

To: T10 Technical Committee
From: Ed D'Avignon, Vitesse Semiconductor (davignon@vitesse.com) and Rob Elliott, HP (elliott@hp.com)
Date: 18 April 2006
Subject: 06-177r1 SAS-2 ZONED BROADCAST processing clarification

Revision history

Revision 0 (30 March 2006) First revision (by Ed D'Avignon) - clarified 4.8.5 meaning
Revision 1 (18 April 2006) Incorporated comments from 30 March 2006 SAS zoning teleconference. Added BROADCAST terminology clarifications throughout the standard to differentiate BROADCAST primitives from Broadcast information (delivered via the BROADCAST primitive or SMP ZONED BROADCAST function) (by Rob Elliott).

Related documents

sas2r03 - Serial Attached SCSI - 2 (SAS-2) revision 3
06-122r2 SAS-2 zoning - phy features (Ralph Weber, ENDL) - incorporated into sas2r03
06-082r2 SAS-2 zoning-related SMP functions (Ralph Weber, ENDL and Steve Johnson, LSI Logic) - incorporated into sas2r03
06-xxxxr0 - SAS-2 No BROADCASTs before link reset sequence completes (Luben Toikov, Vitesse and Rob Elliott, HP)

Overview

sas2r03 describes the behavior of a zoning expander when a BROADCAST primitive or an SMP ORIGINATE BROADCAST PRIMITIVES (renamed ZONED BROADCAST) request is received. The intent was for a BROADCAST primitive or an SMP ZONED BROADCAST request to be transmitted by the zoning expander using a similar set of rules. Unfortunately, 06-082r2 had different assumptions, resulting in different behavior.

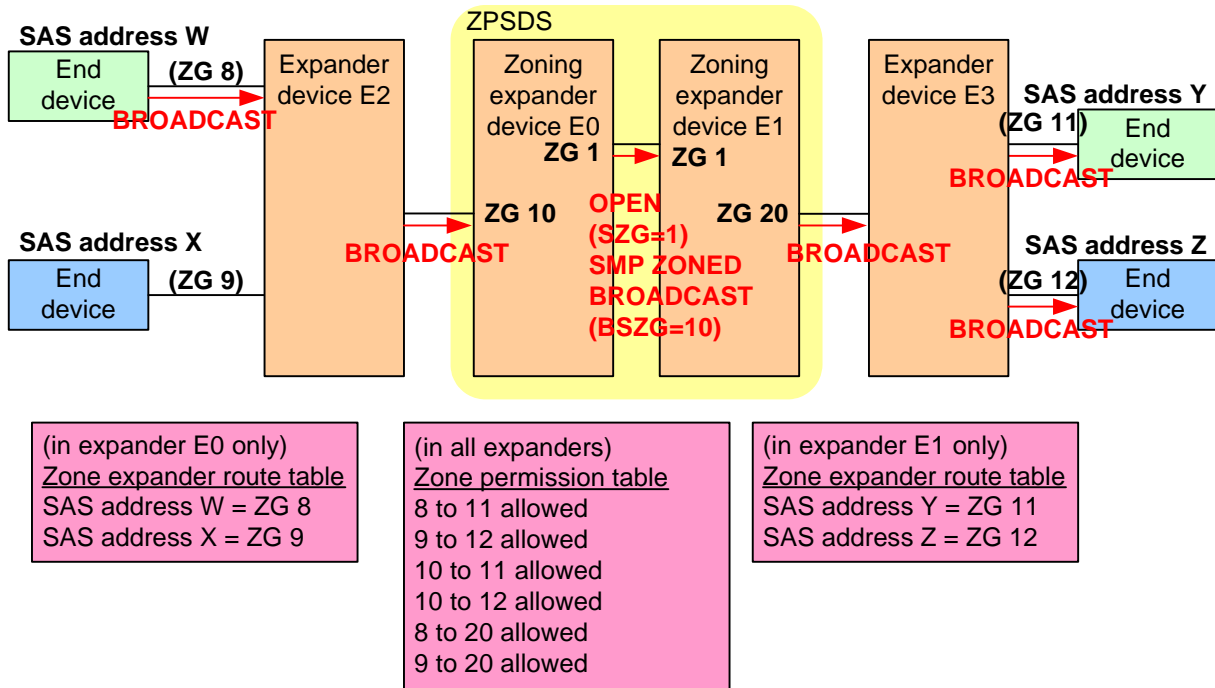
For example, a received BROADCAST primitive will be transmitted on ports with the zone participating bit set to zero only when the group ID of the receiving and transmitting ports match. The intended rule is for the BROADCAST primitive to be transmitted on all ports which have permission to communicate with the receiving port according to the zone permission table.

Editorially, the BROADCAST terminology throughout the standard needs to be overhauled, since BROADCAST or BROADCAST (XYZ) historically both the primitive and the information it conveys. Now, the SMP ZONED BROADCAST request can also carry the broadcast information. Mixed-case "Broadcast" (e.g. (Broadcast (Change))) is used everywhere the broadcast information is being referred to, while BROADCAST refers only to the primitive.

The zone address resolved method complicates matters if avoiding over-broadcasting outside the ZPSDS is a concern. SAS-2 expanders could support the SMP ZONED BROADCAST function to facilitate this to some extent.

Figure 1 shows an example of an ZPSDS attached to SAS-1/SAS-1.1 expanders outside the ZPSDS.

