To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 17 October 2005
Subject: 05-306r1 SAS-2 STP connection time limits and STP/SMP I_T nexus loss

Revision history
Revision 0 (27 August 2005) First revision
Revision 1 (17 October 2005) Incorporated comments from September 2005 SAS protocol WG.

Related documents
sas2r00 - Serial Attached SCSI 2 revision 0
05-305r0 - SAS-2 Maximum SMP connection time (Rob Elliott, HP) - incorporated into sas2r00

Overview
SAS-1.1 provides several tools to help keep SSP target ports (and, indirectly, SSP initiator ports) from keeping connections open too long and causing congestion in the fabric:

a) the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page - limits the time of a connection (even if it is transferring frames)
b) the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page - limits the time of a connection that is not transferring frames
c) the I_T NEXUS LOSS TIME field in the Disconnect-Reconnect mode page - limits the number of OPEN retries in response to OPEN_REJECT (NO DESTINATION)
d) numerous 1 ms timeouts in the link layer - ACK/NAK timeout, Credit timeout, DONE Response timeout, etc.

SAS-1.1 does not define any such tools for STP target ports (nor, indirectly, STP initiator ports and SMP initiator ports), however.

One approach is for the STP target port to close the connection between every frame. This may not be the best approach for all situations - if the STP initiator port or the SATA device can send data frames back-to-back, this adds a lot of overhead. Another approach is to let the STP initiator port always initiate closing the connection. There is no rule guaranteeing it must do so, however; if both sides wait on the other, the connection will never be closed.

Adding fields such as those defined for SSP will let STP target ports, STP initiator ports, and SMP initiator ports make better decisions about connection management and behave more consistently with the SSP target ports in the fabric.

Proposal:

a) Define a new SMP CONFIGURE GENERAL function to handle device-wide settings.
b) Define byte 3 in each SMP request and response frame as a REQUEST/RESPONSE LENGTH field indicating the number of dwords that follow, not counting the CRC field. As new fields are added to the frames, the frame length will simply be increased. Software must tolerate extra bytes it does not understand at the end of a frame. A value of 00h in the field means the frame has the length defined in SAS-1.1.
c) Define an STP MAXIMUM CONNECT TIME LIMIT field (in the SMP REPORT GENERAL/CONFIGURE GENERAL functions) to instruct the STP target port(s) how long before they must try to close a connection (regardless of activity).
d) Define an STP BUS INACTIVITY TIME LIMIT field (in the SMP REPORT GENERAL/CONFIGURE GENERAL functions) to instruct the STP target port(s) how long they may keep a connection open while not sending a frame (e.g. waiting for the SATA device to send another frame)
e) Define an STP SMP I_T NEXUS LOSS TIME field (in the SMP REPORT GENERAL/CONFIGURE GENERAL functions) to direct the handling of OPEN_REJECT (NO DESTINATION) connection responses:
    A) instruct the STP target port(s), if any, how long to keep trying to open an STP initiator port
    B) instruct the SMP initiator port in a self-configuring expander device (if any) how long to keep trying to open an SMP target port.

The timers enforcing these values are implemented by the port layer, like their SSP counterparts (an STP/SATA bridge and an SMP initiator port in an expander device conceptually include a port layer).

Suggested changes

2.3 References under development

ISO/IEC 14776-972 *ATA Attachment - 8 Architecture Model (ATA8-AAM)* (T13/1700-D)

4.5 I_T nexus loss

When a SAS port receives OPEN_REJECT (NO DESTINATION), OPEN_REJECT (PATHWAY BLOCKED), or an open connection timeout occurs in response to a connection request, it shall retry the connection request until:

a) the connection is established;
~~b) for SSP target ports, the time indicated by the I_T NEXUS LOSS field in the Protocol Specific Port mode page (see 10.2.7.2) expires; or~~
~~c) for STP ports, SMP ports, or SSP initiator ports, a vendor specific I_T nexus loss time expires;~~
d) the I_T nexus loss timer, if any, expires (see 8.2.2.1, 10.2.7.2, and 10.4.3.7);

~~An SSP initiator port should retry the connection request for the time indicated by the I_T NEXUS LOSS field in the Protocol Specific Port mode page (see 10.2.7.2) for the SSP target port to which it is trying to establish a connection.~~

I_T nexus loss is handled by the port layer state machines (see 8.2.2.3).

If the I_T nexus loss timer expires in an SSP port, then the port shall send a Nexus Loss event notification to the SCSI application layer (see 10.2.5) and the SCSI device shall perform the actions defined for I_T nexus loss in SAM-3.

Editor's Note 1: reordered the preceding two paragraphs (not shown as changes)

If the I_T nexus loss timer expires in an STP target port, then the port shall abort any commands for the lost STP initiator port and clear the affiliation, if any (see 7.17.5). If the STP target port is in an STP/SATA bridge, the STP/SATA bridge shall originate a link reset sequence to the SATA device.

If the I_T nexus loss timer expires in an STP initiator port, then the port shall consider any commands for the lost STP target port to be completed with an error (see ATA8-ACS).

If the I_T nexus loss timer expires in an SMP initiator port, then the port shall stop attempting to establish connections to the lost SMP target port.

If the I_T nexus loss timer expires in an initiator port, then a management application client should cause a link reset sequence on the phy(s) attached to the lost target port (e.g, if directly attached, the phys in the initiator port; if attached via expander device(s), the phys in the expander device closest to and attached to the target port).

Editor's Note 2: Link reset sequence, not hard reset sequence, is specifically chosen above.

If the target has multiple SSP target ports, the fact that this SSP target port went down should not affect the other SSP target ports. A hard reset would affect all of them. Running an I_T NEXUS RESET task management function is advisable, though.

If the target is an STP target port, then the only difference is that hard reset breaks an affiliation. If some other initiator snuck in, wiping out its affiliation is not appropriate.

7.17.5 Affiliations

Coherent access to the SATA task file registers shall be provided for each STP initiator port. STP target ports that do not track all commands by the STP initiator ports' SAS addresses shall implement affiliations to provide coherency. STP target ports that track all commands by the STP initiator ports' SAS addresses shall not implement affiliations.

An affiliation is a state entered by an STP target port where it refuses to accept connection requests from STP initiator ports other than the one that has established an affiliation.

An STP target port that supports affiliations shall establish an affiliation whenever it accepts a connection request. When an affiliation is established, the STP target port shall reject all subsequent connection requests from other STP initiator ports with OPEN_REJECT (STP RESOURCES BUSY).

An STP target port shall maintain an affiliation until any of the following occurs:

a) power on;
b) the SAS target device receives an SMP PHY CONTROL request specifying the phy with the affiliation and specifying a phy operation of HARD RESET (see 10.4.3.10) from any SMP initiator port;
c) the SAS target device receives an SMP PHY CONTROL request specifying the phy with the affiliation and specifying a phy operation of TRANSMIT SATA PORT SELECTION SIGNAL (see 10.4.3.10) from any SMP initiator port;
d) the SAS target device receives an SMP PHY CONTROL request specifying the phy with the affiliation and specifying a phy operation of CLEAR AFFILIATION (see 10.4.3.10) from the same SAS initiator port that has the affiliation.

   If a connection is already established to the STP target port on one phy while an SMP PHY CONTROL request specifying a phy operation of CLEAR AFFILIATION is processed by an SMP target port on another phy, the affiliation shall be cleared and the STP target port shall respond to new connection attempts with:

   A) AIP (WAITING ON CONNECTION) and/or OPEN_REJECT (RETRY), if the STP target port is in an expander device; or
   B) OPEN_REJECT (RETRY), if the STP target port is in a SAS device;

   rather than OPEN_REJECT (STP RESOURCES BUSY);

e) a connection to the phy with the affiliation is closed with CLOSE (CLEAR AFFILIATION); ~~or~~
f) the STP target port is part of a STP/SATA bridge and a link reset sequence is begun on the SATA physical link that was not requested by an SMP PHY CONTROL request specifying the phy and specifying a phy operation of LINK RESET (see 10.4.3.10); or
g) an I_T nexus loss occurs (see 4.5, 8.2.2.1, and 10.4.3.7).

An affiliation established when a command is transmitted shall be maintained until all frames for the command have been delivered. An STP initiator port implementing command queuing (see SATAII-EXT) shall maintain an affiliation while any commands are outstanding to avoid confusing SATA devices, which only understand one SATA host. STP initiator ports may keep affiliations for longer tenures, but this is discouraged.

An STP target port that implements affiliations shall implement one affiliation per STP target port. Multiple phys on the same STP target port shall use the same affiliation. Support for affiliations is indicated in the SMP REPORT PHY SATA function response (see 10.4.3.7).

8.2.2 PL_OC (port layer overall control) state machine

8.2.2.1 PL_OC state machine overview

A PL_OC state machine:

a) receives requests from the SSP, SMP, and STP transport layers;
b) sends messages to the PL_PM state machine;
c) receives messages from the PL_PM state machine;
d) selects frames to transmit;
e) selects phys on which to transmit frames;

f) receives confirmations from the link layer;
g) sends confirmations to the transport layer;
h) has Arbitration Wait Time timers; and
i) has I_T Nexus Loss timers.

