

T10/07-008 revision 4

Date: February 09, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Expander Notification of Temporary Shutdown

1 Overview

Within the Service Delivery Subsystem in the SAS topology we have intelligent devices that will have to have their software updated. The software is updated by issuing WRITE BUFFER commands to a SES device within the expander device. When this occurs there may be a period of time while the flash is being written when the expander will not permit SMP functions and may limit traffic between it's phys. This period of time can, depending on the implementation, take longer than a minute. After the code is written the expander will cause a reset sequence to occur on all it's phys before allowing traffic to resume.

This may cause timeouts at every layer of SAS for any operations outstanding at the time of the code download. This is a problem that can be handled within the current SAS and SCSI definitions by the application client/initiator device that issued the WRITE BUFFER command. However, any other application client/initiator device may not be happy, as any outstanding operations could without warning just disappear.

This proposal defines a NOTIFY (EXPANDER LIMITING FUNCTION) that expanders could use to notify all devices it is connected to that it will be temporarily (as in seconds) be shutting down the SMP functions and may be limiting traffic.

After a SAS expander sends a NOTIFY (EXPANDER LIMITING FUNCTION) and the time-out expires that SAS expander responds to attempts to open a port with an OPEN REJECT (RETRY LONG) if access to that port is blocked as a result of the action that resulted in the NOTIFY (EXPANDER LIMITING FUNCTION). After the action that resulted in the NOTIFY (EXPANDER LIMITING FUNCTION) is complete the SAS expander initiates a link reset sequence on all phys.

The action a SAS device takes with it receives a NOTIFY (EXPANDER LIMITING FUNCTION) depends on the type of device it is:

- a) Expanders issues a BROADCAST(EXPANDER);
- b) Target devices ignore the BROADCAST (EXPANDER) and NOTIFY (EXPANDER LIMITING FUNCTION); and
- c) Initiator devices may (this behavior is not intended to specified in SAS, this is only informational):
 - A) after receiving a NOTIFY (EXPANDER LIMITING FUNCTION) terminate any outstanding I_T_L_Q nexus;
 - B) after receiving a BROADCAST (EXPANDER) issue a REPORT BROADCAST function to determine the source and cause of the BROADCAST (EXPANDER). If the source of the broadcast was as a result of a NOTIFY (EXPANDER LIMITING FUNCTION), then the Initiator should terminate any outstanding I_T_L_Q nexus; and
 - C) terminate any I_T_L_Q nexus that receive an OPEN REJECT (RETRY LONG) for a length of time longer than indicated by the WRITE BUFFER commands code update timeout.

This requires two new primitives and a new SAS function:

- a) NOTIFY (EXPANDER LIMITING FUNCTION) is needed to allow an expander to identify the port to be marked as going away;
- b) OPEN REJECT (RETRY LONG) Provides a response to a request to open a port that is different than the normal OPEN REJECT (NO DESTINATION). NOTE: Current implementations are already required to treat OPEN_REJECT (RESERVED CONTINUE 0) like an OPEN_REJECT (RETRY); and
- c) REPORT BROADCAST function that provides information on the source and cause of a broadcast.

2 SAS-2 changes

In addition to the changes below the following tables need to be updated with the new primitives.

NOTIFY (EXPANDER LIMITING FUNCTION) replaces NOTIFY (RESERVE 1) in Table 87 — Deletable primitives and Table 91 — Primitive encoding for deletable primitives.

OPEN REJECT (RETRY LONG) replaces OPEN_REJECT (RESERVED CONTINUE 0) in Table 88 — Primitives not specific to type of connection and Table 91 — Primitive encoding for deletable primitives.

4.1.13 Broadcasts

Broadcasts are used to notify all phys in the SAS domain about certain events. Broadcasts are transmitted using BROADCAST (see 7.2.5.5) or the SMP ZONED BROADCAST function (see 10.4.3.17).

Table 1 defines the types of Broadcast supported.

Table 1 — Broadcast types

Broadcast	Primitive	Description
Broadcast (Change)	yes	Originated by an expander device to notify SAS initiator ports that a SAS domain change has occurred (see 7.11). May also be originated by SAS initiator ports. Ignored by SAS target ports.
Broadcast (Reserved Change 0)	yes	Reserved. SAS ports (i.e, SAS initiator ports and SAS target ports) shall process this Broadcast the same as Broadcast (Change).
Broadcast (Reserved Change 1)	yes	Reserved. SAS ports shall process this Broadcast the same as Broadcast (Change).
Broadcast (SES)	yes	Originated by a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-4 and SES-2) accessible through a SAS target port in the SAS domain to notify SAS initiator ports of an asynchronous event. SSP initiator ports should poll all the logical units in the SAS domain with peripheral device types set to 0Dh to determine the source. SAS target ports shall ignore this Broadcast.
Broadcast (Expander)	yes	Originated by an expander device to notify SAS initiator ports that an expander event has occurred, including: <ul style="list-style-type: none"> a) a phy event information peak value detector has reached its threshold value; ef b) a phy event information peak value detector has been cleared by an SMP CONFIGURE PHY EVENT INFORMATION function (see 10.4.3.26); <u>or</u> c) a NOTIFY (EXPANDER LIMITING FUNCTION) was received (see 7.2.5.11.4). Expander events do not include SAS domain changes, which are communicated with Broadcast (Change).
Broadcast (Asynchronous Event)	yes	Originated by an SSP target port when an event occurs that causes one or more unit attention conditions to be established for one or more logical units accessible through the SSP target port. An SSP target port shall only originate one Broadcast (Asynchronous Event) for each event that affects multiple logical units accessible through the SSP target port (e.g., only one Broadcast (Asynchronous Event) is originated when a hard reset occurs).
Broadcast (Reserved 3)	yes	Reserved. SAS ports shall ignore this Broadcast.
Broadcast (Reserved 4)	yes	
Broadcast (Zone Activate)	no	Initiates the zone activate step (see 4.9.6.4). Devices that are not locked zoning expander devices shall ignore this Broadcast.
<p>^a All Broadcasts are supported by the SMP ZONED BROADCAST function (see 10.4.3.17). Broadcasts labeled “yes” are also transmitted via BROADCAST primitive sequences (see 7.2.5.5).</p>		