With legacy expanders outside the ZPSDS that only understand BROADCAST primitives:



If “new” expanders outside the ZPSDS understand the SMP ZONED BROADCAST request, they can partially participate in zoning (while still being considered outside the ZPSDS):

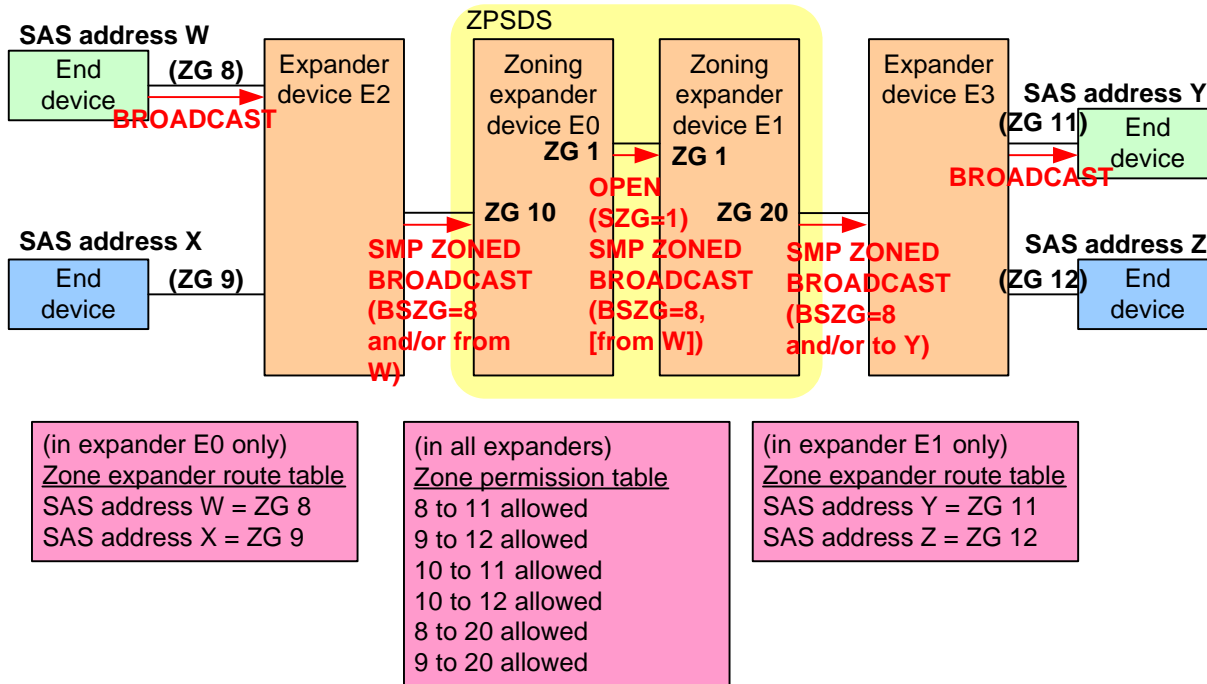
- a) from outside the ZPSDS includes source zone groups and/or source SAS address(es)
- b) inside the ZPSDS includes source zone groups and/or the source SAS address(es)
- c) to outside the ZPSDS includes source zone groups and/or destination source SAS address(es)

Figure 1 — ZPSDS attached to SAS-1/SAS-1.1 expanders outside the ZPSDS

Figure 2 shows an example of a ZPSDS attached to SAS-2 expanders outside the ZPSDS.

If “new” expanders outside the ZPSDS understand the SMP ZONED BROADCAST request, they can partially participate in zoning (while still being considered outside the ZPSDS):

- a) from outside the ZPSDS includes source zone groups and/or source SAS address(es)
- b) inside the ZPSDS includes source zone groups and/or the source SAS address(es)
- c) to outside the ZPSDS includes source zone groups and/or destination source SAS address(es)



They could have zone permission tables and phy zone information and work from the SZG in the ZONED BROADCAST request, or the ZPSDS ingress boundary expander could translate between a source SAS addresses and their SZG, and the ZPSDS egress boundary expander could translate between the SZGs and the destination SAS addresses.

Figure 2 — ZPSDS attached to SAS-2 expanders outside the ZPSDS

Editor’s Note 1: Discussion of those figures may result in ZONED BROADCAST carrying more than just a list of source zone groups

Suggested changes

3.1.xx broadcast: Information about an event in the SAS domain, communicated between phys with the BROADCAST primitive (see 7.2.5.4) and/or the SMP ZONED BROADCAST function (see 10.4.3.11).

3.1.17 broadcast primitive propagation processor (BPP): An object within an expander function that manages broadcasts primitives. See 4.6.5.

3.1.66 expander function: An object within an expander device that contains an expander connection manager, expander connection router, and broadcast primitive propagation processor. See 4.6.1.

3.2 Symbols and abbreviations

Abbreviation Meaning

BPP broadcast primitive propagation processor (see 3.1.17)

4.1.8 Expander devices (edge expander devices and fanout expander devices)

Expander devices are part of the service delivery subsystem and facilitate communication between multiple SAS devices. Expander devices contain two or more external expander ports. Each expander device contains one SMP target port [and a management device server](#) for management, may contain one SMP initiator port [and a management application client for self-configuration](#), and may contain SAS devices (e.g., an expander device may include an SSP target port for access to a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2)).

[\[In figure 18, add a “management application client” box connected to the “SMP initiator port” box and and a “management device server” box connected to the “SMP target port” box.\]](#)

4.1.11 Broadcasts

[Broadcasts are used to notify all phys in the SAS domain about certain events. Broadcasts are transmitted using BROADCAST \(see 7.2.5.4\) or the SMP ZONED BROADCAST function \(see 10.4.3.11\)](#)

[Table 1 defines the types of broadcast supported.](#)

Table 1 — Broadcast types

Broadcast	Description
Broadcast (Change)	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)	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)	Reserved. SAS ports shall process this Broadcast the same as Broadcast (Change).
Broadcast (SES)	<p>Originated by a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2) accessible through a SAS target port in the SAS domain to notify SAS initiator ports of an asynchronous event.</p> <p>SSP initiator ports should poll all the logical units in the SAS domain with peripheral device types set to 0Dh to determine the source.</p> <p>SAS target ports shall ignore this Broadcast.</p>
Broadcast (Expander)	<p>Originated by an expander device to notify SAS initiator ports that an expander event has occurred, including:</p> <ul style="list-style-type: none"> a) a phy event information peak value detector has reached its threshold value; or b) a phy event information peak value detector has been cleared by an SMP CONFIGURE PHY EVENT INFORMATION function (see 10.4.3.15). <p>Expander events do not include SAS domain changes, which are communicated with Broadcast (Change).</p>
Broadcast (Reserved 2)	Reserved. SAS ports shall ignore this Broadcast.
Broadcast (Reserved 3)	Reserved. SAS ports shall ignore this Broadcast.
Broadcast (Reserved 4)	Reserved. SAS ports shall ignore this Broadcast.