This state machine consists of the following states:

a) PL_OC1:Idle (see 8.2.2.2) (initial state); and
b) PL_OC2:Overall_Control (see 8.2.2.3).

After power on this state machine shall start in the PL_OC1:Idle state.

The PL_OC state machine shall maintain:

a) a pool of pending Tx Frame messages for each destination SAS address; and
b) as many pending Tx Open message slots as there are phys in the port.

The PL_OC state machine shall maintain the timers listed in table 1.

Table 1 — PL_OC state machine timers

Timer	Maximum number of timers	Initial value
I_T Nexus Loss timer	One per destination SAS address	Depending on the protocol used by the port: a) for SSP target ports, the value in the I_T NEXUS LOSS TIME field in the Protocol-Specific Port mode page (see 10.2.7.2); b) for SSP initiator ports, the value in the I_T NEXUS LOSS TIME field in the Protocol-Specific Port mode page for the SSP target port with that destination SAS address (see 10.2.7.2); c) for STP target ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function (see 10.4.3.xx); d) for STP initiator ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function for the STP target port with that destination SAS address; or e) for SMP initiator ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function.
Arbitration Wait Time timer	One per pending Tx Open message	0000h, a vendor specific value less than 8000h (see 7.12.3), or the value received with a Retry Open message.

Figure 1 shows the PL_OC state machine.

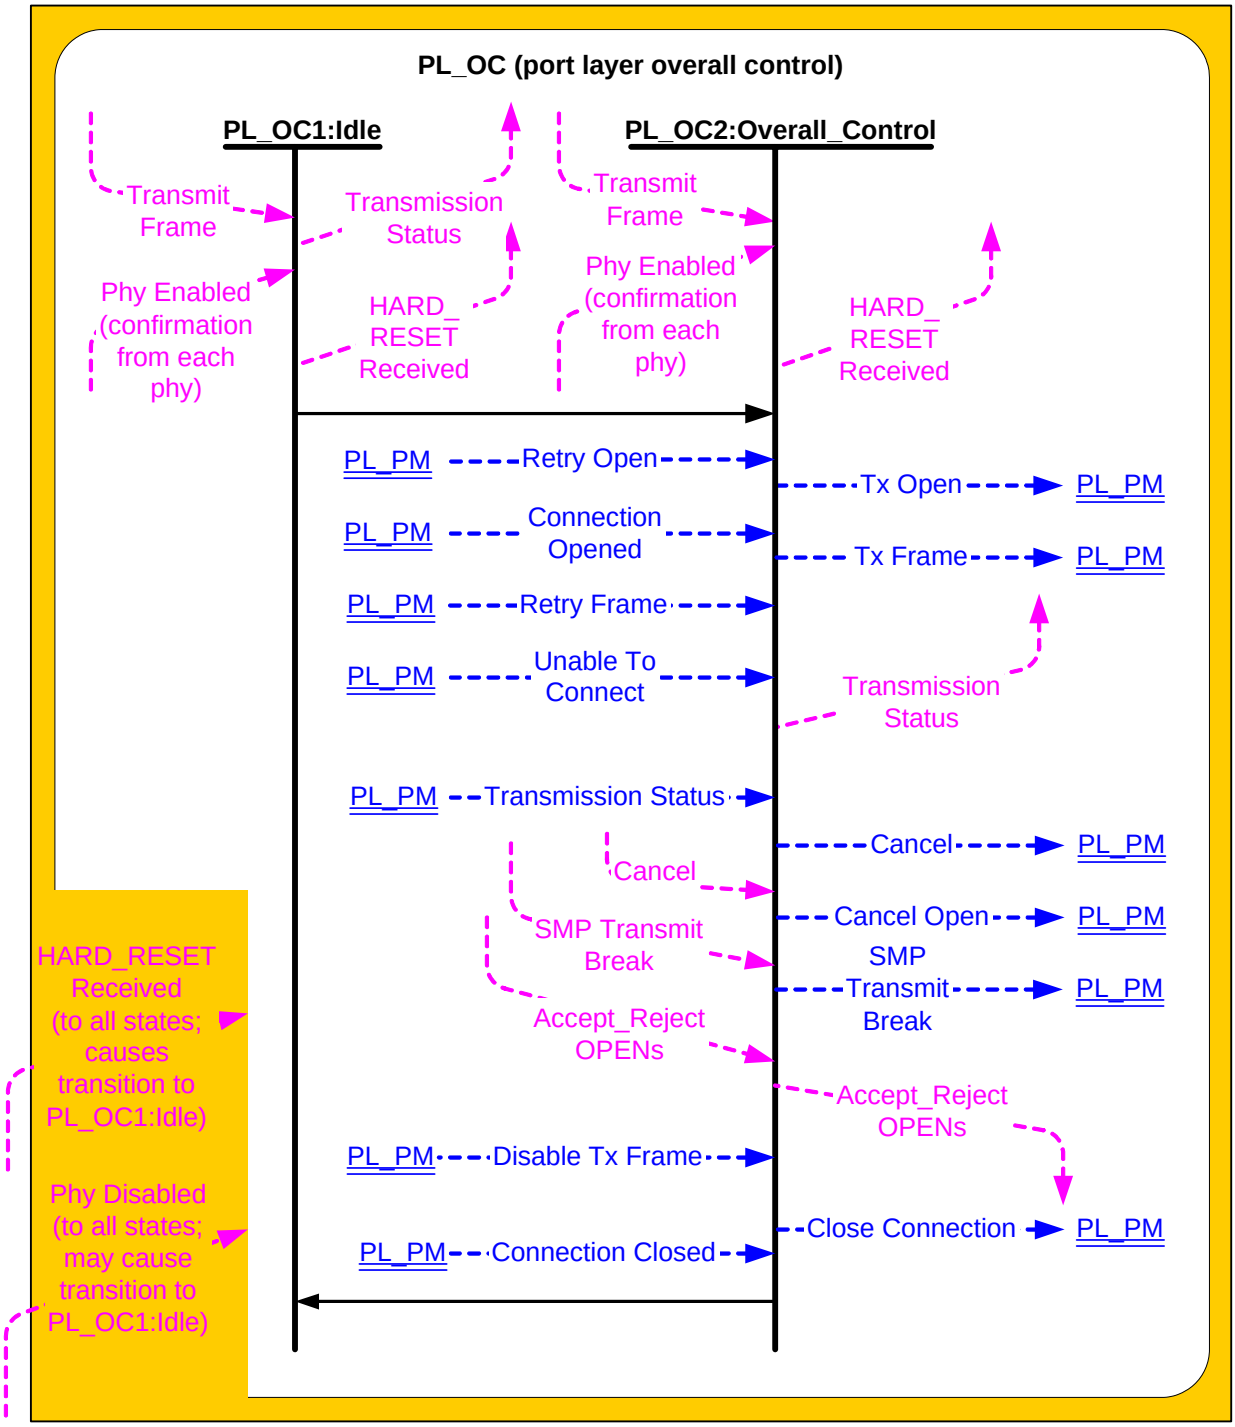

Figure 1 — PL_OC (port layer overall control) state machine

8.2.2.2 PL_OC1:Idle state

8.2.2.2.1 PL_OC1:Idle state description

This state is the initial state of the PL_OC state machine.

If this state receives a HARD_RESET Received confirmation, then this state shall send a HARD_RESET Received confirmation to the transport layer.

If this state receives a Transmit Frame request, then this state shall send a Transmission Status (No Phys In Port) confirmation to the transport layer.

If an I_T Nexus Loss timer expires for a destination SAS address, this state shall perform the following:

a) delete the I_T Nexus Loss timer for the SAS address;
b) send a Transmission Status (I_T Nexus Loss) confirmation for each pending Tx Frame message for the SAS address; and
c) discard each pending Tx Frame message for the SAS address and any corresponding pending Tx Open messages.

If the port is an STP target port or an STP initiator port, the port shall handle all pending commands as described in 4.5.

8.2.2.2.2 Transition PL_OC1:Idle to PL_OC2:Overall_Control

This transition shall occur after a Phy Enabled confirmation is received for at least one phy assigned to the port.

8.2.2.3 PL_OC2:Overall_Control state

8.2.2.3.1 PL_OC2:Overall_Control state overview

This state may receive Transmit Frame requests from the transport layers (i.e., SSP and SMP) and Retry frame messages from PL_PM state machines. This state shall create a pending Tx Frame message for each received Transmit Frame request and Retry Frame message. There may be more than one pending Tx Frame message at a time for each SSP transport layer. There shall be only one pending Tx Frame message at a time for each SMP transport layer.

This state selects PL_PM state machines through which connections are established. This state shall only attempt to establish connections through PL_PM state machines whose phys are enabled. In a vendor-specific manner, this state selects PL_PM state machines on which connections are established to transmit frames. This state shall receive a response to a message from a PL_PM state machine before sending another message to that PL_PM state machine.

This state also:

a) receives connection management requests from the transport layers;
b) sends connection management messages to PL_PM state machines;
c) receives connection management messages from PL_PM state machines; and
d) sends connection management confirmations to the transport layers.

