

Date: 23 January 2006
 To: T10 Technical Committee
 From: Steve Johnson and Ralph Weber
 Subject: SAS-2 Expander Configuration Supervisor

Summary

Changes are proposed that generalize the Supervising Device function proposed in 06-019, making it independent of SAS-2 Zoning and increasing the opportunities for expander set designers. Full functioning of the proposed changes are dependent on approval of the new SMP functions defined in 06-082.

Referenced Documents

SAS-2 revision 2	(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)
Zoning SMP Functions [06-082r0]	(http://www.t10.org/ftp/t10/document.06/06-082r0.pdf)

Revision History

- r0 Initial revision
- r1 Revised in response to discussions in the SAS Protocol Working Group (minutes in 06-038).
See the notes below on page 1.
- r2 Revised in response to discussions in the SAS Zoning conference call on 19 January 2006, including splitting the SMP functions needed by either Zoning proposal into a separate document, 06-028. See page 2 for additional notes regarding changes in r2.

Changes between r1 and r2 are indicated by change bars.

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 2). 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.

Summary of Substantial Issues from January, 2006 SAS Protocol Working Group

N.B. Those issues raised during the January, 2006 SAS Protocol Working Group that can only be addressed by withdrawing this proposal are not discussed or addressed in this section. This section describes substantial issues from the working group and how this r1 of this proposal responds them.

- **No mechanism for limiting BROADCAST (CHANGE) notifications to a single zone group:** The SMP function SEND BROADCAST (CHANGE) has been added and the active supervisor has been instructed to use the function when BROADCAST (CHANGE) primitives should be sent only to the phys in a specific zone group.
- **Incomplete integration of active supervisor configuration with SAS-1.1 configuration:** The SMP REPORT GENERAL bits indicating active supervisor configuration and SAS-1.1 self-configuration have been separated so that the two cases are easily identified. Specific active supervisor requirements have been added

with the goal of maintaining backwards compatibility with SAS-1.1 behavior. Only careful expert review will verify that the goal has been reached. Any errors found will be corrected in r2 or future revisions.

- **Replacing BROADCAST (CHANGE SUPERVISOR) with BROADCAST (CHANGE):** In the working group, it was noted that an infinite loop of supervisor elections could be broken if an election that does not result in a change of active supervisor does not produce another BROADCAST (CHANGE). This idea was considered during the preparation of r1 and omitted because backwards compatibility with SAS-1.1 could not be exhaustively verified. The change can be made in a future revision if a consensus can be reached regarding reliable compatibility.

Because BROADCAST (RESERVED 1) is used in a recently approved proposal, r1 proposes using BROADCAST (RESERVED 2) for BROADCAST (CHANGE SUPERVISOR).

Notes regarding other r1 topics

The ACTIVE SUPERVISOR ELECTED function was changed to ACTIVE SUPERVISOR UPDATE because it is used on occasions other than the completion of an active supervisor election.

Several editorial changes have been made in r1, the most substantial of which is showing all the REPORT GENERAL fields affected by this proposal in the REPORT GENERAL format table.

Summary of Substantial Issues from January 19, 2006 SAS Zoning conference call

N.B. Those issues that can only be addressed by withdrawing this proposal are not discussed or addressed in this section. This section describes substantial issues from the conference call and how this r2 of this proposal responds them.

- **Non-Supervisor Configuration Rules Are Incomplete:** The 4.7a.5 (Configuration changes in devices that are not supervisor devices) subclause totally ignored devices that are not expander devices and provided (at best) ambiguous backward compatibility with SAS-1.1. The subclause has been divided to cover expander devices (see 4.7a.5.1) and end devices (see 4.7a.5.2) and care has been taken to ensure that SAS-1.1 behavior has been covered (which should not be taken as a guarantee that the coverage is complete ... yet).
- **Limiting BROADCAST Primitives to Zone Groups Not Fully Specified:** New requirements were added (see 4.7a.6) to have supervisors intercept BROADCAST primitives and convert them to SMP CHANGE requests that cause the active supervisor to use SMP SEND BROADCAST PRIMITIVES request(s) to limit retransmission of the BROADCAST primitives to the zone groups from which they were received.