Editor’s Note 2: The following text was moved from 7.2.5.4. Most of the paragraphs were merged into the table itself:

When an expander port receives a ~~BROADCASTBroadcast~~ it shall transmit the same ~~BROADCAST~~ on at least one phy in all other expander ports, 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 4.8.5 if zoning is enabled.

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

~~BROADCAST (CHANGE) is sent by an expander device to notify SAS initiator ports and other expander devices that a configuration change has occurred. BROADCAST (CHANGE) may also be transmitted by SAS initiator ports. BROADCAST (CHANGE) shall be ignored by SAS target ports.~~

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

~~BROADCAST (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; or~~
- ~~b) one or more of its phy event information peak value detectors has been cleared by an SMP-CONFIGURE PHY EVENT INFORMATION function (see 10.4.3.15).~~

~~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.~~

~~See 7.11 for details on SAS domain changes. See 10.4.3.3 for details on counting BROADCAST (CHANGE) Broadcast (Change)s generation originated in an expander device. See 4.10 for details on phy event information.~~

Editor's Note 3: End of moved text

4.3.4 State machines and SAS device, SAS port, and SAS phy classes

[In figure 40, change "Broadcast primitive processor" to "Broadcast propagation processor" (twice)]

4.6.1 Expander device model overview

An expander device shall contain the following:

- a) an expander function containing:
 - A) expander connection manager (ECM);
 - B) expander connection router (ECR); and
 - C) broadcast ~~primitive~~propagation processor (BPP);
- b) two or more physical expander phys. For the maximum number of phys, see 4.1.5;
- c) an expander port available per phy; and
- d) an SMP target port with a management device server.

An expander device may contain an SMP initiator port with a management application client. An expander device may contain SAS devices with SSP ports, STP ports, and/or SMP ports and their associated device servers and application clients.

[In figure 42, change "Broadcast primitive processor" to "Broadcast propagation processor"]

4.6.5 Broadcast primitive processor (BPP)

The BPP receives ~~BROADCASTBroadcasts~~ (see 4.1.11) from each expander phy or from the management device server on behalf of an expander phy and requests transmission of those ~~BROADCASTBroadcasts~~ on all expander ports except the expander port from which the ~~BROADCASTBroadcast~~ was received.

In a self-configuring expander device (see 4.1.5), the BPP requests transmission of a ~~BROADCAST~~ ~~(CHANGE)~~ [Broadcast \(Change\)](#) when it completes configuration (see 10.4.3.3).

In a zoning expander device (see 4.8.2), ~~BROADCASTs and SMP ZONED BROADCAST requests are forwarded as BROADCASTs and/or SMP ZONED BROADCAST requests (see 10.4.3.11)~~ [Broadcasts are forwarded as described in 4.8.5.](#)

4.6.6.5 BPP interface

Table 2 describes the requests from an expander phy to the BPP. See 7.11 for more information on broadcasts. [Requests from the management device server about SMP ZONED BROADCAST requests](#)

[received from the SMP target port in zoning expander devices are not described.](#) See 4.8.5 for more information on how zoning expander devices handle broadcasts.

Table 2 — Expander phy to BPP requests

Message	Description
Broadcast Event Notify (Phy Not Ready)	Request to transmit a BROADCAST (CHANGE) on all other ports originate a Broadcast (Change) because an expander phy's SP state machine transitioned from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.8).
Broadcast Event Notify (SATA Spinup Hold)	Request to transmit a BROADCAST (CHANGE) on all other ports originate a Broadcast (Change) because the SATA spinup hold state has been reached (see 6.8 and 6.10).
Broadcast Event Notify (Identification Sequence Complete)	Request to transmit a BROADCAST (CHANGE) on all other ports originate a Broadcast (Change) because an expander phy has completed the identification sequence (see 7.9).
Broadcast Event Notify (SATA Port Selector Change)	Request to transmit a BROADCAST (CHANGE) on all other ports originate a Broadcast (Change) because a SATA port selector appeared or disappeared.
Broadcast Event Notify (CHANGE Change Received)	Request to transmit a BROADCAST (CHANGE) on all other ports forward a Broadcast (Change) because a BROADCAST (CHANGE) Broadcast (Change) was received.
Broadcast Event Notify (RESERVED-CHANGE Reserved Change 0 Received)	Request to transmit a BROADCAST (CHANGE) on all other ports forward a Broadcast (Change) because a BROADCAST (RESERVED-CHANGE 0) Broadcast (Reserved Change 0) was received.
Broadcast Event Notify (Reserved Change 1 Received)	Request to forward a Broadcast (Reserved Change 1) because a Broadcast (Reserved Change 1) was received.
Broadcast Event Notify (SES Received)	Request to transmit a BROADCAST (CHANGE) on all other ports forward a Broadcast (SES) because a BROADCAST (SES) Broadcast (SES) was received.
Broadcast Event Notify (EXPANDER Expander Received)	Request to transmit a BROADCAST (CHANGE) on all other ports forward a Broadcast (Expander) because a BROADCAST (EXPANDER) Broadcast (Expander) was received.
Broadcast Event Notify (Reserved 2 Received)	Request to forward a Broadcast (Reserved 2) because a Broadcast (Reserved 2) was received.
Broadcast Event Notify (Reserved 3 Received)	Request to forward a Broadcast (Reserved 3) because a Broadcast (Reserved 3) was received.
Broadcast Event Notify (Reserved 4 Received)	Request to forward a Broadcast (Reserved 4) because a Broadcast (Reserved 4) was received.

