Date: 19 January 2006
To: T10 Technical Committee
From: Steve Johnson and Ralph Weber
Subject: SAS-2 Zoning-related SMP functions

Summary

Changes are proposed that replace the ZONE BROADCAST address frame proposed in 06-019r1 and 06-029r1 by a combination of two new SMP functions: SMP CHANGE and SMP SEND BROADCAST PRIMITIVES.

WARNING: This proposal only defines the SMP functions. New revisions of 06-019, 06-029, and proposals derived from them are expected to define the models for using these SMP functions.

Referenced Documents

SAS-2 revision 2	(http://www.t10.org/ftp/t10/drafts/sas2/sas2r02.pdf)
SAS-2 Zoning [06-019r1]	(http://www.t10.org/ftp/t10/document.06/06-019r1.pdf)
Zoning Changes [06-029r1]	(http://www.t10.org/ftp/t10/document.06/06-029r1.pdf)

Revision History

Prior to the preparation of this proposal, the SMP functions described herein were part of 06-029r1. This proposal separates the SMP functions so they can be referenced by new revisions of 06-019, 06-029, or proposals spawned from those works.

r0 Initial revision

Only those portions of 06-029r1 concerning SMP function definitions appear in this proposal. Within these constraints changes between 06-029r1 and this proposal are indicated by change bars.

Overview

Many 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 an SMP CHANGE request/response and an SMP SEND BROADCAST PRIMITIVES request/response.

The SMP CHANGE request/response is sent:

- from the device that detects a change in the SAS domain configuration or receives a BROADCAST primitive,
- to a supervising device whose exact nature depends on the configuration model adopted for SAS-2.

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

The SMP SEND BROADCAST PRIMITIVES request/response is designed to allow the origination of BROADCAST primitives to be confined to a zone group.

The SMP SEND BROADCAST PRIMITIVES request/response is sent:

- from a supervising device whose exact nature depends on the configuration model adopted for SAS-2,
- to an expander to which devices (both end devices and non-zoning expanders) that should receive the specified BROADCAST primitive are connected.

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

For a summary of substantial changes prior to the creation of this proposal please see 06-029r1.

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

This section describes substantial issues from the conference call and how this proposal responds them.

- **Modify the CHANGE request/response to make it more useful in zoned configurations:** A ZONE GROUP field has been added to each Phy Descriptor in the CHANGE request that indicates the zone group in which the Phy is participating.
- **Modify the CHANGE request/response to make allow its use in zoning BROADCAST primitives:** A field indicating that a CHANGE request results from the receipt of a BROADCAST has been added to the request frame. The field contains a coded value and the coded value allows for the possibility that the CHANGE request is not the result of receiving a BROADCAST primitive.
- **Modify the SEND BROADCAST (CHANGE) function to align its use with zoning:** Several changes have been made:
  - The function name has been changed to SEND BROADCAST PRIMITIVES;
  - A coded value has been added to the request that specifies which BROADCAST primitive is to be sent; and
  - The list of Phy Identifiers has been replaced with a list of Zone Groups (to which the specified BROADCAST primitive is to be originated).

Changes Proposed in SAS-2 revision 2

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 192.

Table 190 — SMP functions (FUNCTION field)

Code	SMP Function	Description	Reference
…			
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		{usage defined in 06-019r1}	
84h		{usage defined in 06-029r1}	
85h	SEND BROADCAST PRIMITIVES	Send specified BROADCAST primitive on the phy or phys in the specified zone group or zone groups	10.4.3.cc
8~~3~~6h - 8Fh	Reserved for general SMP output functions		
…			

10.4.3.bb CHANGE function

The CHANGE function is a request that is sent to the elected active supervisor device (see xxx) 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								
12	Reserved						BROADCAST RECEIVED	
13	Reserved							
14	(MSB)							
15	EXPANDER EXCHANGE COUNT							(LSB)
16	(MSB)							
17	EXPANDER ROUTE INDEXES							(LSB)
18	Reserved							
19	NUMBER OF PHYS							
20	Reserved							
21								
22	Reserved							
23	PHY DESCRIPTORS LENGTH ((n - 28) / 4)							
Phy descriptors								
24	first phy descriptor							
	. . .							
	last phy descriptor							
n-4								
n-3	(MSB)							
n	CRC							(LSB)

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 BROADCAST RECEIVED field (see table bb2) specifies the type of BROADCAST primitive (see 7.2.5.4), if any, whose receipt caused the CHANGE request to be sent.

Table bb2 — BROADCAST RECEIVED field

Code	Description
000b	No BROADCAST primitives received
001b	BROADCAST (CHANGE) primitive received
010b	BROADCAST (RESERVED CHANGE 0) primitive received
011b	BROADCAST (RESERVED CHANGE 1) primitive received
100b	BROADCAST (SES) primitive received
101b	BROADCAST (EXPANDER) primitive received
110b	BROADCAST (RESERVED 3) primitive received
111b	BROADCAST (RESERVED 4) primitive received

ROW Note 1: N.B. BROADCAST (RESERVED 2) – proposed to be BROADCAST (CHANGE SUPERVISIOR) in 06-029 – does not appear in table bb2. This is intentional.

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 bb3.

Table bb3 — Expander phy descriptor format

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

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 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 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 ZONE GROUP field specifies the zone group that contains the phy to which the descriptor applies.

The CHANGE request CRC field is defined in 10.4.3.1.

Table bb4 defines the response format.

Table bb4 — 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.

10.4.3.cc SEND BROADCAST PRIMITIVES function

The SEND BROADCAST PRIMITIVES function requests that the specified BROADCAST primitive be originated on all the phys that are in one or more specified zone groups.

Table cc1 defines the request format.

Table cc1 — SEND BROADCAST PRIMITIVES request

Byte/ Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (85h)							
2	Reserved							
3	REQUEST LENGTH ((n - 7) / 4)							
4	Reserved						BROADCAST PRIMITIVE	
5	NUMBER OF ZONE GROUPS (n - 4)							
Zone groups								
6	Reserved	ZONE GROUP (first)						
	:							
n-5	Reserved	ZONE GROUP (last)						
n-3	(MSB)	CRC						
n								(LSB)

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

The FUNCTION field shall be set to 85h.

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

The BROADCAST PRIMITIVE field (see table cc2) specifies the type of BROADCAST primitive (see 7.2.5.4) that shall be originated on each phy in the specified zone groups.

Table cc2 — BROADCAST PRIMITIVE field

Code	Description
000b	Reserved
001b	BROADCAST (CHANGE) primitive
010b	BROADCAST (RESERVED CHANGE 0) primitive
011b	BROADCAST (RESERVED CHANGE 1) primitive
100b	BROADCAST (SES) primitive
101b	BROADCAST (EXPANDER) primitive
110b	BROADCAST (RESERVED 3) primitive
111b	BROADCAST (RESERVED 4) primitive

The NUMBER OF ZONE GROUPS field specifies the number of zone groups in which the specified BROADCAST primitive is to be originated.

Each ZONE GROUP field specifies one zone group in which the specified BROADCAST primitive shall be originated on each participating phy.

The CRC field is defined in 10.4.3.1.

Table cc3 defines the response format.

Table cc3 — SEND BROADCAST PRIMITIVES response

Byte/ Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (85h)							
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 85h.

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.