After receiving a Transmit Frame request for a destination SAS address for which there is no connection established and for which no I_T Nexus Loss timer has been created, this state shall create an I_T Nexus Loss timer for that SAS address if:

~~a) the protocol is SSP;~~
~~b) this state machine is in an SSP target port;~~
~~c) the Protocol-Specific Port mode page is implemented by the SSP target port; and~~
~~d) the I_T nexus loss time is not 0000h.~~
a) the protocol is SSP, the port is an SSP target port, the Protocol-Specific Port mode page is implemented, and the I_T NEXUS LOSS TIME field in the Disconnect-Reconnect mode page (see 10.2.7.1) is not set to 0000h;
b) the protocol is STP, the port is an STP target port, and the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function is not set to 0000h; or
c) the protocol is SMP, the port is an SMP initiator port, and the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function is not set to 0000h.

This state may create an I_T Nexus Loss timer for that SAS address if:

a) the protocol is SSP and the port is ~~this state machine is in~~ an SSP initiator port; or ~~then this state may create an I_T Nexus Loss timer for the SAS address. If a state machine in an SSP initiator port and creates an I_T Nexus Loss timer, then the state machine should use the value in the I_T NEXUS LOSS~~

~~TIME field in the Protocol-Specific Port mode page for the SSP target port (see 10.2.7.2) as the initial value for its I_T Nexus Loss timer.~~

b) the protocol is STP and the port is an STP initiator port.

~~Other SAS ports may detect an I_T nexus loss in a vendor-specific manner.~~

When this state creates an I_T Nexus Loss timer it shall:

a) initialize the I_T Nexus Loss timer as specificied in table 1 (see 8.2.2.1); and
b) not start the I_T Nexus Loss timer.

Editor's Note 3: moved the When... paragraph down a few paragraphs from its current location

---

If there are no pending Tx Frame messages for a destination SAS address and an I_T Nexus Loss timer has been created for that destination SAS address, then this state shall delete the I_T Nexus Loss timer for that destination SAS address.

If this state receives a HARD_RESET Received confirmation, then this state shall discard all pending Tx Frame messages and delete all I_T Nexus Loss timers and send a HARD_RESET Received confirmation to the transport layer.

8.2.2.3.2 PL_OC2:Overall_Control state establishing connections

This state receives Phy Enabled confirmations indicating when a phy is available.

This state receives Retry Open messages from a PL_PM state machine.

This state creates pending Tx Open messages based on pending Tx Frame messages and Retry Open messages. Pending Tx Open messages are sent to a PL_PM state machine as Tx Open messages.

If this state receives a Retry Open (Retry) message, then this state shall process the Retry Open message.

If this state receives a Retry Open (No Destination) or a Retry Open (Open Timeout Occurred) message and an I_T Nexus Loss timer has not been created for the destination SAS address (e.g., an SSP target port does not support the I_T NEXUS LOSS TIME field in the Protocol-Specific Port mode page or the field is set to 0000h), then this state shall process the Retry Open message as either a Retry Open message or an Unable To Connect message. This selection is vendor-specific.

If this state receives a Retry Open (Pathway Blocked) message and an I_T Nexus Loss timer has not been created for the destination SAS address, then this state shall process the Retry Open message.

If this state receives a Retry Open (No Destination), Retry Open (Open Timeout Occurred), or Retry Open (Pathway Blocked) message, and an I_T Nexus Loss timer has been created for the destination SAS address with an initial value of FFFFh, then this state shall process the Retry Open message (i.e., the Retry Open message is never processed as an Unable to Connect message).

If this state receives a Retry Open (No Destination) or a Retry Open (Open Timeout Occurred) message, an I_T Nexus Loss timer has been created for the destination SAS address, and there is no connection established with the destination SAS address, then this state shall check the I_T Nexus Loss timer, and:

a) if the I_T Nexus Loss timer is not running and the I_T nexus loss time is not set to FFFFh, then this state shall start the timer;
b) if the I_T Nexus Loss timer is running, then this state shall not stop the timer; and
c) if the I_T Nexus Loss timer has expired, then this state shall process the Retry Open message as if it were an Unable To Connect message (see 8.2.2.3.4).

If this state receives a Retry Open (Pathway Blocked) message, an I_T Nexus Loss timer has been created for the destination SAS address, and there is no connection established with the destination SAS address, then this state shall check the I_T Nexus Loss timer, and:

a) if the I_T Nexus Loss timer is running, then this state shall not stop the timer; and

b) if the I_T Nexus Loss timer has expired, then this state shall process the Retry Open message as if it were an Unable To Connect message (see 8.2.2.3.4).

If this state receives a Retry Open (Retry) and an I_T Nexus Loss timer is running for the destination SAS address, then this state shall:

a) stop the I_T Nexus Loss timer (if the timer has been running); and
b) initialize the I_T Nexus Loss timer.

This state shall create a pending Tx Open message if:

a) this state has a pending Tx Frame message or has received a Retry Open message;
b) this state has fewer pending Tx Open messages than the number of PL_PM state machines (i.e., the number of phys in the port);
c) there is no pending Tx Open message for the destination SAS address; and
d) there is no connection established with the destination SAS address.

This state may create a pending Tx Open message if:

a) this state has a pending Tx Frame message, or this state has received a Retry Open message and has not processed the message by sending a confirmation; and
b) this state has fewer pending Tx Open messages than the number of PL_PM state machines.

This state shall have no more pending Tx Open messages than the number of PL_PM state machines.

If this state receives a Retry Open message and there are pending Tx Frame messages for which pending Tx Open messages have not been created, then this state should create a pending Tx Open message from the Retry Open message.

If this state does not create a pending Tx Open message from a Retry Open message (e.g., the current number of pending Tx Open messages equals the number of phys), then this state shall discard the Retry Open message. This state may create a new pending Tx Open message at a later time for the pending Tx Frame message that resulted in the Retry Open message.

If this state receives a Retry Open (Opened By Destination) message and the initiator port bit and protocol arguments match those in the Tx Open messages that resulted in the Retry Open message, then this state may discard the Retry Open message and use the established connection to send pending Tx Frame messages as Tx Frame messages to the destination SAS address. If this state receives a Retry Open (Opened By Destination) message, then, if this state has a pending Tx Open slot available, this state may create a pending Tx Open message from the Retry Open message.

NOTE 1 - If a connection is established by another port as indicated by a Retry Open (Opened By Destination) message, credit may not be granted for frame transmission. In this case this state may create a pending Tx Open message from a Retry Open message in order to establish a connection where credit is granted.

This state shall send a pending Tx Open message as a Tx Open message to a PL_PM state machine that has an enabled phy and does not have a connection established. If there is more than one pending Tx Open message, this state should send a Tx Open message for the pending Tx Open message that has been pending for the longest time first.

If this state creates a pending Tx Open message from one of the following messages:

a) a Retry Open (Opened By Destination);
b) a Retry Open (Opened By Other);
c) a Retry Open (Collided); or
d) a Retry Open (Pathway Blocked),

then this state shall:

a) create an Arbitration Wait Time timer for the pending Tx Open message;
b) set the Arbitration Wait Time timer for the pending Tx Open message to the arbitration wait time argument from the Retry Open message; and
c) start the Arbitration Wait Time timer for the pending Tx Open message.

When a pending Tx Open message is sent to a PL_PM state machine as a Tx Open message, the Tx Open message shall contain the following arguments to be used in an OPEN address frame:

a) initiator port bit from the Transmit Frame request;
b) protocol from the Transmit Frame request;
c) connection rate from the Transmit Frame request;
d) initiator connection tag from the Transmit Frame request;
e) destination SAS address from the Transmit Frame request;
f) source SAS address from the Transmit Frame request;
g) pathway blocked count; and
h) arbitration wait time.

If this state creates a pending Tx Open message from one of the following:

a) a Transmit Frame request;
b) a Retry Open (No Destination) message;
c) a Retry Open (Open Timeout Occurred) message; or
d) a Retry Open (Retry) message,

then this state shall:

a) set the pathway blocked count argument in the Tx Open message to zero; and
b) set the arbitration wait time argument in the Tx Open message to zero or a vendor-specific value less than 8000h (see 7.12.3).

If a pending Tx Open message was created as the result this state receiving a Retry Open (Pathway Blocked) message, then this state shall set the pathway blocked count argument in the Tx Open message to the value of the pathway blocked count argument received with the message plus one, unless the pathway blocked count received with the argument is FFh.

If a pending Tx Open message was created as the result of this state receiving one of the following:

a) a Retry Open (Opened By Destination) message;
b) a Retry Open (Opened By Other) message;
c) a Retry Open (Collided) message; or
d) a Retry Open (Pathway Blocked) message;

then this state shall set the arbitration wait time argument in the Tx Open message to be the value from the Arbitration Wait Time timer created as a result of the Retry Open message.

After this state sends a Tx Open message, this state shall discard the pending Tx Open message from which the Tx Open messages was created. After this state discards a pending Tx Open message, this state may create a new pending Tx Open message.

If this state receives a Connection Opened message and the initiator port bit and protocol arguments match those in any pending Tx Frame messages, then this state may use the established connection to send pending Tx Frame messages as Tx Frame messages to the destination SAS address.