Table 3 describes the indications from the BPP to an expander phy. [Indications to the management application client to generate SMP ZONED BROADCAST functions from the SMP initiator port in a zoning](#)

[expander device are not described. See 4.8.5 for more information on how zoning expander devices handle broadcasts.](#)

Table 3 — BPP to expander phy indications

Message	Description
Transmit Broadcast (type)	Indication to transmit a BROADCAST with the specified type.

4.8 Zoning

4.8.1 Zoning overview

SAS zoning is implemented by a set of zoning expander devices that define the zoned portion of a service delivery subsystem (ZPSDS). The zoning expander devices control whether a phy is permitted to participate in a connection to another phy.

There shall be at most one ZPSDS in a SAS domain. The ZPSDS may encompass some or all of the SAS domain.

Every phy in a ZPSDS belongs to a zone group. All phys in a wide port shall belong to the same zone group. Each zoning expander contains a zone permission table that controls whether a connection is allowed between the zone group of a source port and the zone group of a target port. There are 128 zone groups.

A requested connection shall only be established if the zone permission table indicates that access between the zone group of the source port and the zone group of the target port is allowed.

The zoning expander route table is an extended version of the expander route table (see 4.6.7.3) that also includes zone phy information.

The treatment of OPEN address frames received from expander devices that are outside the boundary of the ZPSDS depends on:

- a) whether the OPEN address frame enters the ZPSDS via a table routed phy; and
- b) the value of the ZONE ADDRESS RESOLVED bit for the receiving phy (see 4.8.3.6).

The following cases apply:

- a) if an OPEN address frame is received from a non-zoning expander device via a table routed phy for which the ZONE ADDRESS RESOLVED bit is set to one, then the source zone group is determined based on the SAS address in the OPEN address frame. In this case, proper management of the zoning expander device extends the zoning capability to beyond the ZPSDS; or
- b) otherwise, the source zone group is a fixed value that is associated with the receiving phy. In this case, non-zoning expander devices are treated in the same manner as end devices and all phys in a non-zoning expander device belong to the zone group specified for the phy that attaches the non-zoning expander device to the ZPSDS.

...

4.8.3.3 Zone permission table

The zone permission table specifies access permission between zone groups. If a bit in the zone permission table is set to one then connection requests shall be permitted between phys in the zone groups. If a bit in the zone permission table is set to zero then connection requests between phys in the zone groups shall be rejected with OPEN_REJECT (ZONE VIOLATION) as described in 4.8.3.5.

The zone permission table structure is shown in table 4.

Table 4 — Zone permission table

Destination zone group d	Source zone group s				
	0	1	2	3 to 7	8 to 127
0	0	1	0	0	0
1	1	1	1	1	1
2	0	1	ZP[2, 2]	Reserved	ZP[s = 8 to 127, 2]
3 to 7	0	1	Reserved	Reserved	Reserved
8 to 127	0	1	ZP[s = 2, d = 8 to 127]	Reserved	ZP[s = 8 to 127, d = 8 to 127]

Note: Shading identifies configurable zone groups.

A ZP[s, d] bit set to one specifies that zone group s has permission to access zone group d. A ZP[s, d] bit set to zero specifies that zone group s does not have permission to access zone group d.

If ZP[s, d] is set to a value, ZP[d,s] shall be set to the same value.

Zone group 0 shall not have access to any other group except zone group 1 (e.g., bits ZP[0, 0] and bits ZP[2 to 127, 0] are set to zero).

Zone group 1 shall have access to all other zone groups (e.g., bits ZP[0 to 127, 1] are set to one).

Regardless of the zoning expander phy attributes, OPEN address frames addressed to a zoning expander device shall be processed as if the destination is zone group 1.

Zone groups 3 through 7 are reserved zone groups. All reserved ZP bits shall be set to zero (e.g., bits ZP[3 to 7, 3 to 127] are set to zero).

If a expander phy has the ZONE ADDRESS RESOLVED METHOD bit set to one, then the zone group of that phy shall be configured to have permission to access all the zone groups to which the SAS addresses in the zoning expander route table for that expander phy have access.

4.8.5 ~~BROADCAST~~ Broadcast processing in a zoning expander device

Editor's Note 4: first pass of changes to 4.8.5 shown for reference:

If a zoning expander device receives a ~~BROADCAST or an SMP ZONED BROADCAST request~~ Broadcast, then the zoning expander device shall process the ~~broadeast~~ Broadcast on each other expander port as follows, where d is the zone group of the destination port:

- a) Transmit an SMP ZONED BROADCAST request from its SMP initiator port if the expander port has its ZONE PARTICIPATING bit set to one and has ~~at least one zone permission bit is set to one in~~:
 - A) at least one ZP[s,d] is set to one where s is any of the zone groups listed in the received SMP ZONED BROADCAST request; or
 - B) the ZP[s,d] is set to one where s is the zone group of the expander port on which the BROADCAST was received;

or

- b) transmit a BROADCAST if the expander port has its ZONE PARTICIPATING bit set to zero and one of the following conditions is met:
 - A) at least one ~~zone permission bit~~ ZP[s,d] is set to one where s is in any of the zone groups listed in the received SMP ZONED BROADCAST request; or

- B) ZP[s,d] is set to one where s is the zone group of the expander port ~~matches the zone group of the expander port~~ on which the BROADCAST was received.

Editor's Note 5: replacement text for section 4.8.5 follows:

The BPP determines the source zone group(s) of the Broadcast as follows:

- a) if the BPP receives a Broadcast Event Notify message (i.e., a zoning expander phy received a BROADCAST) from an expander phy, the Broadcast has a single source zone group set to the zone group of that expander phy; and
- b) if the BPP receives a message from the management device server indicating that it received an SMP ZONED BROADCAST request from an SMP initiator port that has access to zone group 3, the Broadcast has:
 - [Pick one of the following:]
 - A) each of the source zone groups specified in the SMP ZONED BROADCAST request that has a SAS address in the zone expander route table for the expander phy that received the SMP ZONED BROADCAST request.
 - B) the source zone group of the expander phy that received the SMP ZONED BROADCAST request.