When an expander port receives a Broadcast, the BPP (see 4.6.5) shall forward the Broadcast on at least one phy in each other expander port if zoning is disabled, or forward the Broadcast as described in 4.9.5 if zoning is enabled.

An expander device is not required to queue multiple identical Broadcasts for the same expander port. If a second identical Broadcast is requested before the first Broadcast has been transmitted, the second Broadcast may be ignored.

See 10.4.3.3 for details on counting Broadcast (Change)s originated in an expander device. See 4.11 for details on phy event information.

7.2.5.11 NOTIFY

7.2.5.11.1 NOTIFY overview

NOTIFY may be transmitted in place of any ALIGN (see 7.2.5.2) being transmitted for physical link rate tolerance management (see 7.3) and rate matching (see 7.13). Substitution of a NOTIFY for an ALIGN may or may not affect the ALIGN rotation (i.e., the NOTIFY may take the place of one of the ALIGNs in the rotation through ALIGN (0), ALIGN (1), ALIGN (2), and ALIGN (3), or it may delay the rotation). A specific NOTIFY shall not be transmitted in more than three consecutive dwords until at least three other dwords have been transmitted.

NOTIFYs are deletable primitives (see 7.3).

The forwarding of NOTIFY through expander devices is as specified in table 2.

NOTIFY shall not be forwarded through expander devices. Expander devices shall substitute an ALIGN for a NOTIFY if necessary.

SAS target devices are not required to detect every transmitted NOTIFY.

The versions of NOTIFY representing different reasons are defined in table 2.

Table 2 — NOTIFY primitives

Primitive	Description	Reference
NOTIFY (ENABLE SPINUP)	Specify to a SAS target device that it may temporarily consume additional power while transitioning into the active or idle power condition state.	7.2.5.11.2
NOTIFY (POWER LOSS EXPECTED)	Specify to a SAS target device that power loss may occur within the time specified by the POWER LOSS TIMEOUT field in the Protocol-Specific Port mode page Shared Port Control subpage (see 10.2.7.2.4).	7.2.5.11.3
NOTIFY (EXPANDER LIMITING FUNCTION)	Specify to a SAS device that an expander device is going to temporarily have reduced function (e.g., disable SMP access, reduced performance, disable communication between SAS devices).	7.2.5.11.4
NOTIFY (RESERVED 1)	Reserved.	

NOTIFY (RESERVED 1) ~~and NOTIFY (RESERVED 2)~~ shall be ignored by all devices.

7.2.5.11.2 NOTIFY (ENABLE SPINUP)

...

7.2.5.11.3 NOTIFY (POWER LOSS EXPECTED)

...

7.2.5.11.4 NOTIFY (EXPANDER LIMITING FUNCTION)

NOTIFY (EXPANDER LIMITING FUNCTION) is transmitted by an expander port to specify that the expander device is going to temporarily have reduced function (e.g., disable SMP access, reduced performance, disable phy to phy communication) for a period of time. The period of time the expander device is going to be offline may be determined:

- a) by requesting write buffer command information using a REPORT SUPPORTED OPERATION CODES command (see SPC-4) issued to a SES device within the expander device; and
- b) from the contents of the MAXIMUM REDUCED FUNCTION TIME field (see 10.4.3.3).

If supported, a NOTIFY (EXPANDER LIMITING FUNCTION) shall be transmitted at least three times on each expander port and shall only transmit a NOTIFY (EXPANDER LIMITING FUNCTION) on one phy of each expander port. After the last NOTIFY (EXPANDER LIMITING FUNCTION) is transmitted the expander device shall wait the time specified in the TIME TO REDUCED FUNCTION field in the CONFIGURE GENERAL function (see 10.4.3.15): before reducing any expander functions.

After any expander functions have been reduced:

- a) on any attempt to open a connection to a phy which is not accessible because of the reduced function, respond with an OPEN REJECT (RETRY LONG) until the operation that caused the reduced function is complete; and
- b) if access to any SMP functions or virtual phys is blocked the expander shall not generate or propagate any broadcasts until the reduced function is complete.

After the operation that caused the reduced function is complete the expander shall initiate a link reset sequence on all phys.