8.2.2.3.3 PL_OC2:Overall_Control state connection established

If this state receives a Connection Opened message or a Retry Open (Opened By Destination) message for a SAS address, and an I_T Nexus Loss timer has been created for the SAS address, then this state shall:

a) stop the I_T Nexus Loss timer for the SAS address, if the timer has been running; and
b) initialize the I_T Nexus Loss timer.

8.2.2.3.4 PL_OC2:Overall_Control state unable to establish a connection

If this state receives a Retry Open (No Destination), Retry Open (Open Timeout Occurred), or Retry Open (Pathway Blocked) message and the I_T Nexus Loss timer for the SAS address has expired, then this state shall perform the following:

a) delete the I_T Nexus Loss timer for the SAS address;
b) discard the Retry Open message;

c) send a Transmission Status (I_T Nexus Loss) confirmation for the pending Tx Frame message from which the Retry Open message resulted;
d) discard the pending Tx Frame message from which the Retry Open message resulted;
e) if this state has any pending Tx Frame messages with the same destination SAS address and protocol as the Retry Open message, and this state has not sent a Tx Open message to a PL_PM state machine for the messages, then this state shall send a Transmission Status (I_T Nexus Loss) confirmation for each pending Tx Frame message and discard the pending Tx Frame messages and any corresponding pending Tx Open messages; and
f) if this state has any pending Tx Frame messages with the same destination SAS address and protocol as the Retry Open message, and this state has sent a Tx Open message to a PL_PM state machine for a message, then this state shall send a Cancel Open message to each PL_PM state machine to which it has sent a Tx Open message. After receiving an Unable To Connect (Cancel Acknowledge) message from a PL_PM state machine in response to the Cancel Open message, then this state shall send a Transmission Status (I_T Nexus Loss) confirmation for each pending Tx Frame message and discard the pending Tx Frame messages and any corresponding pending Tx Open messages.

If this state receives a Retry Open (No Destination), Retry Open (Open Timeout Occurred), or Retry Open (Pathway Blocked) message and processes it as an Unable To Connect message, or this state receives an Unable To Connect message, then this state shall send a Transmission Status confirmation as defined in table 2.

Table 2 — Confirmations from Unable To Connect or Retry Open messages

Message received	Confirmation to be sent to transport layer
Retry Open (No Destination)	Transmission Status (I_T Nexus Loss) if the I_T Nexus Loss timer for the SAS address has expired, or Transmission Status (No Destination) if it has not
Retry Open (Open Timeout Occurred)	Transmission Status (I_T Nexus Loss) if the I_T Nexus Loss timer for the SAS address has expired, or Transmission Status (Open Timeout Occurred) if it has not
Retry Open (Pathway Blocked)	Transmission Status (I_T Nexus Loss) if the I_T Nexus Loss timer for the SAS address has expired
Unable to Connect (Bad Destination)	Transmission Status (Bad Destination)
Unable To Connect (Break Received)	Transmission Status (Break Received)
Unable To Connect (Connection Rate Not Supported)	Transmission Status (Connection Rate Not Supported)
Unable To Connect (Port Layer Request)	Transmission Status (Cancel Acknowledge)
Unable To Connect (Protocol Not Supported)	Transmission Status (Protocol Not Supported)
Unable To Connect (STP Resources Busy)	Transmission Status (STP Resources Busy)
Unable To Connect (Wrong Destination)	Transmission Status (Wrong Destination)

If this state receives an Unable To Connect (Connection Rate Not Supported), Unable To Connect (Protocol Not Supported), or Unable To Connect (STP Resources Busy) message and an I_T Nexus Loss timer is running for the SAS address, then this state shall:

a) stop the I_T Nexus Loss timer, if the timer has been running; and
b) initialize the I_T Nexus Loss timer.

This state shall discard the pending Tx Frame message for which the Transmission Status confirmation was sent.

8.2.2.3.5 PL_OC2:Overall_Control state connection management

If this state receives an Accept_Reject Opens request, then this state shall send an Accept_Reject Opens message to all phys in the port.

If this state receives an SMP Transmit Break request, then this state shall send an SMP Transmit Break message to the PL_PM state machine associated with the corresponding SMP transport state machine. If there is no PL_PM state machine associated with the request, the PM_OC state shall ignore the request.

If this state receives one of the following:

a) a Connection Closed (Close Timeout) message;
b) a Connection Closed (Break Requested) message; or
c) a Connection Closed (Break Received) message,

then this state shall not send a Tx Open or Tx Frame message to the PL_PM state machine that sent the message until this state receives a Connection Closed (Transition to Idle) message from that PL_PM state machine.

If this state receives a Connection Closed (Normal) message or a Connection Closed (Transition to Idle) message indicating that a connection with a destination SAS address is no longer open and this state has pending Tx Open messages, then this state may send a Tx Open message to the PL_PM state machine that sent the Connection Closed message.

If this state is in a wide SSP port, then this state shall not reject an incoming connection request on one phy because it has an outgoing connection request on another phy.

If this state is in an SSP port, there are no pending Tx Frame messages for a destination SAS address with which a PL_PM state machine has established a connection, and the connection was established by a message from this state, then this state should send a Close Connection message to the PL_PM state machine.

If this state is in an SSP port, has no pending Tx Frame messages for a destination SAS address with which a PL_PM state machine has established a connection, and the connection was established by the destination, then this state may wait a vendor-specific time and then shall send a Close Connection message to the PL_PM state machine.

If this state has received a Disable Tx Frame message from a PL_PM state machine, then this state should send a Close Connection message to the PL_PM state machine.

NOTE 2 - The PL_PM state machine sends a Close Connection request to the link layer upon receipt of a Close Connection message or on expiration of the Bus Inactivity Time Limit timer (see 8.2.3.4.1).

8.2.2.3.6 PL_OC2:Overall_Control state frame transmission

In order to prevent livelocks, If this state is in a wide SSP port, has multiple connections established, and has a pending Tx Frame message, then this state shall send at least one Tx Frame message to a PL_PM state machine before sending a Close Connection message to the PL_PM state machine.

After this state receives a Connection Opened message from a PL_PM state machine, this state selects pending Tx Frame messages for the destination SAS address with the same initiator port bit and protocol arguments, and, as an option, the same connection rate argument, and sends the messages to the PL_PM state machine as Tx Frame messages.

This state may send a Tx Frame message to any PL_PM state machine that has established a connection with the destination SAS address when the initiator port bit and protocol arguments match those in the Tx Frame message.

After this state sends a Tx Frame message to a PL_PM state machine, it shall not send another Tx Frame message to that PL_PM state machine until it receives a Transmission Status (Frame Transmitted) message.

This state may send a Tx Frame message containing a COMMAND frame for a destination SAS address to a PL_PM state machine while waiting for one of the following messages for Tx Frame messages containing COMMAND frames for the same destination SAS address from different PL_PM state machines:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);
c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

This state shall not send a Tx Frame message containing a TASK frame for a task that only affects an I_T_L_Q nexus (e.g., an ABORT TASK or QUERY TASK task management function (see SAM-3)) until this state has received one of the following messages for each Tx Frame message with the same I_T_L_Q nexus:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);
c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

This state shall not send a Tx Frame message containing a TASK frame for a task that only affects an I_T_L nexus (e.g., an ABORT TASK SET, CLEAR TASK SET, CLEAR ACA, or LOGICAL UNIT RESET task management function (see SAM-3)) until this state has received one of the following messages for each Tx Frame message with the same I_T_L nexus:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);
c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

This state shall not send a Tx Frame message containing a TASK frame for a task that only affects an I_T nexus until this state has received one of the following messages for each Tx Frame message with the same I_T nexus:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);
c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

Once this state has sent a Tx Frame message containing a DATA frame to a PL_PM state machine, this state shall not send a Tx Frame message containing a DATA frame with the same I_T_L_Q to another PL_PM state machine until this state has received one of the following messages for each Tx Frame message containing a DATA frame for the same I_T_L_Q nexus:

a) Transmission Status (ACK Received);
b) Transmission Status (NAK Received);
c) Transmission Status (ACK/NAK Timeout); or
d) Transmission Status (Connection Lost Without ACK/NAK).

Read DATA frames and write DATA frames for the same I_T_L_Q nexus may be transmitted and received simultaneously on the same or different phys.

If this state is in an SMP initiator port, then this state shall send the Tx Frame message containing the SMP REQUEST frame to the PL_PM state machine on which the connection was established for the Tx Open message. If this state is in an SMP target port, then this state shall send the Tx Frame message containing the SMP RESPONSE frame to the PL_PM state machine on which the connection was established for the Tx Open message. See 7.18 for additional information about SMP connections.

Characteristics of STP connections are defined by SATA (also see 7.17).

The following arguments shall be included with the Tx Frame message:

a) the frame to be transmitted; and
b) Balance Required or Balance Not Required.

A Balance Not Required argument shall only be included if:

a) the request was a Transmit Frame (Non-Interlocked) request (i.e., the request included a DATA frame); and
b) the last Tx Frame message sent to this PL_PM state machine while this connection has been established was for a DATA frame having the same logical unit number and tag value as the DATA frame in this Tx Frame message.

If a Balance Not Required argument is not included in the Tx Frame message, then a Balance Required argument shall be included.

If this state receives a Disable Tx Frames message from a PL_PM state machine, then this state should send no more Tx Frame messages to that state machine until a new connection is established.