The BPP forwards the Broadcast to each expander port other than the one on which the Broadcast was received (i.e., the expander port that received the BROADCAST or SMP ZONED BROADCAST request) if any of the source zone groups have access to the zone group of the expander port;

To forward a Broadcast to an expander port:

- a) if the expander port's ZONE PARTICIPATING bit is set to one, the BPP shall request that the SMP initiator port establish a connection on at least one phy in the expander port to the SMP target port of the attached expander device and transmit an SMP ZONED BROADCAST request specifying the source zone groups.
- b) if the expander port's ZONE PARTICIPATING bit is set to zero, the BPP shall send a Transmit Broadcast message to at least one phy in the expander port, causing it to transmit a BROADCAST.

7.2.5.4 BROADCAST

BROADCASTs are used to ~~notify all SAS ports in a domain of an event~~ notify phys that a Broadcast event has occurred in the SAS domain.

The versions of BROADCAST representing different ~~reasons~~[Broadcast types](#) are defined in table 5.

Table 5 — BROADCAST primitives

Primitive	Description
BROADCAST (CHANGE)	Notification of a configuration change.
BROADCAST (RESERVED CHANGE 0)	Reserved. Processed the same as BROADCAST (CHANGE) by SAS ports (i.e., SAS initiator ports and SAS target ports).
BROADCAST (RESERVED CHANGE 1)	Reserved. Processed the same as BROADCAST (CHANGE) by SAS ports (i.e., SAS initiator ports and SAS target ports).
BROADCAST (SES)	Notification of an asynchronous event from a logical unit with a peripheral device type set to 0Dh (i.e., enclosure services device) (see SPC-3 and SES-2) in the SAS domain.
BROADCAST (EXPANDER)	Notification of an expander event, including: a) a phy event information peak value detector reaching its threshold value; and b) a phy event information peak value detector being cleared. These expander events do not include SAS domain changes.
BROADCAST (RESERVED 2)	Reserved.
BROADCAST (RESERVED 3)	Reserved.
BROADCAST (RESERVED 4)	Reserved.
BROADCAST (CHANGE)	Broadcast (Change)
BROADCAST (RESERVED CHANGE 0)	Broadcast (Reserved Change 0)
BROADCAST (RESERVED CHANGE 1)	Broadcast (Reserved Change 1)
BROADCAST (SES)	Broadcast (SES)
BROADCAST (EXPANDER)	Broadcast (Expander)
BROADCAST (RESERVED 2)	Broadcast (Reserved 2)
BROADCAST (RESERVED 3)	Broadcast (Reserved 3)
BROADCAST (RESERVED 4)	Broadcast (Reserved 4)

BROADCAST shall only be transmitted outside of connections after the link reset sequence has completed.

Editor's Note 6: This assumes 06-xxx has been incorporated, which moves the previous sentence to that location and changes the wording to that shown.

~~When an expander port receives a BROADCAST it shall transmit the same BROADCAST on at least one phy in all other expander ports.~~

~~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.~~

~~BROADCAST (CHANGE) is sent by an expander device to notify SAS initiator ports and other expander devices that a configuration change has occurred. BROADCAST (CHANGE) may also be transmitted by SAS initiator ports. BROADCAST (CHANGE) shall be ignored by SAS target ports.~~

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

~~BROADCAST (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 10.4.3.15).~~

~~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.~~

~~See for details on SAS domain changes. See 10.4.3.3 for details on counting BROADCAST (CHANGE) generation in an expander device. See 4.10 for details on phy event information.~~

See 4.1.11 for definitions of broadcasts.

7.11 SAS domain changes

After power on or receiving ~~BROADCAST (CHANGE)~~[Broadcast \(Change\)](#), the management application client in each SAS initiator port should scan the SAS domain using the discover process (see 4.7) to search for SAS initiator devices, SAS target devices, and expander devices.

The expander device shall ~~transmit~~[originate BROADCAST \(CHANGE\)](#)[Broadcast \(Change\)](#) from at least one phy in each expander port other than the expander port that is the cause for ~~transmitting~~[originating BROADCAST \(CHANGE\)](#)[Broadcast \(Change\)](#).

Expander devices shall ~~transmit~~[originate BROADCAST \(CHANGE\)](#)[Broadcast \(Change\)](#) for the following reasons:

- a) after an expander phy's SP state machine transitions from the SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready state to the SP0:OOB_COMINIT state (see 6.8);

NOTE 1 - This occurs when the expander phy is reset or disabled with the SMP PHY CONTROL function DISABLE, LINK RESET, HARD RESET, or TRANSMIT SATA PORT SELECTION SIGNAL phy operations (see 10.4.3.13) as well as when dword synchronization is unexpectedly lost;

- b) after a virtual phy has been disabled with the SMP PHY CONTROL function DISABLE phy operation or started processing a reset requested by the LINK RESET or HARD RESET phy operations (see 10.4.3.13);
- c) after an expander phy's SP state machine reaches the SATA spinup hold state (see 6.8.7 and 6.10);
- d) after an expander phy's SP state machine receives a COMWAKE Detected message in states SP0:OOB_COMINIT, SP1:OOB_AwaitCOMX, SP3:OOB_AwaitCOMINIT_Sent, or SP4:OOB_COMSAS if the value of the ATTACHED SATA PORT SELECTOR bit in the DISCOVER response is zero prior to receiving the COMWAKE detected message (see 6.8.3 and table 206 in 10.4.3.5);
- e) after an expander phy's SP state machine transitions from the SP1:OOB_AwaitCOMX state to the SP0:OOB_COMINIT state if the value of the ATTACHED SATA PORT SELECTOR bit was one in the DISCOVER response upon entry to SP1:OOB_AwaitCOMX, and if no COMWAKE detected message was received while in SP1:OOB_AwaitCOMX before the transition to SP0:OOB_COMINIT (see 6.8.3.3.2);
- f) after the link reset sequence completes (see 7.9);
- g) after a virtual phy has been enabled or completed processing a reset requested by the SMP PHY CONTROL function LINK RESET or HARD RESET phy operations (see 10.4.3.13);
- h) after a self-configuring expander device has completed configuration and has changed its CONFIGURING bit from one to zero in the SMP REPORT GENERAL function (see 10.4.3.3); [and](#)
- i) after an STP/SATA bridge receives an initial Register - Device to host FIS (see 9.3.1); ~~and~~

[Expander devices shall forward Broadcast \(Change\) for the following reasons:](#)

j) after an expander phy receives BROADCAST (CHANGE).