If an expander device supports NOTIFY (EXPANDER LIMITING FUNCTION) and receives NOTIFY (EXPANDER LIMITING FUNCTION) on one of its SAS expander ports, then that expander device shall transmit BROADCAST (EXPANDER) on all expander ports except the expander port from which the NOTIFY (EXPANDER LIMITING FUNCTION) was received. If the phy that received the NOTIFY (EXPANDER LIMITING FUNCTION) is inside a ZPSDS then the source zone group shall be sent to the zone group of the phy that received the NOTIFY (EXPANDER LIMITING FUNCTION);

For the SAS initiator device rules on determining the cause of a BROADCAST (EXPANDER) 7.9.2.

SAS target devices shall ignore NOTIFY (EXPANDER LIMITING FUNCTION).

7.2.5.13 OPEN_REJECT

OPEN_REJECT specifies that a connection request has been rejected and specifies the reason for the rejection. The result of some OPEN_REJECTs is to abandon (i.e., not retry) the connection request and the result of other OPEN_REJECTs is to retry the connection request.

All of the OPEN_REJECT versions defined in table 3 shall result in the originating port abandoning the connection request.

Table 3 — OPEN_REJECT abandon primitives

Primitive	Originator	Description
OPEN_REJECT (BAD DESTINATION)	Expander phy	A connection request arrives through an expander phy using the direct routing or table routing method and the expander device determines the connection request would have to be routed to the same expander port as the expander port through which the connection request arrived (e.g., the destination SAS address equals the source SAS address), and the expander device has not chosen to return OPEN_REJECT (NO DESTINATION) (see 7.12.4.3).
OPEN_REJECT (CONNECTION RATE NOT SUPPORTED)	Any phy	The requested connection rate is not supported on some physical link on the pathway between the source phy and destination phy. When a SAS initiator phy is directly attached to a SAS target phy, the requested connection rate is not supported by the destination phy. The connection request may be modified and reattempted as described in 7.12.2.2.
OPEN_REJECT (PROTOCOL NOT SUPPORTED)	Destination phy	Phy with destination SAS address exists but the destination phy does not support the requested initiator/target role, protocol, initiator connection tag, or features (i.e., the values in the INITIATOR PORT bit, the PROTOCOL field, the INITIATOR CONNECTION TAG field, and/or the FEATURES field in the OPEN address frame are not supported).
OPEN_REJECT (RESERVED ABANDON 1)	Unknown	Reserved. Process the same as OPEN_REJECT (WRONG DESTINATION).
OPEN_REJECT (RESERVED ABANDON 2)		
OPEN_REJECT (RESERVED ABANDON 3)		
OPEN_REJECT (STP RESOURCES BUSY)	Destination phy	STP target port with destination SAS address exists but the STP target port has an affiliation with another STP initiator port or all of the available task file registers have been allocated to other STP initiator ports (see 7.17.4). Process the same as OPEN_REJECT (WRONG DESTINATION) for non-STP connection requests.
OPEN_REJECT (WRONG DESTINATION)	Destination phy	The destination SAS address does not match the SAS address of the SAS port to which the connection request was delivered.
OPEN_REJECT (ZONE VIOLATION)	Zoning expander phy	The connection request is from a zone group that does not have permission to access the zone group that contains the destination phy according to the zone permission table of an unlocked zoning expander device.

All of the OPEN_REJECT versions defined in table 4 shall result in the originating port retrying the connection request.

Table 4 — OPEN_REJECT retry primitives

Primitive	Originator	Description
OPEN_REJECT (NO DESTINATION) ^a	Expander phy	Either: a) No such destination phy; b) the expander device determines the connection request would have to be routed to the same expander port as the expander port through which the connection request arrived (e.g., the destination SAS address equals the source SAS address) and the expander device has not chosen to return OPEN_REJECT (BAD DESTINATION) (see 7.12.4.3); or c) the SAS address is valid for an STP target port in an STP/SATA bridge, but the initial Register - Device to Host FIS has not been successfully received (see 10.4.3.9).
OPEN_REJECT (PATHWAY BLOCKED) ^b	Expander phy	An expander device determined the pathway was blocked by higher priority connection requests.
OPEN_REJECT (RETRY LONG) ^c	Expander phy	Phy with destination SAS address exists but is not able to accept connections (see 7.2.5.11.4).
OPEN_REJECT (RESERVED CONTINUE 1) ^c	Unknown	Reserved. Process the same as OPEN_REJECT (RETRY).
OPEN_REJECT (RESERVED INITIALIZE 0) ^a	Unknown	Reserved. Process the same as OPEN_REJECT (NO DESTINATION).
OPEN_REJECT (RESERVED INITIALIZE 1) ^a		
OPEN_REJECT (RESERVED STOP 0) ^b	Unknown	Reserved. Process the same as OPEN_REJECT (PATHWAY BLOCKED).
OPEN_REJECT (RESERVED STOP 1) ^b		
OPEN_REJECT (RETRY) ^c	Destination phy or zoning expander phy	Phy with destination SAS address exists but is not able to accept connections, or the connection request is from a zone group that does not have permission to access the zone group that contains the destination phy according to the zone permission table of a locked zoning expander device.
<p>^a If the I_T Nexus Loss timer is already running, it continues running; if it is not already running, it is initialized and started. Stop retrying the connection request if the I_T Nexus Loss timer expires.</p> <p>^b If the I_T Nexus Loss timer is already running, it continues running. Stop retrying the connection request if the I_T Nexus Loss timer expires.</p> <p>^c If the I_T Nexus Loss timer (see 8.2.2) is already running, it is stopped.</p>		

NOTE 1 - Some SAS logical phys compliant with earlier versions of this standard also transmit OPEN_REJECT (RETRY) if they receive an OPEN address frame while their SL_CC state machines are in the SL_CC5:BreakWait state (see 7.14.4.7).