8.2.2.3.7 PL_OC2:Overall_Control state frame transmission cancellations

Cancel requests cause this state to cancel previous Transmit Frame requests. A Cancel request includes the following arguments:

a) destination SAS address; and
b) tag.

If this state receives a Cancel request and has not already sent a Tx Frame message for the Transmit Frame request to a PL_PM state machine for the Transmit Frame request specified by the Cancel request, then this state shall:

a) discard all Transmit Frame requests for the specified destination SAS address and tag; and
b) send a Transmission Status (Cancel Acknowledge) confirmation to the transport layer.

If this state receives a Cancel request and has already sent a Tx Frame message to a PL_PM state machine for the Transmit Frame request specified by the Cancel request, then this state shall send a Cancel message to the PL_PM state machine to which the Tx Frame message was sent. The Cancel message shall include the tag.

8.2.2.3.8 Transition PL_OC2:Overall_Control to PL_OC1:Idle

This transition shall occur after:

a) sending a HARD_RESET Received confirmation to the transport layer; or
b) a Phy Disabled confirmation is received from all of the link layers in the port;

8.2.3 PL_PM (port layer phy manager) state machine

8.2.3.1 PL_PM state machine overview

A PL_PM state machine:

a) receives messages from the PL_OC state machine;
b) sends requests to the link layer;
c) receives confirmations from the link layer;
d) sends confirmations to the transport layer;
e) sends messages to PL_OC state machine;
f) has an Arbitration Wait Time timer;
g) may have a Bus Inactivity Time Limit timer; and
h) may have Maximum Connect Time Limit timer.

This state machine consist of the following states:

a) PL_PM1:Idle (see 8.2.3.2) (initial state);
b) PL_PM2:Req_Wait (see 8.2.3.3);
c) PL_PM3:Connected (see 8.2.3.4); and
d) PL_PM4:Wait_For_Close (see 8.2.3.5).

After power on this state machine shall start in the PL_PM1:Idle state.

The PL_PM state machine shall maintain the timers listed in table 3.

Table 3 — PL_PM state machine timers

Timer	Initial value
Arbitration Wait Time timer	The arbitration wait time argument from a Retry Open message (see 8.2.2.3.1).
Bus Inactivity Time Limit timer	Depending on the protocol used by the port: a) for SSP target ports, the value in the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page; or b) for STP target ports, the value in the STP BUS INACTIVITY TIME LIMIT field in the SMP REPORT GENERAL function.
Maximum Connect Time Limit timer	Depending on the protocol used by the port: a) for SSP target ports, the value in the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1); or b) for STP target ports, the value in the STP MAXIMUM CONNECT TIME LIMIT field in the SMP REPORT GENERAL function.

Figure 2 shows part 1 of the PL_PM state machine.

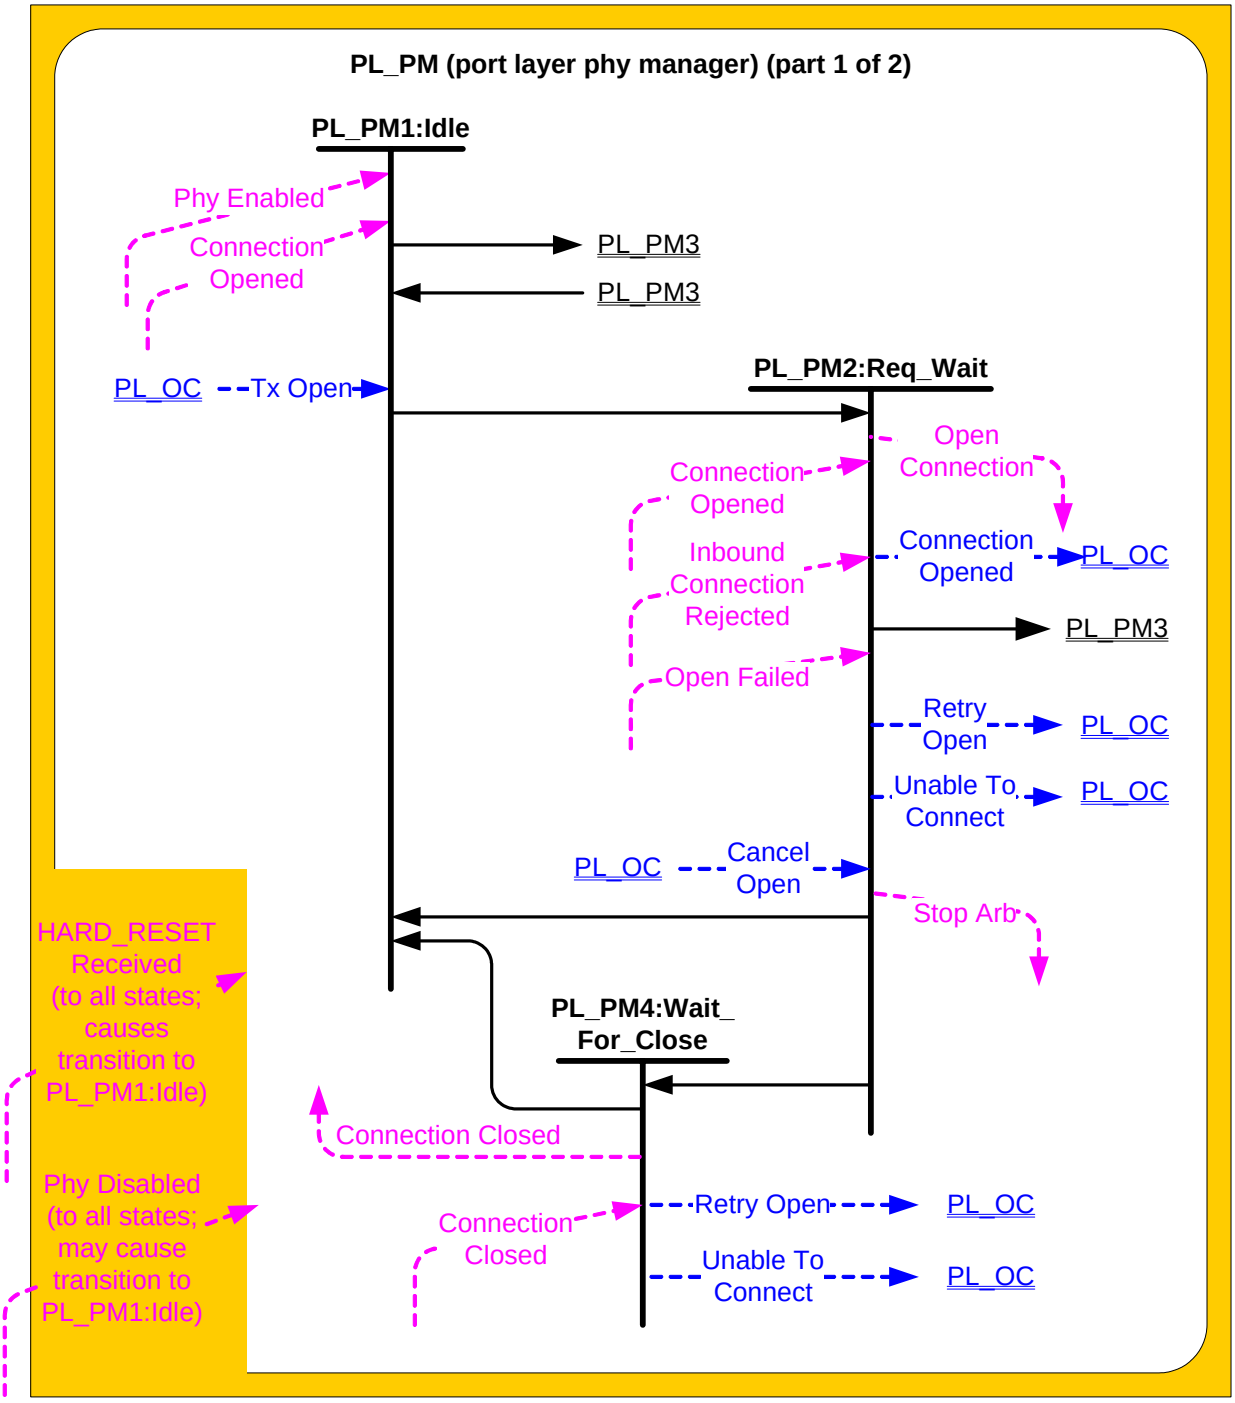

Figure 2 — PL_PM (port layer phy manager) state machine (part 1)

Figure 3 shows part 2 of the PL_PM state machine.