For a virtual phy, if there is any time after a reset is originated during which connection requests to the attached SAS address result in connection responses of OPEN_REJECT (NO DESTINATION), the expander device shall ~~transmit~~originate the ~~BROADCAST (CHANGE)~~Broadcast (Change) twice, once at the start of the reset (i.e., when the SAS address becomes unavailable) and once at its completion (i.e., when the SAS address becomes available). If there is no such time window, the expander device shall ~~transmit~~originate the ~~BROADCAST (CHANGE)~~Broadcast (Change) once.

~~BROADCAST (CHANGE)~~Broadcast (Change) may be ~~sent~~originated by SAS initiator ports to force other SAS initiator ports and expander ports to re-run the discover process, but should not be sent by SAS target ports.

A SAS initiator port that ~~detects BROADCAST (CHANGE)~~receives Broadcast (Change) shall follow the SAS initiator device rules (see 7.9.2) to discover the topology.

A fanout expander device that ~~detects BROADCAST (CHANGE)~~receives Broadcast (Change) shall follow the fanout device rules (see 7.9.3) to discover the topology.

An edge expander device that ~~detects BROADCAST (CHANGE)~~receives Broadcast (Change) shall follow the edge device rules (see 7.9.4).

See 10.4.3.3 for details on counting ~~BROADCAST (CHANGE) generation~~Broadcast (Change) origination in an expander device.

7.14 SL (link layer for SAS phys) state machines

7.14.1 SL state machines overview

...

Editor's Note 7: In figure 149 (SL state machines part 2), in two places, change BROADCAST (Change) Received to BROADCAST Received.

7.14.4.2 SL_CC0:Idle state

7.14.4.2.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

Upon entry into this state, this state shall send:

- a) a) an Enable Disable SSP (Disable) message to the SSP link layer state machines;
- b) b) an Enable Disable SMP (Disable) message to the SMP link layer state machines;
- c) c) an Enable Disable STP (Disable) message to the STP link layer state machines; and
- d) d) a Connection Closed (Transition to Idle) confirmation to the port layer.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter (see 7.4).

If a BROADCAST Received (Change) message, BROADCAST Received (Reserved Change 0) message, or BROADCAST Received (Reserved Change 1) message is received, this state shall send a Change Received confirmation to the management layer.

If a Transmit Broadcast request is received with any argument, this state shall send a Transmit BROADCAST message with the same argument to the SL transmitter.

7.14.4.3 SL_CC1:ArbSel state

7.14.4.3.1 State description

This state is used to make a connection request.

Upon entry into this state, this state shall:

- 1) request an OPEN address frame be transmitted by sending a Transmit SOAF/Data Dwords/EOAF message to the SL transmitter with the dwords containing the OPEN address frame with its fields set to the arguments received with the Open Connection request;
- 2) initialize and start the Open Timeout timer; and
- 3) request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the SL transmitter.

This state shall ignore OPEN_REJECT Received and OPEN_ACCEPT Received messages from the time a Transmit SOAF/Data Dwords/EOAF message is sent to the SL transmitter until an SOAF/Data Dwords/EOAF Transmitted message is received from the SL transmitter.

If a BROADCAST Received (Change) message, [BROADCAST Received \(Reserved Change 0\) message, or BROADCAST Received \(Reserved Change 1\) message is received](#), is received this state shall send a Change Received confirmation to the management layer.

...

7.15.3 XL0:Idle state

7.15.3.1 State description

This state is the initial state and is the state that is used when there is no connection pending or established.

If a Phy Layer Not Ready confirmation is received, this state shall send a Broadcast Event Notify (Phy Not Ready) request to the BPP.

If a SATA Spinup Hold confirmation is received, this state shall send a Broadcast Event Notify (SATA Spinup Hold) request to the BPP.

If an Enable Disable SAS Link (Enable) message is received, this state shall send a Broadcast Event Notify (Identification Sequence Complete) request to the BPP.

If a SATA Port Selector Change confirmation is received, this state shall send a Broadcast Event Notify (SATA Port Selector Change) request to the BPP.

If a BROADCAST Received message is received, this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., ~~CHANGE-Received~~[Change Received](#)).

If a Transmit Broadcast indication is received, this state shall send a Transmit BROADCAST message to the XL transmitter with an argument specifying the specific type from the Transmit Broadcast indication.

Otherwise, this state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

7.15.8 XL5:Forward_Open state

7.15.8.1 State description

This state is used to transmit an OPEN address frame passed with the transition into this state.

If a BROADCAST Received message is received, this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., ~~CHANGE-Received~~[Change Received](#)).

Upon entry into this state, this state shall send a Transmit OPEN Address Frame message to the XL transmitter with the fields set to the values specified with the transition into this state.

This state shall request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter.

7.15.9 XL6:Open_Response_Wait state

7.15.9.1 State description

This state waits for a response to a transmitted OPEN address frame and determines the appropriate action to take based on the response.

This state shall either:

- a) a) request idle dwords be transmitted by repeatedly sending Transmit Idle Dword messages to the XL transmitter, honoring ALIGN insertion rules for rate matching and clock skew management; or
- b) b) send Transmit Dword messages to the XL transmitter to transmit all dwords received with Forward Dword indications.

If a BROADCAST Received message is received before an AIP Received message is received this state shall send a Broadcast Event Notify request to the BPP with the argument indicating the specific BROADCAST primitive received (e.g., ~~CHANGE Received~~[Change Received](#)).

