE STP BUS INACTIVITY TIME LIMIT
	...							

...

If the UPDATE SUPERVISOR PRIORITY bit is set to one, the SUPERVISOR PRIORITY field specifies the new supervisor device priority to be used by the supervisor device (see 4.7a). Otherwise, the contents of the SUPERVISOR PRIORITY field are ignored.

If the UPDATE SUPERVISOR PRIORITY bit is set to one the device shall change its supervisor priority to the value contained in the SUPERVISOR PRIORITY field. If the UPDATE SUPERVISOR PRIORITY bit is set to zero, device shall not change its supervisor priority and shall ignore contents of the SUPERVISOR PRIORITY field.

{no other changes in 10.4.3.9}

10.4.3.aa ACTIVE SUPERVISOR ELECTED function

The ACTIVE SUPERVISOR ELECTED function is requested by the elected active supervisor device (see 4.7a.3.3) to inform other supervisor devices and zoning expander devices in the SAS domain of the election's completion.

Table aa1 defines the request format.

Table aa1 — ACTIVE SUPERVISOR ELECTED request

Byte/ Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (84h)								
2	Reserved								
3	REQUEST LENGTH (02h)								
4	ACTIVE SUPERVISOR SAS ADDRESS								
11	ACTIVE SUPERVISOR SAS ADDRESS								
12	(MSB)	CRC							
15								(LSB)	

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

The FUNCTION field shall be set to 84h.

The REQUEST LENGTH field shall be set to 02h.

The ACTIVE SUPERVISOR SAS ADDRESS field shall be set to the SAS address of the active supervisor device (i.e., to the SAS address of the device that is sending the request).

The CRC field is defined in 10.4.3.1.

Table aa2 defines the response format.

Table aa2 — ACTIVE SUPERVISOR ELECTED response

Byte/ Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (84h)								
2	FUNCTION RESULT								
3	REQUEST LENGTH (00h)								
4	(MSB)	CRC							
7								(LSB)	

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

The FUNCTION field shall be set to 84h.

The REQUEST LENGTH field shall be set to 00h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The CRC field is defined in 10.4.3.1.

10.4.3.bb CHANGE function

The CHANGE function is a request that is sent to the elected active supervisor device (see 4.7a.3.3) that describes change in the SAS domain that another device has detected.

Table bb1 defines the request format.

Table bb1 — CHANGE request

Byte/ Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (81h)							
2	Reserved							
3	REQUEST LENGTH ((n - 7) / 4)							
4	SAS ADDRESS							
11	SAS ADDRESS							
12	Reserved							
13	Reserved							
14	(MSB)	EXPANDER EXCHANGE COUNT						(LSB)
15	EXPANDER EXCHANGE COUNT							
16	(MSB)	EXPANDER ROUTE INDEXES						(LSB)
17	EXPANDER ROUTE INDEXES							
18	Reserved							
19	NUMBER OF PHYS							
20	Reserved							
21	Reserved							
22	Reserved							
23	PHY DESCRIPTORS LENGTH ((n - 28) / 4)							
24	first phy descriptor							
	⋮							
	last phy descriptor							
n-4	last phy descriptor							
n-3	(MSB)	CRC						(LSB)
n	CRC							

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

The FUNCTION field shall be set to 81h.

The REQUEST LENGTH field specifies the number of dwords that follow not including the CRC field.

The SAS ADDRESS field shall contain the SAS address of the expander device about which the change is being reported.

The EXPANDER EXCHANGE COUNT field shall contain the number of BROADCAST (CHANGE)s originated by the expander device about which the change is being reported as determined from the most recently available REPORT GENERAL response (see 10.4.3.3) from that expander device.

The EXPANDER ROUTE INDEXES field shall contain the maximum number of route indexes per phy for the expander device about which the change is being reported as determined from the most recently available REPORT GENERAL response from that expander device.

The NUMBER OF PHYS field the number of phys in the expander device about which the change is being reported as determined from the most recently available REPORT GENERAL response from that expander device, including any virtual phys and any vacant phys.

The PHY DESCRIPTORS LENGTH field specifies the number of dwords of phy descriptors that follow not including the CRC field.

Each phy descriptor shall have the format shown in table bb2.

Table bb2 — Expander phy descriptor format

Byte/ Bit	7	6	5	4	3	2	1	0
0	VIRTUAL PHY	Reserved			DESCRIPTOR LENGTH (4h)			
1	PHY IDENTIFIER							
2	PHY CHANGE COUNT							
3	Reserved				ROUTING ATTRIBUTE			
4	ATTACHED SAS ADDRESS							
11	ATTACHED SAS ADDRESS							
12	ATTACHED PHY IDENTIFIER							
13	Reserved				ATTACHED DEVICE TYPE			
14	Reserved							
15	Reserved							

A VIRTUAL PHY bit set to one indicates the phy is part of an internal port and the attached device is contained within the expander device. A VIRTUAL PHY bit set to zero indicates the phy is a physical phy and the attached device is not contained within the expander device.

The DESCRIPTOR LENGTH field specifies the number of dwords in the descriptor and shall be set to four.

The PHY IDENTIFIER field specifies the phy to which the descriptor applies.

The PHY CHANGE COUNT field the number of configuration changes originated by the phy to which the descriptor applies.

The ROUTING ATTRIBUTE field specifies the routing attribute supported by the phy and shall be set to the value returned by phy in response to a DISCOVER function (see 10.4.3.5).

The ATTACHED SAS ADDRESS field contains the value of the SAS ADDRESS field received in the IDENTIFY address frame during the identification sequence and shall be set to the value returned by phy in response to a DISCOVER function.

The ATTACHED PHY IDENTIFIER field contains a phy identifier for the attached phy and shall be set to the value returned by phy in response to a DISCOVER function.

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and shall be set to the value returned by phy in response to a DISCOVER function.

The CHANGE request CRC field is defined in 10.4.3.1.

Table bb3 defines the response format.

Table bb3 — CHANGE response

Byte/ Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (81h)								
2	FUNCTION RESULT								
3	REQUEST LENGTH (00h)								
4	(MSB)	CRC							
7								(LSB)	

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

The FUNCTION field shall be set to 81h.

The REQUEST LENGTH field shall be set to 00h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The CRC field is defined in 10.4.3.1.

Annex J

...

Table J.1 — Primitives with Hamming distance of 8 ...

1 st	2 nd	3 rd	4 th	Assignment
...				
K28.5	D04.7	D01.4	D24.0	BROADCAST (RESERVED CHANGE SUPERVISOR)
...				

{no other changes in annex J}