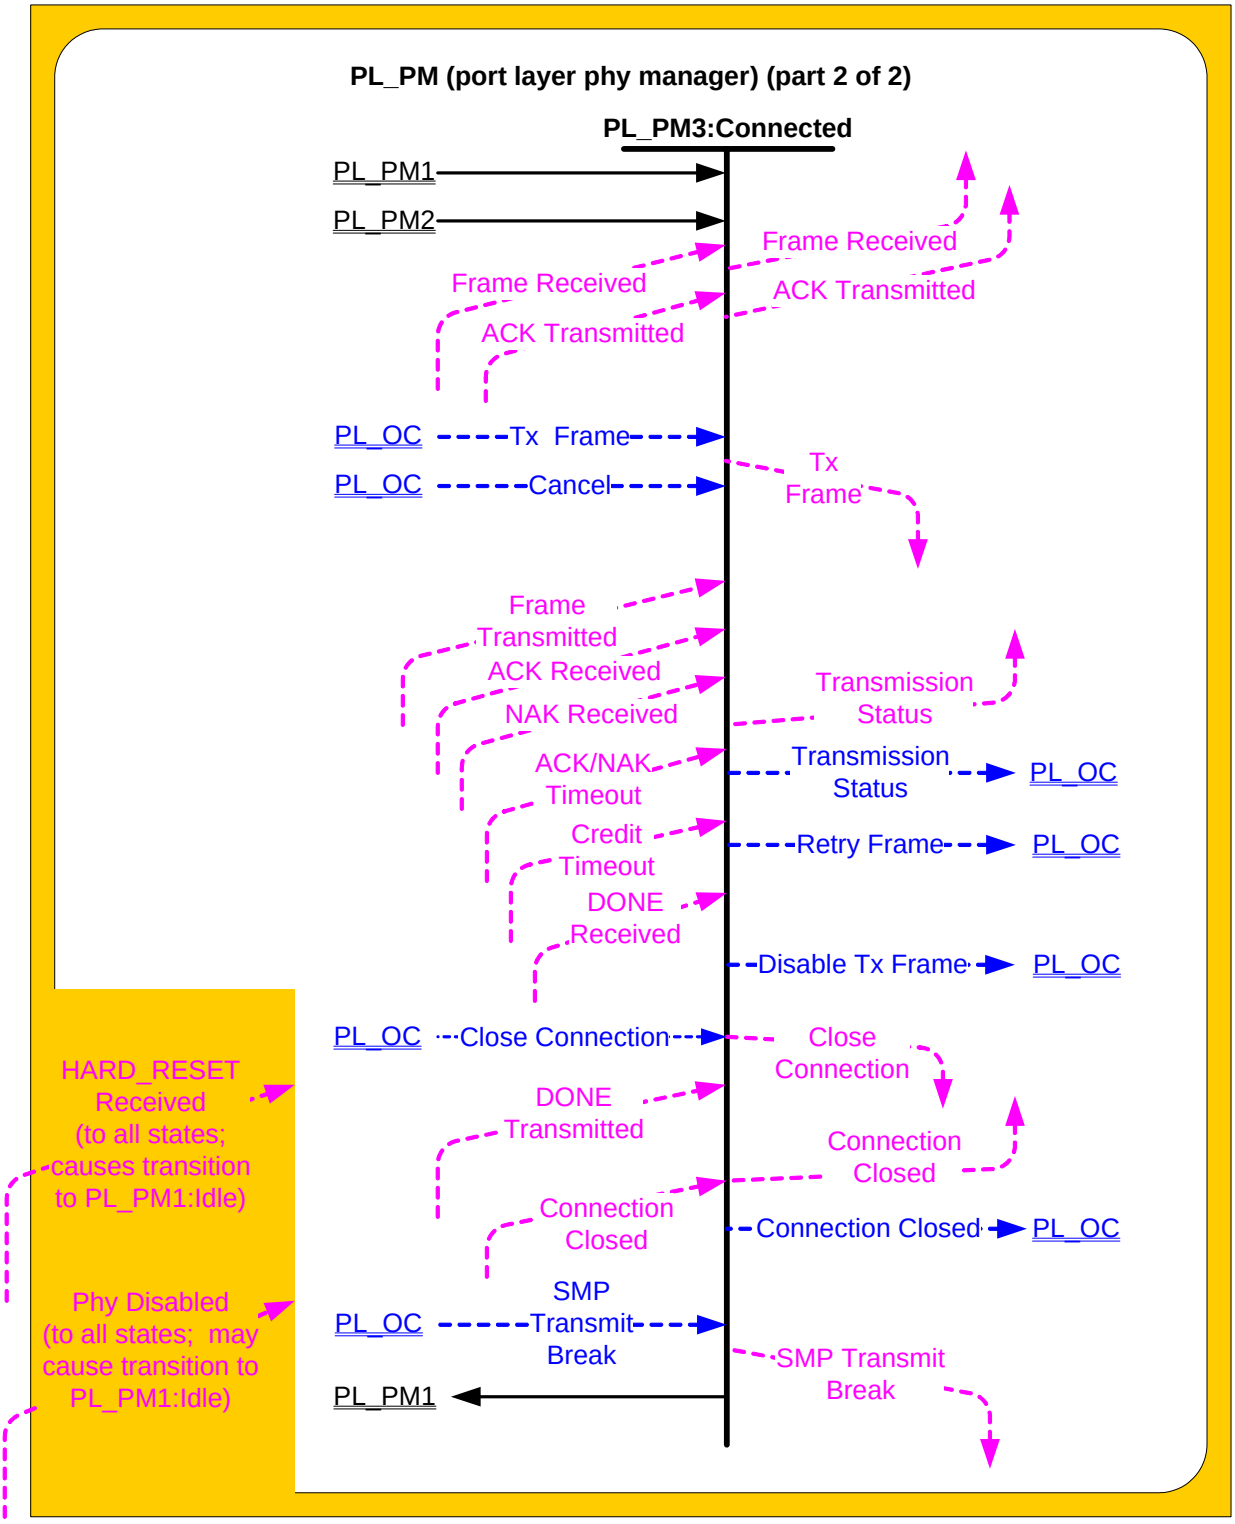

Figure 3 — PL_PM (port layer phy manager) state machine (part 2)

8.2.3.2 PL_PM1:Idle state

8.2.3.2.1 PL_PM1:Idle state description

This is the initial state of the PL_PM state machine.

8.2.3.2.2 Transition PL_PM1:Idle to PL_PM2:Req_Wait

This transition shall occur after:

a) a Phy Enabled confirmation is received; and
b) a Tx Open message is received.

8.2.3.2.3 Transition PL_PM1:Idle to PL_PM3:Connected

This transition shall occur after a Connection Opened confirmation is received.

8.2.3.3 PL_PM2:Req_Wait state

8.2.3.3.1 PL_PM2:Req_Wait state overview

This state sends an Open Connection request to the link layer and waits for a confirmation. This state sends and receives connection management messages to and from the PL_OC state machine.

If this state receives a HARD_RESET Received confirmation, then this state shall terminate all operations.

8.2.3.3.2 PL_PM2:Req_Wait establishing a connection

Upon entry into this state, this state shall:

a) create an Arbitration Wait Time timer;
b) initialize the Arbitration Wait Time timer to the arbitration wait time argument received with the Tx Open message;
c) start the Arbitration Wait Time timer; and
d) send an Open Connection request to the link layer.

The Open Connection request shall contain the following arguments from the Tx Open message to be used in an OPEN address frame:

a) initiator port bit;
b) protocol;
c) connection rate;
d) initiator connection tag;
e) destination SAS address;
f) source SAS address;
g) pathway blocked count; and
h) arbitration wait time.

8.2.3.3.3 PL_PM2:Req_Wait connection established

If this state receives a Connection Opened confirmation, then this state shall send a Connection Opened message to the PL_OC state machine.

If this state receives a Connection Opened confirmation and the confirmation was not in response to an Open Connection request from this state (i.e., the connection was established in response to an OPEN address frame from another device), then this state shall discard any Open Connection request and send a Retry Open message to the PL_OC state machine. If the Connection Opened confirmation was from the destination of the Open Connection request, then this state shall send a Retry Open (Opened By Destination) message. If the Connection Opened confirmation was from a destination other than the destination of the Open Connection request, then this state shall send a Retry Open (Opened By Other) message.

A Retry Open (Opened By Destination) or Retry Open (Opened By Other) message shall contain the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;

f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

8.2.3.3.4 PL_PM2:Req_Wait unable to establish a connection

If this state receives one of the Open Failed confirmations listed in table 4, then this state shall send either a Retry Open message or an Unable To Connect message to the PL_OC state machine.

Table 4 defines the message to be sent for each Open Failed confirmation.

Table 4 — Messages from Open Failed confirmations

Confirmation received	Message to be sent to PL_OC
Open Failed (Pathway Blocked)	Retry Open (Pathway Blocked)
Open Failed (Retry)	Retry Open (Retry)
Open Failed (No Destination)	Retry Open (No Destination)
Open Failed (Bad Destination)	Unable To Connect (Bad Destination)
Open Failed (Connection Rate Not Supported)	Unable To Connect (Connection Rate Not Supported)
Open Failed (Protocol Not Supported)	Unable To Connect (Protocol Not Supported)
Open Failed (STP Resources Busy)	Unable To Connect (STP Resources Busy)
Open Failed (Wrong Destination)	Unable To Connect (Wrong Destination)

If this state receives an Inbound Connection Rejected confirmation after sending an Open Connection request, then this state shall discard the Open Connection request and send a Retry Open (Collided) message to the PL_OC state machine.

A Retry Open message shall include the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;
f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count argument set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

An Unable To Connect message shall include the following arguments:

a) initiator connection tag set to the value received with the Tx Open message;
b) destination SAS address set to the value received with the Tx Open message; and
c) source SAS address set to the value received with the Tx Open message.