When a SAS logical phy detects more than one reason to transmit an OPEN_REJECT, the SL_CC state machine determines the priority in the SL_CC2:Selected state (see 7.14.4.4).

When an expander logical phy detects more than one reason to transmit an OPEN_REJECT, the ECM determines the priority (see 7.12.4).

See 7.12 for details on connection requests.

7.9.2 SAS initiator device rules

After a link reset sequence, or after receiving a Broadcast (Change), a management application client behind an SMP initiator port should perform a discover process (see 4.7).

When a discover process is performed after a link reset sequence, the management application client discovers all the devices in the SAS domain. When a discover process is performed after a Broadcast (Change), the management application client determines which devices have been added to or removed from the SAS domain.

The discover information may be used to select connection rates for connection requests (see 7.8.3).

After receiving a BROADCAST (EXPANDER), a management application client behind an SMP initiator port should issue a REPORT BROADCAST function to all expander devices to determine

- a) the event that caused the BROADCAST (EXPANDER); and
- b) which expander port to which the event is associated.

10.4.3.3 REPORT GENERAL function

The REPORT GENERAL function returns general information about the SAS device (e.g., a SAS device contained in an expander device). This SMP function shall be implemented by all management device servers.

Table 5 defines the request format.

Table 5 — REPORT GENERAL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (00h)							
2	Reserved							
3	REQUEST LENGTH (00h)							
4	(MSB)							
7	CRC (LSB)							

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

The FUNCTION field shall be set to 00h.

The REQUEST LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.1.

Table 6 defines the response format.

Table 6 — REPORT GENERAL response (part 1 of 2)

Byte/Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (00h)								
2	FUNCTION RESULT								
3	RESPONSE LENGTH (0Ch)								
4	(MSB)	EXPANDER CHANGE COUNT						(LSB)	
5									
6	(MSB)	EXPANDER ROUTE INDEXES						(LSB)	
7									
8	Reserved								
9	NUMBER OF PHYS								
10	TABLE TO TABLE SUPPORTED	Reserved		ZONE ADDRESS RESOLVED SUPPORTED	CONFIGURES OTHERS	CONFIGURING	EXTERNALLY CONFIGURABLE ROUTE TABLE		
11	Reserved								
12	ENCLOSURE LOGICAL IDENTIFIER								
19									
20	Reserved								
29									
30	(MSB)	STP BUS INACTIVITY TIME LIMIT						(LSB)	
31									
32	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						(LSB)	
33									
34	(MSB)	STP SMP I_T NEXUS LOSS TIME						(LSB)	
35									
36	Reserved			ZONE LOCKED	PHYSICAL PRESENCE SUPPORTED	PHYSICAL PRESENCE ASSERTED	ZONING SUPPORTED	ZONING ENABLED	
37	Reserved								
38	(MSB)	MAXIMUM NUMBER OF ROUTED SAS ADDRESSES						(LSB)	
39									

Table 6 — REPORT GENERAL response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
40	ACTIVE ZONE MANAGER SAS ADDRESS							
47								
48	(MSB)	ZONE LOCK INACTIVITY TIME LIMIT						(LSB)
49								
50	TIME TO REDUCED FUNCTION Reserved							
51	MAXIMUM REDUCED FUNCTION TIME RESERVED							
52	(MSB)	CRC						(LSB)
55								

....

[The TIME TO REDUCED FUNCTION field contains the time the expander waits from transmitting a NOTIFY \(EXPANDER LIMITING FUNCTION\) to reducing any expander functions. The content of the TIME TO REDUCED FUNCTION field is set by the CONFIGURE GENERAL function \(see 10.4.3.15\).](#)

[The MAXIMUM REDUCED FUNCTION TIME field contains the maximum time the expander device shall issue OPEN REJECT \(RETRY LONG\) in response to attempts to open a port if access to that port is blocked as a result of the action that resulted in the NOTIFY \(EXPANDER LIMITING FUNCTION\) to be sent \(see 7.2.5.11.4\). This timer shall start after the time to reduced function expires This timer value shall be in 1 s increments.](#)

10.4.3.15 CONFIGURE GENERAL function

The CONFIGURE GENERAL function requests actions by the device containing the management device server. This SMP function may be implemented by any management device server. In zoning expander devices, if zoning is enabled then this function shall only be processed from SMP initiator ports that have access to zone group 2 (see 4.9.3.2).

Table 7 defines the request format.