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 ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one or 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 domain, and limit retransmission of BROADCAST primitives to the zone groups in which they are originated.

...

3.1.263 zone-supervising supervisor device priority: The priority of a zone supervisor device for election to the role 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 2 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 SAS domain 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 every configurable expander device with the necessary expander route entries to allow complete routing of connection requests through the SAS domain, and configures every 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 their 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;
- f) perform the discovery process defined in 4.7 with the additions described in 4.7a.3.2; and
- g) act to limit distribution of BROADCAST primitives to zone groups as described in 4.7a.6.

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 *in 06-082r0*) 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, if supervisor device 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, then it shall originate a BROADCAST (CHANGE SUPERVISOR) primitive.

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.

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 GERNERAL 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 the values of the following bits the REPORT GENERAL response:
 - A) the ZONING bit;
 - B) the CONFIGURE ROUTE TABLE bit; and
 - C) the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit;
- 3) 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.

ROW Note 2: 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 and BROADCAST primitives may not be limited to the zone groups that originated them.

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 3) 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) configure the SAS domain as follows:
 - A) if any REPORT GENERAL response had the CONFIGURE ROUTE TABLE bit set to one or the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one, then the expander 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); and
 - B) 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);
- 2) perform an SMP ACTIVE SUPERVISOR UPDATE function (see 10.4.3.aa) with each device that reported:
 - A) a non-zero SUPERVISOR PRIORITY field in its REPORT GENERAL response to update its ACTIVE SUPERVISOR SAS ADDRESS field and indicate the end of the election process;
 - B) a ZONING bit set to one in its REPORT GENERAL response to indicate that its zone route expander table information and zone permission table have been updated; and
 - C) an ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one to indicate that its expander route table has been updated;
- 3) set the ACTIVE SUPERVISOR SAS ADDRESS field in its REPORT GENERAL response data (see 10.4.3.3) to its SAS address;
- 4) set the CONFIGURING bit in its REPORT GENERAL response data to zero; and
- 5) originate a BROADCAST (CHANGE) primitive.

Upon receipt of an SMP ACTIVE SUPERVISOR UPDATE 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 UPDATE 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 3: One BROADCAST (CHANGE) originated 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 *in 06-082r0*); 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 UPDATE function (see 10.4.3.aa) as necessary (e.g., to notify a device that sent an SMP CHANGE request that its expander route table information, zone route expander table information, and zone permission table information has been updated in response to the change);
- e) start new device polling by SAS initiators, as needed, by performing one or more SMP SEND BROADCAST PRIMITIVES functions (see 10.4.3.cc *in 06-082r0*);
- f) perform the functions described in 4.7a.6 to limit the distribution of BROADCAST primitives to the zone groups that originate them; and

- g) restart the active supervisor election process in response to receipt of a BROADCAST (CHANGE SUPERVISOR) primitive.

If any devices in the SAS domain have the CONFIGURE ROUTE TABLE bit set to one, then the active supervisor shall poll the expander route tables in those devices to verify that their contents are compatible with the SAS domain configuration known to the active supervisor.

4.7a.5 Configuration change handling by devices that are not supervisor devices

4.7a.5.1 Expander devices that are not supervisor devices

If a non-supervisor expander device (i.e., an expander device that is not a supervisor (see 3.1.264)) with the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one in its REPORT GENERAL response determines that a change has occurred in its ability to communicate with any device, then it shall:

- a) perform an SMP CHANGE function (see 10.4.3.bb in 06-082r0) with the active supervisor device indicating the nature of the configuration change that it has detected;
- b) set the CONFIGURING bit in its REPORT GENERAL response data to one; and
- c) not originate a BROADCAST (CHANGE) primitive.