8.2.3.3.5 PL_PM2:Req_Wait connection management

If this state receives a Cancel Open message and a Connection Opened confirmation has not been received, then this state shall send a Stop Arb request to the link layer.

8.2.3.3.6 Transition PL_PM2:Req_Wait to PL_PM1:Idle

This transition shall occur after:

a) a Retry Open message is sent to the PL_OC state machine;
b) an Unable To Connect message is sent to the PL_OC state machine;

c) all operations have been terminated after a HARD_RESET Received confirmation is received; or
d) a Phy Disabled confirmation is received.

8.2.3.3.7 Transition PL_PM2:Req_Wait to PL_PM3:Connected

This transition shall occur after a Connection Opened confirmation is received.

8.2.3.3.8 Transition PL_PM2:Req_Wait to PL_PM4:Wait_For_Close

This transition shall occur after one of the following confirmations is received:

a) an Open Failed (Open Timeout Occurred);
b) an Open Failed (Break Received); or
c) an Open Failed (Port Layer Request).

8.2.3.4 PL_PM3:Connected state

8.2.3.4.1 PL_PM3:Connected state description

If:

a) the protocol for the connection is SSP, the port is an SSP target port, the Disconnect-Reconnect mode page is implemented, and the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is not set to zero;
b) the protocol for the connection is SMP and the port is an SMP target port, or
c) the protocol for the connection is STP, the port is an STP target port, and the STP MAXIMUM CONNECT TIME LIMIT field in the SMP REPORT GENERAL function (see 10.x.x) is not set to zero,

then, upon entry into this state, this state shall:

a) create a Maximum Connect Time Limit timer;
b) initialize the Maximum Connect Time Limit timer as specified in table 115 (see 8.2.3.1); and
c) start the Maximum Connect Time Limit timer.

If:

a) the protocol for the connection is SSP, the port is an SSP initiator port, and the MAXIMUM CONNECT TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) for the destination SSP target port is not set to zero; or
b) the protocol for the connection is STP, the port is an STP initiator port, and the STP MAXIMUM CONNECT TIME LIMIT field in the SMP REPORT GENERAL function (see 10.x.x) for the destination STP target port is not set to zero,

then, upon entry into this state, this state may:

a) create a Maximum Connect Time Limit timer;
b) initialize the Maximum Connect Time Limit timer as specified in table 115 (see 8.2.3.1); and
c) start the Maximum Connect Time Limit timer.

~~Other SAS ports may implement a Maximum Connect Time Limit timer in a vendor specific manner.~~

Editor's Note 4: Specifics are provided in the new "may" section above

---

If:

a) the protocol for the connection is SSP, the port is an SSP target port, the Disconnect-Reconnect mode page is implemented, and the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) is ~~set to a non-zero value~~ not set to zero; or
b) the protocol for the connection is STP, the port is an STP target port, and the STP BUS INACTIVITY TIME LIMIT field in the SMP REPORT GENERAL function (see 10.x.x) is not set to zero,

then, upon entry into this state, this state shall:

a) create a Bus Inactivity Time Limit timer;
b) initialize the Bus Inactivity Time Limit timer as specified in table 115 (see 8.2.3.1); and
c) start the Bus Inactivity Time Limit timer.

If:

a) the protocol for the connection is SSP, the port is an SSP initiator port, and the BUS INACTIVITY TIME LIMIT field in the Disconnect-Reconnect mode page (see 10.2.7.1) for the destination SSP target port is not set to zero; or
b) the protocol for the connection is STP, the port is an STP initiator port, and the STP BUS INACTIVITY TIME LIMIT field in the SMP REPORT GENERAL function (see 10.x.x) for the destination STP target port is not set to zero,

then, upon entry into this state, this state may:

a) create a Bus Inactivity Time Limit timer;
b) initialize the Bus Inactivity Time Limit timer as specified in table 115 (see 8.2.3.1); and
c) start the Bus Inactivity Time Limit timer.

~~Other SAS ports may implement a Bus Inactivity Time Limit timer in a vendor-specific manner.~~

Editor's Note 5: Specifics are provided in the new "may" section above

---

If a Bus Inactivity Time Limit timer has been created and:

a) the connection is SSP or SMP and this state receives a Tx Frame message; or
b) the connection is STP and the phy is not both transmitting and receiving SATA_SYNC,

then this state shall:

a) stop the Bus Inactivity Time Limit timer, if it is running; and
b) initialize the Bus Inactivity Time Limit timer as specified in table xx (see 8.2.3.1).

Editor's Note 6: This timer is not intended to clean up errors where SYNC/SYNC never appears - it is just to prevent bus hogging from compliant devices. Thus, for STP, waiting for SYNC/SYNC should suffice. The Maximum Connect Time Limit would normally kick in next.

If this state receives a Tx Frame message, this state shall send a Tx Frame request to the link layer. The following arguments from the Tx Frame message shall be included with the Tx Frame request:

a) the frame to be transmitted; and
b) if this port is an SSP port, Balance Required or Balance Not Required.

For STP connections, this state connects the STP transport layer to the STP link layer.

If a Bus Inactivity Time Limit timer expires:

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer;
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation; and
c) if the connection is STP, then this state shall send a Close Connection request to the link layer.

Editor's Note 7: Per earlier text, the timer is stopped when there is no SYNC/SYNC, so the wait is

not mentioned in c).

If a Maximum Connect Time Limit timer expires:

a) if the connection is SSP and there is no Tx Frame request outstanding (i.e., this state is not waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer;
b) if the connection is SSP and there is a Tx Frame request outstanding (i.e., this state is waiting for an ACK Received or NAK Received confirmation), then this state shall send a Close Connection request to the link layer after receiving an ACK Received or NAK Received confirmation;
c) if the connection is SMP, then this state shall send an SMP Transmit Break request to the link layer; and
d) if the connection is STP, then this state shall send a Close Connection request to the link layer after the phy is both transmitting and receiving SATA_SYNC; and

Editor's Note 8: This timer is not intended to clean up errors where SYNC/SYNC never appears - it is just to prevent bus hogging from compliant devices. Thus, for STP, waiting for SYNC/SYNC should suffice. Another timer may be appropriate to clean up connections that are not reaching SYNC/SYNC (started after this timer expires?). It would have to use BREAKs and link resets to clean up.

If this state receives a Tx Frame message after sending a Close Connection request but before receiving a Connection Closed confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

If this state receives a Frame Received confirmation, then this state shall send a Frame Received confirmation to the transport layer. The confirmation shall include the arguments received with the confirmation (e.g., the frame).

If this state receives an ACK Transmitted confirmation, then this state shall send an ACK Transmitted confirmation to the transport layer including the tag of the frame that was ACKed.

If this state receives a Frame Transmitted confirmation, then this state shall send a Transmission Status (Frame Transmitted) confirmation to the transport layer.

If this state receives an ACK Received confirmation, then this state shall send a Transmission Status (ACK Received) confirmation to the transport layer.

If this state receives a NAK Received confirmation, then this state shall send a Transmission Status (NAK Received) confirmation to the transport layer.

If this state receives an ACK/NAK Timeout confirmation, then this state shall send a Transmission Status (ACK/NAK Timeout) confirmation to the transport layer.

If this state receives a Cancel message, then this state shall:

a) discard all Tx Frame requests for the specified tag;
b) send a Transmission Status (Cancel Acknowledge) confirmation to the transport layer including the destination SAS address and the tag as arguments; and
c) discard any subsequent confirmations for previous Tx Frame requests sent for the tag.

If this state receives a Close Connection message from the PL_OC state machine, then this state shall send a Close Connection request to the link layer.

If this state receives one of the following:

a) a Connection Closed (Normal) confirmation;
b) a Connection Closed (Close Timeout) confirmation;
c) a Connection Closed (Break Requested) confirmation;
d) a Connection Closed (Break Received) confirmation; or

e) a Connection Closed (Transition to Idle) confirmation,

then this state shall send a Connection Closed message to the PL_OC state machine including the argument received with the confirmation.

If this state receives a Connection Closed (Transition to Idle) confirmation after receiving:

a) a Connection Closed (Break Received) confirmation; or
b) a Connection Closed (Break Requested) confirmation,

then this state shall send a Transmission Status (Break Received) confirmation to the transport layer.

If this state receives a Connection Closed (Normal) confirmation, a Connection Closed (Transition to Idle) confirmation, or a Phy Disabled confirmation after sending a Transmission Status (Frame Transmitted) confirmation, but before this state receives an ACK Received or NAK Received confirmation, then this state shall send a Transmission Status (Connection Lost Without ACK/NAK) confirmation to the transport layer.

If this state receives a Connection Closed (Normal) confirmation, a Connection Closed (Transition to Idle) confirmation, or a Phy Disabled confirmation after sending a Tx Frame request but before receiving a Frame Transmitted confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

If this state receives a Connection Closed confirmation during an SMP connection, this state shall send a Connection Closed confirmation to the transport layer.

If this state receives a Credit Timeout confirmation, then this state shall send a Retry Frame message to the PL_OC state machine.