Table 7 — CONFIGURE GENERAL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (80h)							
2	Reserved							
3	REQUEST LENGTH (034h)							
4	(MSB)	EXPECTED EXPANDER CHANGE COUNT						(LSB)
5								
6	Reserved							
7								
8	Reserved				UPDATE TIME TO REDUCED FUNCTION	UPDATE STP SMP I_T NEXUS LOSS TIME	UPDATE STP MAXIMUM CONNECT TIME LIMIT	UPDATE STP BUS INACTIVITY TIME LIMIT
9	Reserved							
10	(MSB)	STP BUS INACTIVITY TIME LIMIT						(LSB)
11								
12	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						(LSB)
13								
14	(MSB)	STP SMP I_T NEXUS LOSS TIME						(LSB)
15								
16	TIME TO REDUCED FUNCTION							
17	Reserved							
19								
16 20	(MSB)	CRC						(LSB)
19 23								

...

An [UPDATE TIME TO REDUCED FUNCTION](#) bit set to one specifies that the [TIME TO REDUCED FUNCTION](#) field shall be honored. An [UPDATE TIME TO REDUCED FUNCTION](#) bit set to zero specifies that the [TIME TO REDUCED FUNCTION](#) field shall be ignored.

...

The TIME TO REDUCED FUNCTION field contains the minimum period that an expander device shall wait from the transmission of a NOTIFY (EXPANDER LIMITING FUNCTION) to reducing expander functions (see 7.2.5.11.4).

This timer value shall be in 100 ms increments. When this timer is exceeded, the expander device may initiate the reduced function. This value is reported in the TIME TO REDUCED FUNCTION field in the SMP REPORT GENERAL response (see 10.4.3.3). The expander device should set the default value for the TIME TO REDUCED FUNCTION field to 2 s (i.e., 14h).

10.4.3 SMP functions

10.4.3.1 SMP function request frame format

An SMP request frame is sent by a management application client via an SMP initiator port to request an SMP function be performed by a management device server. Table 8 defines the SMP request frame format.

Table 8 — SMP request frame format

Byte/Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION								
2	Reserved								
3	REQUEST LENGTH $((n - 7) / 4)$								
4	ADDITIONAL REQUEST BYTES								
m									
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

The SMP FRAME TYPE field is included in each frame format defined in this clause, although that field is parsed by the SMP transport layer (see 9.4). The SMP FRAME TYPE field is set to 40h.

The FUNCTION field specifies which SMP function is being requested and is defined in table 9. If the value in the FUNCTION field is not supported by the management device server, it shall return a function result of UNKNOWN SMP FUNCTION as described in table 13.

Table 9 — SMP functions (FUNCTION field) (part 1 of 3)

Code	SMP function	Description	Reference
00h	REPORT GENERAL	Return general information about the device	10.4.3.3
01h	REPORT MANUFACTURER INFORMATION	Return vendor and product identification	10.4.3.4
02h	READ GPIO REGISTER	See SFF-8485	
03h	REPORT SELF-CONFIGURATION STATUS	Return status of the discover process in a self-configuring expander device	10.4.3.5

Table 9 — SMP functions (FUNCTION field) (part 2 of 3)

Code	SMP function	Description	Reference
04h	REPORT ZONE PERMISSION	Return zone permission table active or shadow values	10.4.3.6
05h	REPORT BROADCAST	Return information on broadcasts	10.4.3.3
056h - 0Fh	Reserved for general SMP input functions		
10h	DISCOVER	Return information about the specified phy	10.4.3.7
11h	REPORT PHY ERROR LOG	Return error logging information about the specified phy	10.4.3.8
12h	REPORT PHY SATA	Return information about a phy currently attached to a SATA phy	10.4.3.9
13h	REPORT ROUTE INFORMATION	Return phy-based expander route table information	10.4.3.10
14h	REPORT PHY EVENT INFORMATION	Return phy event information for the specified phy	10.4.3.11
15h	REPORT PHY BROADCAST COUNTS	Return Broadcast counts	10.4.3.12
16h	DISCOVER LIST	Return information about the specified phys	10.4.3.13
17h	REPORT EXPANDER ROUTE TABLE	Return contents of the expander-based expander route table	10.4.3.14
18h - 1Fh	Reserved for phy-based SMP input functions		
20h - 3Fh	Reserved for SMP input functions		
40h - 7Fh	Vendor specific		
80h	CONFIGURE GENERAL	Configure the device	10.4.3.15
81h	ENABLE DISABLE ZONING	Enable or disable zoning	10.4.3.16
82h	WRITE GPIO REGISTER	See SFF-8485	
83h - 84h	Reserved for general SMP output functions		
85h	ZONED BROADCAST	Transmit the specified Broadcast on the expander ports in the specified zone group(s)	10.4.3.17
86h	ZONE LOCK	Lock a zoning expander device	10.4.3.18
87h	ZONE ACTIVATE	Set the zoning expander active values equal to the zoning expander shadow values	10.4.3.19
88h	ZONE UNLOCK	Unlock a zoning expander device	10.4.3.20
89h	Reserved for a zoning function		
8Ah	CONFIGURE ZONE PHY INFORMATION	Configure zone phy information	10.4.3.21
8Bh	CONFIGURE ZONE PERMISSION TABLE	Configure the zone permission table	10.4.3.22
8Ch - 8Fh	Reserved for general SMP output functions		
90h	CONFIGURE ROUTE INFORMATION	Change phy-based expander route table information	10.4.3.23

Table 9 — SMP functions (FUNCTION field) (part 3 of 3)

Code	SMP function	Description	Reference
91h	PHY CONTROL	Request actions by the specified phy	10.4.3.24
92h	PHY TEST FUNCTION	Request a test function by the specified phy	10.4.3.25
93h	CONFIGURE PHY EVENT INFORMATION	Configure phy event information for the specified phy	10.4.3.26
94h - 9Fh	Reserved for phy-based SMP output functions		
A0h - BFh	Reserved for SMP output functions		
C0h - FFh	Vendor specific		

The REQUEST LENGTH field specifies the number of dwords that follow, not including the CRC field. For compatibility with previous versions of this standard, a REQUEST LENGTH field set to 00h sometimes specifies a non-zero number of dwords; this is defined in the function description.