9.3 STP transport layer

9.3.1 Initial FIS

A SATA device phy transmits a Register - Device to Host FIS after completing the link reset sequence (see G.5 for exceptions to this). The expander device shall update a set of shadow registers with the contents of this FIS and shall not deliver it to any STP initiator port. SMP initiator ports may read the shadow register contents using the SMP REPORT PHY SATA function (see 10.4.3.7). The expander device ~~generates a BROADCAST (CHANGE)~~[originates a Broadcast \(Change\)](#) after receiving the Register - Device to Host FIS (see 7.11).

After the Register - Device to Host FIS is accepted, if the SATA device sends a SATA_X_RDY before an affiliation is established, the expander device shall not send SATA_R_RDY.

10.4.3.2 SMP function request frame format

...

Table 6 — SMP functions (FUNCTION field)

Code	SMP function	Description	Reference
...	...		
85h	ZONED BROADCAST	Transmit the specified BROADCAST Broadcast on the expander ports in the specified zone group(s)	

10.4.3.3 REPORT GENERAL function

...

The EXPANDER CHANGE COUNT field counts the number of ~~BROADCAST (CHANGE)~~[Broadcast \(Change\)](#)s originated by an expander device (see 7.11). SMP target ports in expander devices shall support this field. SMP target ports in other device types (e.g., end devices) shall set the EXPANDER CHANGE COUNT field to 0000h. This field shall be set to 0000h at power on. If the SMP target port has [originated a Broadcast \(Change\)](#) ~~transmitted BROADCAST (CHANGE)~~ for any reason described in 7.11 ~~other than forwarding a BROADCAST (CHANGE)~~ since transmitting a REPORT GENERAL response, it shall increment this field at least once from the value in the previous REPORT GENERAL response. [It shall not increment this field when forwarding a Broadcast \(Change\)](#). This field shall not be incremented when forwarding a ~~BROADCAST (CHANGE)~~[Broadcast \(Change\)](#) from another expander device. This field shall wrap to zero after the maximum value (i.e., FFFFh) has been reached.

NOTE 2 - Application clients that use the EXPANDER 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.

...

A CONFIGURING bit set to one indicates that a self-configuring expander device has not completed configuring its expander route table. 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)Broadcast (Change)** being originated. 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.

...

10.4.3.5 DISCOVER function

...

Table 7 — ATTACHED SATA PORT SELECTOR **and** ATTACHED SATA DEVICE **bits**

ATTACHED SATA PORT SELECTOR bit value ^a _b	ATTACHED SATA DEVICE bit value ^c _d	Description
0	0	Neither a SATA port selector nor a SATA device is attached and ready on the selected phy.
0	1	The attached phy is a SATA device phy. No SATA port selector is present (i.e., the SP state machine did not detect COMWAKE in response to the initial COMINIT, but sequenced through the normal (non-SATA port selector) SATA device OOB sequence).
1	0	The attached phy is a SATA port selector host phy, and either: a) the attached phy is the inactive host phy, or b) the attached phy is the active host phy and a SATA device is either not present or not ready behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT).
1	1	The attached phy is a SATA port selector's active host phy and a SATA device is present behind the SATA port selector (i.e., the SP state machine detected COMWAKE while waiting for COMINIT, timed out waiting for COMSAS, and exchanged COMWAKE with an attached SATA device).
<p>^a The ATTACHED SATA PORT SELECTOR bit is invalid if the NEGOTIATED PHYSICAL LINK RATE field is set to UNKNOWN (i.e., 0h) or DISABLED (i.e., 1h).</p> <p>^b Whenever the ATTACHED SATA PORT SELECTOR bit changes, the phy shall generate a BROADCAST(CHANGE)originate a Broadcast (Change) notification.</p> <p>^c For the purposes of the ATTACHED SATA DEVICE bit, the SATA port selector is not considered a SATA device.</p> <p>^d The ATTACHED SATA DEVICE bit shall be updated at SATA spin-up hold time (see 6.8.7 and 6.10).</p>		

...

The PHY CHANGE COUNT field counts the number of **BROADCAST(CHANGE)Broadcast (Change)**s originated by an expander phy. Expander devices shall support this field. Other device types shall not support this field. This field shall be set to zero at power on. The expander device shall increment this field at least once when it transmits a **BROADCAST(CHANGE)originates a Broadcast (Change)** for any reason described in 7.11 originating from the expander phy **other than forwarding a BROADCAST (CHANGE)and shall not increment this field when forwarding a Broadcast (Change) from the expander phy.**

The expander device is not required to increment the PHY CHANGE COUNT field again unless a DISCOVER response is transmitted. This field shall not be incremented when forwarding a BROADCAST (CHANGE) from another expander device. The PHY CHANGE COUNT field shall wrap to zero after the maximum value (i.e., FFh) has been reached.

NOTE 3 - Application clients that use the PHY CHANGE COUNT field should read it often enough to ensure that it does not increment a multiple of 256 times between reading the field.

10.4.3.9 REPORT PHY EVENT INFORMATION

...

If the phy event information source is a peak value detector, the PEAK VALUE DETECTOR THRESHOLD field contains the value of the peak value detector that ~~triggers a BROADCAST (EXPANDER)~~ [causes a Broadcast \(Expander\) to be originated](#) (see 7.2.5.4). If the phy event information source is not a peak value detector, then the PEAK VALUE DETECTOR THRESHOLD field is reserved.

...

10.4.3.11 ZONED BROADCAST function

The ZONED BROADCAST function requests that the specified ~~BROADCAST~~ [Broadcast](#) be ~~transmitted~~ [forwarded](#) on all the [expander](#) ports that are in ~~one or more specified~~ [zone groups accessible to the specified zone groups or that have expander routing tables with SAS addresses in zone groups accessible to the specified zone groups](#), with the exception of the [expander](#) port on which the ZONED BROADCAST function was received (see 4.8.5). This SMP function shall be supported by SMP target ports in zoning expander devices (see 4.8). Other SMP target ports shall not support this SMP function.