Upon receipt of an SMP ACTIVE SUPERVISOR UPDATE request, an expander device with the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one in its REPORT GENERAL response shall:

- a) note the SAS address of the active supervisor for use in future SMP CHANGE requests;
- b) set the CONFIGURING bit in its REPORT GENERAL response data to zero; and
- c) not originate a BROADCAST (CHANGE) primitive.

If an expander device with the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to zero in its REPORT GENERAL response determines that a change has occurred in its ability to communicate with any device, then it originates a BROADCAST (CHANGE) primitive as defined in SAS-1.1.

4.7a.5.2 End devices that are not supervisor devices

End devices that set the SMP TARGET PORT bit and SMP INITIATOR PORT bit to one in their IDENTIFY address frame (see 7.8.2) should assist in device discovery as described in this subclause. End devices that do not set the SMP TARGET PORT bit and SMP INITIATOR PORT bit to one in their IDENTIFY address frame are unable to assist in device discovery.

If an end device that is assisting with device discovery determines that a change has occurred in its ability to communicate with any device, then it should perform an SMP CHANGE function (see 10.4.3.bb in 06-082r0) with the active supervisor device indicating the nature of the configuration change that it has detected.

Upon receipt of an SMP ACTIVE SUPERVISOR UPDATE request, an end device that is assisting with device discovery should note the SAS address of the active supervisor for use in future SMP CHANGE requests.

4.7a.6 Limiting BROADCAST primitives to zone groups

If any device in the SAS domain supports zoning (see 4.9), the active supervisor device (see 3.1.4) shall set the ZONE BROADCASTS bit to one in the SMP ACTIVE SUPERVISOR UPDATE (see 10.4.3.aa) all requests that it sends.

If an SMP ACTIVE SUPERVISOR UPDATE request is received in which the ZONE BROADCASTS bit is set to one:

- a) Supervisor expander devices (see 3.1.264) shall process received BROADCAST primitives as defined in this subclause; and
- b) Other expander devices with the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to one in their REPORT GENERAL responses should process received BROADCAST primitives as defined in this subclause.

If the most recently received SMP ACTIVE SUPERVISOR UPDATE request had the ZONE BROADCASTS bit is set to one, then any received BROADCAST primitive except BROADCAST (CHANGE SUPERVISOR) shall:

- a) not be transmitted on other phys by the BPP (see 4.6.5); and
- b) instead an SMP CHANGE function (see 10.4.3.bb in 06-082r0) shall be performed with the active supervisor device in which the BROADCAST RECEIVED field is set to indicate which BROADCAST primitive has been received and the phy descriptor set to indicate the phy or phys on which the BROADCAST primitive was received.

When an SMP CHANGE request is received indicating that a BROADCAST primitive has been received, the active supervisor device shall perform as many SMP SEND BROADCAST PRIMITIVES functions (see 10.4.3.cc in 06-082r0) as are needed to originate the received BROADCAST primitive on every phy in the zone group within which the original BROADCAST primitive was received.

{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 95 in SAS-2 revision 2}

...

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 4: 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
	...							
10	ZONE DEVICE	Reserved				ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE	CONFIGURING	CONFIGURABLE ROUTE TABLE
11	ACTIVE ZONE SUPERVISOR PRIORITY Reserved				ZONE SUPERVISING SUPERVISOR PRIORITY			
12	ENCLOSURE LOGICAL IDENTIFIER							
19								
20	ACTIVE ZONE SUPERVISOR SAS ADDRESS							
27								
	...							

...

The **ZONE DEVICE** bit shall be set to one if the device supports zoning (see 4.9). If the **ZONE DEVICE** bit is set to one, then the **CONFIGURABLE ROUTE TABLE** bit shall be set to zero and the **ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE** bit shall be set to one. If the device does not support zoning, the **ZONE DEVICE** bit shall be set to zero.

The ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit shall be set to one if the device supports configuration of its expander route table by an active supervisor (see 3.1.4). If the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit is set to one, then the CONFIGURABLE ROUTE TABLE bit shall be set to zero. If the device does not support configuration of its expander route table by an active supervisor, the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE 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 2 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 device with the ACTIVE SUPERVISOR CONFIGURABLE ROUTE TABLE bit set to zero 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;
- b) the device is not a supervisor device; or
- c) 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 PARTICIPATING	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 COMPLETE ^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 2

4.6.5 Broadcast primitive processor (BPP)

The BPP receives broadcast primitive requests from each expander phy and processes them as follows:

- a) if an active supervisor device is limiting BROADCAST primitives to zone groups, then the BPP processes the broadcast primitive requests as defined in 4.7a.6; or
- b) otherwise, the BPP requests transmission of those requests on all expander ports except the expander port from which the broadcast primitive request was received.

In a self-configuring expander device (see 4.1.5), the BPP requests transmission of a BROADCAST (CHANGE) when it completes configuration (see 10.4.3.3).

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 (CHANGE SUPERVISOR Received)	Request to transmit a BROADCAST (CHANGE SUPERVISOR) on all other ports because a BROADCAST (CHANGE SUPERVISOR) was received.

ROW Note 5: Table 15 mentions BROADCAST (RESERVED 1) but not BROADCAST (RESERVED 2) which is used by this proposal. Therefore, it is assumed that BROADCAST (CHANGE SUPERVISOR) needs to be added. Proper references to BROADCAST (RESERVED 3) and BROADCAST (RESERVED 4) are left as an exercise for the editor.

{no other changes in 4.6.6.5}

7.2.2 Primitive summary

Table 71 defines the primitives not specific to the type of connection.

Table 71 — 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 1)	NoConn				I	E	T	Redundant
BROADCAST (RESERVED 2 CHANGE SUPERVISOR)	NoConn	I	E		I	E	T	Redundant
...								

{no other changes in 7.2.2}

7.2.3 Primitive encodings

Table 74 defines the primitive encoding for primitives not specific to type of connection.

Table 74 — Primitive encoding for primitives not specific to type of connection ...

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

{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 80.

Table 80 — BROADCAST primitives

Primitive	Description
...	
BROADCAST (RESERVED 2 CHANGE SUPERVISOR)	Reserved. Notification of need to start active supervisor device election.
...	

{Several paragraphs skipped}

BROADCAST (EXPANDER) is sent by an expander device to notify SAS initiator ports that an expander event has occurred, including:

- a) one or more of its phy event information peak value detectors has reached its threshold value; and
- b) one or more of its phy event information peak value detectors has been cleared by an SMP CONFIGURE PHY EVENT INFORMATION function (see 10.4.3.14).

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 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 190. 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 190 — 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		
...			
82h	WRITE GPIO REGISTER	See SFF-8485	
83h	CONFIGURE ZONE PERMISSION	Change zone permissions table	10.4.3.14
84h	ACTIVE SUPERVISOR UPDATE	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 214 — CONFIGURE GENERAL request

Byte/ Bit	7	6	5	4	3	2	1	0
	...							
9	SUPERVISOR PRIORITY				UPDATE SUPERVI- SOR PRI- ORITY	UPDATE STP SMP I_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 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 UPDATE function

The ACTIVE SUPERVISOR UPDATE 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 UPDATE request

Byte/ Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (84h)							
2	Reserved							
3	REQUEST LENGTH (03h)							
4	ACTIVE SUPERVISOR SAS ADDRESS							
11								
12	Reserved							ZONE BROAD- CASTS
13	Reserved							
15								
16	(MSB)	CRC						
19								(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 03h.

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).

If the ZONE BROADCASTS bit is set to one, an expander device shall convert any received BROADCAST primitive to an SMP CHANGE request as defined in (see 4.7a.6). If the ZONE BROADCASTS bit is set to zero, received BROADCAST primitives should not be converted to SMP CHANGE requests.

The CRC field is defined in 10.4.3.1.

Table aa2 defines the response format.

Table aa2 — ACTIVE SUPERVISOR UPDATE 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.

Annex J

...

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

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

{no other changes in annex J}