The ADDITIONAL REQUEST BYTES field definition and length are based on the SMP function. The maximum size of the ADDITIONAL REQUEST BYTES field is 1 024 bytes, making the maximum size of the frame 1 032 bytes (i.e., 1 024 bytes of data + 4 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the ADDITIONAL REQUEST BYTES field so the CRC field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field is included in each request frame format defined in this clause, although that field is defined by the SMP transport layer (see 9.4.1) and parsed by the SMP link layer (see 7.18).

10.4.3.2 SMP function response frame format

An SMP response frame is sent by a management device server via an SMP target port in response to an SMP request frame. Table 10 defines the SMP response frame format.

Table 10 — SMP response frame format

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION								
2	FUNCTION RESULT								
3	RESPONSE LENGTH $((n - 7) / 4)$								
4	ADDITIONAL RESPONSE BYTES								
m									
	Fill bytes, if needed								
n - 3	(MSB)	CRC							
n								(LSB)	

...

Table 11 defines the priority of the SMP function results defined in table 13.

Table 11 — Function result priority (part 1 of 3)

SMP function	SMP function result priority
REPORT GENERAL (see 10.4.3.3)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
REPORT MANUFACTURER INFORMATION (see 10.4.3.4)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
READ GPIO REGISTER (see SFF-8485)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
REPORT SELF-CONFIGURATION STATUS (see 10.4.3.5)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
REPORT ZONE PERMISSION TABLE (see 10.4.3.6)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
REPORT BROADCAST (see 10.4.3.3.3)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
DISCOVER (see 10.4.3.7)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED
REPORT PHY ERROR LOG (see 10.4.3.8)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED
REPORT PHY SATA (see 10.4.3.9)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) PHY DOES NOT SUPPORT SATA; 5) SMP FUNCTION FAILED; and 6) SMP FUNCTION ACCEPTED
REPORT ROUTE INFORMATION (see 10.4.3.10)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) INDEX DOES NOT EXIST; 5) SMP FUNCTION FAILED; and 6) SMP FUNCTION ACCEPTED
REPORT PHY EVENT INFORMATION (see 10.4.3.11)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED

Table 11 — Function result priority (part 2 of 3)

SMP function	SMP function result priority
REPORT PHY BROADCAST COUNTS (see 10.4.3.12)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
DISCOVER LIST (see 10.4.3.13)	1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) UNKNOWN DESCRIPTOR TYPE; 4) UNKNOWN PHY FILTER; 5) SMP FUNCTION FAILED; and 6) SMP FUNCTION ACCEPTED
REPORT EXPANDER ROUTE TABLE (see 10.4.3.14)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
CONFIGURE GENERAL (see 10.4.3.15)	1) INVALID REQUEST FRAME LENGTH; 2) SMP ZONE VIOLATION; 3) INVALID EXPANDER CHANGE COUNT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED
WRITE GPIO REGISTER (see SFF-8485)	1) INVALID REQUEST FRAME LENGTH; 2) SMP FUNCTION FAILED; and 3) SMP FUNCTION ACCEPTED
ENABLE DISABLE ZONING (see 10.4.3.16)	1) INVALID REQUEST FRAME LENGTH; 2) ZONE LOCK VIOLATION; 3) UNKNOWN ENABLE DISABLE ZONING VALUE; 4) NO MANAGEMENT ACCESS RIGHTS; 5) INVALID EXPANDER CHANGE COUNT; 6) SMP FUNCTION FAILED; and 7) SMP FUNCTION ACCEPTED
ZONED BROADCAST (see 10.4.3.17)	1) INVALID REQUEST FRAME LENGTH; 2) SMP ZONE VIOLATION; 3) SMP FUNCTION FAILED; and 4) SMP FUNCTION ACCEPTED
ZONE LOCK (see 10.4.3.18)	1) INVALID REQUEST FRAME LENGTH; 2) ZONE LOCK VIOLATION; 3) NO MANAGEMENT ACCESS RIGHTS; 4) INVALID EXPANDER CHANGE COUNT; 5) SMP FUNCTION FAILED; and 6) SMP FUNCTION ACCEPTED
ZONE ACTIVATE (see 10.4.3.19)	1) INVALID REQUEST FRAME LENGTH; 2) ZONE LOCK VIOLATION; 3) INVALID EXPANDER CHANGE COUNT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED
ZONE UNLOCK (see 10.4.3.20)	1) INVALID REQUEST FRAME LENGTH; 2) ZONE LOCK VIOLATION; 3) NOT ACTIVATED; 4) BUSY; 5) SMP FUNCTION FAILED; and 6) SMP FUNCTION ACCEPTED