Table 8 defines the request format.

Table 8 — ZONED BROADCAST 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 TYPE				
5	Reserved								
6	Reserved								
7	NUMBER OF BROADCAST SOURCE ZONE GROUPS								
Zone group list									
8	BROADCAST SOURCE ZONE GROUP (first)								
	...								
	BROADCAST SOURCE ZONE GROUP (first)								
	PAD (if needed)								
n - 4									
n - 3	(MSB)	CRC							
n									(LSB)

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

The FUNCTION field shall be set to 85h.

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

The BROADCAST TYPE field specifies the type of ~~BROADCAST~~Broadcast that shall be ~~transmitted~~forwarded ~~on-~~each expander port in the specified zone groups and is defined in Table 9.

Table 9 — BROADCAST TYPE field

Code	Description
000b	BROADCAST (CHANGE)
001b	BROADCAST (RESERVED CHANGE 0)
010b	BROADCAST (RESERVED CHANGE 1)
011b	BROADCAST (SES)
100b	BROADCAST (EXPANDER)
101b	BROADCAST (RESERVED 2)
110b	BROADCAST (RESERVED 3)
111b	BROADCAST (RESERVED 4)
<u>000b</u>	<u>Broadcast (Change)</u>
<u>001b</u>	<u>Broadcast (Reserved Change 0)</u>
<u>010b</u>	<u>Broadcast (Reserved Change 1)</u>
<u>011b</u>	<u>Broadcast (SES)</u>
<u>100b</u>	<u>Broadcast (Expander)</u>
<u>101b</u>	<u>Broadcast (Reserved 2)</u>
<u>110b</u>	<u>Broadcast (Reserved 3)</u>
<u>111b</u>	<u>Broadcast (Reserved 4)</u>

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

~~Each ZONE GROUP field specifies one zone group in which the specified BROADCAST~~Broadcast shall be originated, once on each expander port in that zone group. Each BROADCAST SOURCE ZONE GROUP field specifies a source zone group for the Broadcast. The expander device forwards the Broadcast to each destination zone group accessible to that source zone group. Zone group values between 128 and 255, inclusive, are reserved.

The PAD field contains zero, one, two, or three bytes set to 00h such that the total length of the SMP request is a multiple of four.

The CRC field is defined in 10.4.3.1.

Table 10 defines the response format.

Table 10 — ZONED BROADCAST response

Byte\Bit	7	6	5	4	3	2	1	0	
0	SMP FRAME TYPE (41h)								
1	FUNCTION (85h)								
2	FUNCTION RESULT								
3	RESPONSE 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 FUNCTION RESULT field is defined in 10.4.3.2.

The RESPONSE LENGTH field shall be set to 00h.

The CRC field is defined in 10.4.3.2.

10.4.3.13 PHY CONTROL function

...

Table 11 — PHY OPERATION field (part 1 of 2)

Code	Operation	Description
00h	NOP	No operation.
01h	LINK RESET	<p>If the specified phy is not a virtual phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the specified phy is a virtual phy, perform an internal reset and enable the specified phy. See 7.11 for BROADCAST (CHANGE)Broadcast (Change) requirements related to this phy operation in an expander device.</p> <p>Any affiliation (see 7.17.5) shall continue to be present. The phy shall bypass the SATA spinup hold state, if implemented (see 6.8.3.9).</p> <p>The SMP response shall be returned without waiting for the link reset to complete.</p>
02h	HARD RESET	<p>If the specified phy is not a virtual phy, perform a link reset sequence (see 4.4) on the specified phy and enable the specified phy. If the attached phy is a SAS phy or an expander phy, the link reset sequence shall include a hard reset sequence (see 4.4.2). If the attached phy is a SATA phy, the phy shall bypass the SATA spinup hold state. See 7.11 for BROADCAST (CHANGE)Broadcast (Change) requirements related to this phy operation in an expander device.</p> <p>If the specified phy is a virtual phy, perform an internal reset and enable the specified phy.</p> <p>Any affiliation (see 7.17.5) shall be cleared.</p> <p>The SMP response shall be returned without waiting for the hard reset to complete.</p>
03h	DISABLE	Disable the specified phy (i.e., stop transmitting valid dwords and receiving dwords on the specified phy). The LINK RESET and HARD RESET operations may be used to enable the phy. See 7.11 for BROADCAST (CHANGE) Broadcast (Change) requirements related to this phy operation in an expander device.
04h	Reserved	
05h	CLEAR ERROR LOG	Clear the error log counters reported in the REPORT PHY ERROR LOG function (see 10.4.3.6) for the specified phy.
06h	CLEAR AFFILIATION	Clear an affiliation (see 7.17.5) from the STP initiator port with the same SAS address as the SMP initiator port that opened this SMP connection. If there is no such affiliation, the SMP target port shall return a function result of SMP FUNCTION FAILED in the response frame.

Table 11 — PHY OPERATION field (part 2 of 2)

Code	Operation	Description
07h	TRANSMIT SATA PORT SELECTION SIGNAL	<p>This function shall only be supported by phys in an expander device.</p> <p>If the expander phy incorporates an STP/SATA bridge and supports SATA port selectors, the phy shall transmit the SATA port selection signal (see 6.6) which causes the SATA port selector to select the attached phy as the active host phy and make its other host phy inactive. See 7.11 for BROADCAST <u>Broadcast (Change)</u> requirements related to this phy operation in an expander device.</p> <p>Any affiliation (see 7.17.5) shall be cleared.</p> <p>If the expander phy does not support SATA port selectors, then the SMP target port shall return a function result of PHY DOES NOT SUPPORT SATA.</p> <p>If the expander phy supports SATA port selectors but is attached to a SAS phy or an expander phy, the SMP target port shall return a function result of SMP FUNCTION FAILED.</p>
08h	CLEAR STP I_T NEXUS LOSS	<p>The STP I_T NEXUS LOSS OCCURRED bit in the REPORT PHY SATA function (see 10.4.3.7) shall be set to zero.</p>
All others	Reserved	

...