Table 11 — Function result priority (part 3 of 3)

SMP function	SMP function result priority
CONFIGURE ZONE PHY INFORMATION (see 10.4.3.21)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) ZONE LOCK VIOLATION; 4) UNKNOWN ZONE PHY INFORMATION VALUE; 5) INVALID EXPANDER CHANGE COUNT; 6) SMP FUNCTION FAILED; and 7) SMP FUNCTION ACCEPTED
CONFIGURE ZONE PERMISSION (see 10.4.3.22)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) ZONE LOCK VIOLATION; 3) INVALID EXPANDER CHANGE COUNT; 4) SMP FUNCTION FAILED; and 5) SMP FUNCTION ACCEPTED
CONFIGURE ROUTE INFORMATION (see 10.4.3.23)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) INDEX DOES NOT EXIST; 5) INVALID EXPANDER CHANGE COUNT; 6) SMP FUNCTION FAILED; and 7) SMP FUNCTION ACCEPTED
PHY CONTROL (see 10.4.3.24)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP ZONE VIOLATION; 5) LOGICAL LINK RATE NOT SUPPORTED; 6) UNKNOWN PHY OPERATION; 7) PHY DOES NOT SUPPORT SATA; 8) INVALID EXPANDER CHANGE COUNT; 9) SMP FUNCTION FAILED; and 10) SMP FUNCTION ACCEPTED
PHY TEST FUNCTION (see 10.4.3.25)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP ZONE VIOLATION; 5) UNKNOWN PHY TEST FUNCTION; 6) PHY TEST FUNCTION IN PROGRESS; 7) INVALID EXPANDER CHANGE COUNT; 8) SMP FUNCTION FAILED; and 9) SMP FUNCTION ACCEPTED
CONFIGURE PHY EVENT INFORMATION (see 10.4.3.26)	<ol style="list-style-type: none"> 1) INVALID REQUEST FRAME LENGTH; 2) PHY DOES NOT EXIST; 3) PHY VACANT; 4) SMP ZONE VIOLATION; 5) UNKNOWN PHY EVENT INFORMATION SOURCE; 6) INVALID EXPANDER CHANGE COUNT; 7) SMP FUNCTION FAILED; and 8) SMP FUNCTION ACCEPTED

The RESPONSE LENGTH field indicates the number of dwords that follow, not including the CRC field. For compatibility with previous versions of this standard, a RESPONSE LENGTH field set to 00h sometimes indicates a non-zero number of dwords; this is defined in the function description.

The ADDITIONAL RESPONSE BYTES field definition depends on the SMP function requested. The maximum size of the ADDITIONAL RESPONSE BYTES field is 1 024 bytes, making the maximum size of the frame 1 032 bytes (i.e., 1 024 bytes of data + 4 bytes of header + 4 bytes of CRC).

Fill bytes shall be included after the ADDITIONAL RESPONSE BYTES field so the CRC field is aligned on a four byte boundary. The contents of the fill bytes are vendor specific.

The CRC field is included in each response frame format defined in this clause, although that field is defined by the SMP transport layer (see 9.4.1) and parsed by the SMP link layer (see 7.18).

10.4.3.3 REPORT BROADCAST function

[Editor's Note 1: All new in this section](#)

10.4.3.3.1 REPORT BROADCAST function overview

The REPORT BROADCAST function returns information about broadcasts that were issued from this expander device. This SMP function may implemented by any management device server. An expander device is not required to maintain broadcast information in non-volatile storage or across events that cause the expander device to be reset.

10.4.3.3.2 REPORT BROADCAST request

Table 12 defines the request format.

Table 12 — REPORT BROADCAST request

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (40h)								
1	FUNCTION (05h)								
2	Reserved								
3	REQUEST LENGTH (00h)								
12	(MSB)	CRC							
15								(LSB)	

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

The FUNCTION field shall be set to 05h.

The REQUEST LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.1.

10.4.3.3.3 REPORT BROADCAST response

Table 239 defines the response format.

Table 13 — REPORT BROADCAST response

Byte/Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (05h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH							
4	(MSB)	Reserved						(LSB)
6								
7	NUMBER OF BROADCAST DESCRIPTORS							
Broadcast descriptor list								
8	Broadcast descriptor (first) (see table 14)							
23								
...	...							
n - 19	Broadcast descriptor (last) (see table 14)							
n - 4								
n - 3	(MSB)	CRC						(LSB)
n								

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

The FUNCTION field shall be set to 05h.

The FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field contains the number of dwords that follow, not including the CRC field.

The NUMBER OF BROADCAST DESCRIPTORS field indicates how many broadcast descriptors follow.

NOTE 2 - The number of broadcast descriptors is limited to 63.

The broadcast descriptor list contains broadcast descriptors as defined in 10.4.3.3.4.

The CRC field is defined in 10.4.3.2.

10.4.3.3.4 REPORT BROADCAST response phy event descriptor

Table 14 defines the broadcast descriptor.

Table 14 — Broadcast descriptor

Byte/Bit	7	6	5	4	3	2	1	0
0	BRDC_SCR	Reserved			BROADCAST TYPE			
1	Reserved			BROADCAST REASON CODE				
2	Reserved							
3	(MSB)	CHANGE COUNT						(LSB)
4								
5	(MSB)	Reserved						(LSB)
6								
7	PHY IDENTIFIER							
8	(MSB)	ATTACHED SAS ADDRESS						(LSB)
15								

The BROADCAST TYPE field, defined in table 264 in 10.4.3.17, indicates the last broadcast of this type that was received by the expander device.

A broadcast source (BRDC_SCR) bit set to zero indicates that the broadcast originated from this expander device. A BRDC_SCR bit set to one indicates that the broadcast originated from another SAS device.

The BROADCAST REASON CODE field indicates the reason the broadcast indicated in the BROADCAST TYPE field was transmitted as defined in table 15.

Table 15 — BROADCAST REASON CODE field

Broadcast type	Code	Description
Broadcast (Change)	0h	Broadcast (Change)
	1h - Fh	Reserved
Broadcast (Reserved Change 0)	0h	Broadcast (Reserved Change 0)
	1h - Fh	Reserved
Broadcast (Reserved Change 1)	0h	Broadcast (Reserved Change 1)
	1h - Fh	Reserved
Broadcast (SES)	0h	Broadcast (SES)
	1h - Fh	Reserved
Broadcast (Expander)	0h	Unknown
	1h	A phy event information peak value detector has reached its threshold value.
	2h	A phy event information peak value detector has been cleared by an SMP CONFIGURE PHY EVENT INFORMATION function (see 10.4.3.26).
	3h	A NOTIFY (EXPANDER LIMITING FUNCTION) was received (see 7.2.5.11.4)
	4h	A BROADCAST (EXPANDER) was received
	5h - Fh	Reserved
Broadcast (Asynchronous Even)	0h	Broadcast (Asynchronous Even)
	1h - Fh	Reserved
Broadcast (Reserved 3)	0h	Broadcast (Reserved 3)
	1h - Fh	Reserved
Broadcast (Reserved 4)	0h	Broadcast (Reserved 4)
	1h - Fh	Reserved
Broadcast (Zone Activate)	0h	Broadcast (Zone Activate)
	1h - Fh	Reserved

The CHANGE COUNT field counts the number of broadcasts of the type specified in the BROADCAST TYPE field originated by an expander device. This field shall be set to at least 0001h at power on. If the expander device has originated the broadcast for any reason since transmitting a REPORT BROADCAST response, it shall increment this field at least once from the value in the previous REPORT BROADCAST response. It shall not increment this field when forwarding a broadcast. This field shall wrap to at least 0001h after the maximum value (i.e., FFFFh) has been reached.

NOTE 3 - Application clients that use the CHANGE COUNT field should read it often enough to ensure that it does not increment a multiple of 65 536 times between reading the field.

If the BRDC_SCR bit is set to one, then the PHY IDENTIFIER field indicates the phy through which the broadcast was received. If the BRDC_SCR bit is set to zero, then the PHY IDENTIFIER field shall be ignored.

If the BRDC_SCR bit is set to one, then the SAS ADDRESS field contains the value of the SAS ADDRESS field received in the IDENTIFY address frame during the identification sequence associated with the phy through which the broadcast was received. If the BRDC_SCR bit is set to zero, then the SAS ADDRESS field contains the SAS address of this expander device.

3 SPC-4 changes

Add the following section to the WRITE BUFFERS command. This change assumes that 05-284 is accepted.

6.36.1 WRITE BUFFER commands command processing times descriptor

[Editor's Note 2: All new in this section](#)

The command processing times descriptor (see table 16) reported by the REPORT SUPPORTED OPERATION CODES command (x.x) for a WRITE BUFFER command indicates timeout information specific to the WRITE BUFFER command.

Table 16 — Command Processing Times descriptor

Byte/Bit	7	6	5	4	3	2	1	0	
0	(MSB)	DESCRIPTOR LENGTH (0Ah)							
1								(LSB)	
2	Reserved								
3	MAXIMUM REDUCED FUNCTION TIME								
4	(MSB)	NOMINAL COMMAND PROCESSING TIME							
7								(LSB)	
8	(MSB)	RECOMMENDED COMMAND TIMEOUT							
11								(LSB)	

The MAXIMUM REDUCED FUNCTION TIME field contains the maximum time access to the SCSI device is limited or not possible through any SCSI ports associated with the logical unit that receives the WRITE BUFFER command. This timer value shall be in 1 s increments. A value of zero in the MAXIMUM REDUCED FUNCTION TIME field indicates that there is no maximum time specified. The MAXIMUM REDUCED FUNCTION TIME field shall only apply when the following modes are specified:

- a) Download microcode mode (04h);
- b) Download microcode and save mode (05h);
- c) Download microcode with offsets mode (06h);
- d) Download microcode with offsets and save mode (07h);
- e) Download microcode with offsets and defer activation mode (0Eh) only if the microcode is activated by an event other than an activate deferred microcode mode; and
- f) Activate deferred microcode mode (0Fh).

The NOMINAL COMMAND PROCESSING TIME field and RECOMMENDED COMMAND TIMEOUT field are defined in x.x.