A Retry Frame message shall include the following arguments from the Tx Frame message:

a) initiator port bit;
b) protocol;
c) connection rate;
d) initiator connection tag;
e) destination SAS address;
f) source SAS address; and
g) frame.

After this state receives a DONE Received (Normal) or DONE Received (Credit Blocked) confirmation, if it does not receive a Tx Frame message within 1 ms, then this state shall send a Disable Tx Frames message to the PL_OC state machine.

If this state receives a DONE Received (ACK/NAK Timeout) or DONE Transmitted confirmation, then this state shall send a Disable Tx Frames message to the PL_OC state machine.

If this state receives an SMP Transmit Break message, then this state shall send an SMP Transmit Break request to the link layer.

If this state receives a HARD_RESET Received confirmation, then this state machine shall terminate all operations.

8.2.3.4.2 Transition PL_PM3:Connected to PL_PM1:Idle

This transition shall occur after:

a) a Connection Closed (Transition to Idle) message is sent to the PL_OC state machine; or
b) all operations are terminated after a HARD_RESET Received confirmation is received.

8.2.3.5 PL_PM4:Wait_For_Close state

8.2.3.5.1 PL_PM4:Wait_For_Close state description

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered as the result of the PL_PM2:Req_Wait state receiving an Open Failed (Open Timeout Occurred) confirmation, then this state

shall send a Retry Open (Open Timeout Occurred) message to the PL_OC state machine. The Retry Open message shall include the following arguments:

a) initiator port bit set to the value received with the Tx Open message;
b) protocol set to the value received with the Tx Open message;
c) connection rate set to the value received with the Tx Open message;
d) initiator connection tag set to the value received with the Tx Open message;
e) destination SAS address set to the value received with the Tx Open message;
f) source SAS address set to the value received with the Tx Open message;
g) pathway blocked count argument set to the value received with the Tx Open message; and
h) arbitration wait time set to the value of the Arbitration Wait Time timer.

If this state receives a Connection Closed confirmation and the connection request was for an SMP connection, this state shall send a Connection Closed confirmation to the transport layer.

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered after the PL_PM2:Req_Wait state received an Open Failed (Port Layer Request) confirmation (i.e., as the result of the PL_PM2:Req_Wait state sending a Stop Arb request), then this state shall send an Unable to Connect (Port Layer Request) message to the PL_OC state machine.

After receiving a Connection Closed (Transition to Idle) confirmation, if this state was entered as the result of the PL_PM2:Req_Wait state receiving an Open Failed (Break Received) confirmation, then this state shall send an Unable to Connect (Break Received) message to the PL_OC state machine.

The Unable To Connect message shall include the following arguments:

a) initiator connection tag set to the value received with the Tx Open message;
b) destination SAS address set to the value received with the Tx Open message; and
c) source SAS address set to the value received with the Tx Open message.

If this state receives a HARD_RESET Received confirmation, then this state shall terminate all operations.

8.2.3.5.2 Transition PL_PM4:Wait_For_Close to PL_PM1:Idle

This transition shall occur after:

a) a Retry Open or Unable To Connect message is sent to the PL_OC state machine; or
b) all operations are terminated after a HARD_RESET Received confirmation is received.

10.4.3.1 SMP function request frame format

...

Table 5 — SMP functions (FUNCTION field)

Code	SMP function	Description	~~Request frame size (in bytes)~~	~~Response frame size (in bytes)~~	Reference
00h	REPORT GENERAL	Return general information about the device	~~8~~	~~32~~	10.4.3.2
01h	REPORT MANUFACTURER INFORMATION	Return vendor and product identification	~~8~~	~~64~~	10.4.3.3
02h	READ GPIO REGISTER	See SFF-8485			
03h - 0Fh	Reserved for general SMP input functions				
10h	DISCOVER	Return information about the specified phy	~~16~~	~~56~~	10.4.3.4
11h	REPORT PHY ERROR LOG	Return error logging information about the specified phy	~~16~~	~~32~~	10.4.3.5
12h	REPORT PHY SATA	Return information about a phy currently attached to a SATA phy	~~16~~	~~60~~	10.4.3.7
13h	REPORT ROUTE INFORMATION	Return route table information	~~16~~	~~44~~	10.4.3.7
14h - 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	~~20~~	~~8~~	10.4.3.x
~~80h -~~ 81h	Reserved for general SMP output functions				
82h	WRITE GPIO REGISTER	See SFF-8485			
83h - 8Fh	Reserved for general SMP output functions				
90h	CONFIGURE ROUTE INFORMATION	Change route table information	~~44~~	~~8~~	10.4.3.8
91h	PHY CONTROL	Request actions by the specified phy	~~44~~	~~8~~	10.4.3.10
92h	PHY TEST FUNCTION	Request a test function by the specified phy	~~44~~	~~8~~	10.4.3.10
92h - 9Fh	Reserved for phy-based SMP output functions				
A0h - BFh	Reserved for SMP output functions				
C0h - FFh	Vendor specific				

10.4.3.2 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 SMP target ports.

Table 6 defines the request format.

Table 6 — 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 contains the number of dwords that follow, not including the CRC field (i.e., 0).

The CRC field is defined in 10.4.3.1.

Table 7 defines the response format.

Table 7 — REPORT GENERAL response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (00h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (09h)							
4	(MSB)	EXPANDER CHANGE COUNT						
5								(LSB)
6	(MSB)	EXPANDER ROUTE INDEXES						
7								(LSB)
8	Reserved							
9	NUMBER OF PHYS							
10	Reserved						CONFIGURING	CONFIGURABLE ROUTE TABLE
11	Reserved							
12		ENCLOSURE LOGICAL IDENTIFIER						
19								
20		Reserved						
27								
28	(MSB)	Reserved						
29								(LSB)
30	(MSB)	STP BUS INACTIVITY TIME LIMIT						
31								(LSB)
32	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						
33								(LSB)
34	(MSB)	STP SMP I_T NEXUS LOSS TIME						
35								(LSB)
~~28~~ 34	(MSB)	CRC						
~~31~~ 37								(LSB)

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

The FUNCTION field shall be set to 00h.

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 (i.e., 9). A RESPONSE LENGTH field set to 00h indicates there are 6 additional dwords (i.e., 24 additional bytes) before the CRC field in the response frame.

The EXPANDER CHANGE COUNT field counts the number of 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 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. This field shall not be incremented when forwarding a BROADCAST (CHANGE) from another expander device. This field shall wrap to zero after the maximum value (i.e., FFFFh) has been reached.

NOTE 3 - 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.

The EXPANDER ROUTE INDEXES field contains the maximum number of route indexes per phy for the expander device (see 4.6.7.3). 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 ROUTE INDEXES field to zero. Not all phys in an edge expander device are required to support the maximum number indicated by this field.

The NUMBER OF PHYS field contains the number of phys in the device, including any virtual phys and any vacant phys.

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) 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.

The CONFIGURABLE ROUTE TABLE bit indicates whether the expander device has an expander route table that is required to be configured with the SMP CONFIGURE ROUTE INFORMATION function (see 4.6.7.3). An expander device with a configurable route table shall set the CONFIGURABLE ROUTE TABLE bit to one. An expander device without a configurable route table or a device with any other device type shall set the CONFIGURABLE ROUTE TABLE bit to zero.

The ENCLOSURE LOGICAL IDENTIFIER field identifies the enclosure, if any, in which the device is located, and is defined in SES-2. The ENCLOSURE LOGICAL IDENTIFIER field shall be set to the same value reported by the enclosure services process, if any, for the enclosure. An ENCLOSURE LOGICAL IDENTIFIER field set to zero indicates no enclosure information is available.

The STP BUS INACTIVITY TIME LIMIT field contains the bus inactivity time limit for STP connections which is set by the CONFIGURE GENERAL function (see 10.4.3.xx).

The STP MAXIMUM CONNECT TIME LIMIT field contains the maximum connect time limit for STP connections which is set by the CONFIGURE GENERAL function (see 10.4.3.xx).

The STP SMP I_T NEXUS LOSS TIME field contains the time that an STP target port and an SMP initiator report retry certain connection requests which is set by the CONFIGURE GENERAL function (see 10.4.3.xx).

The CRC field is defined in 10.4.3.2.

10.4.3.3 REPORT MANUFACTURER INFORMATION function

The REPORT MANUFACTURER INFORMATION function returns vendor and product identification. This SMP function may be implemented by any SMP target port.

Table 8 defines the request format.

Table 8 — REPORT MANUFACTURER INFORMATION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (01h)							
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 01h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 0).

The CRC field is defined in 10.4.3.1.

Table 9 defines the response format.

Table 9 — REPORT MANUFACTURER INFORMATION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (01h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (0Eh)							
4	Reserved							
7								
8	Reserved							SAS-1.1 FORMAT
9	Reserved							
11								
12	(MSB) VENDOR IDENTIFICATION							
19								(LSB)
20	(MSB) PRODUCT IDENTIFICATION							
35								(LSB)
36	(MSB) PRODUCT REVISION LEVEL							
39								(LSB)
40	(MSB) COMPONENT VENDOR IDENTIFICATION							
47								(LSB)
48	(MSB) COMPONENT ID							
49								(LSB)
50	COMPONENT REVISION ID							
51	Reserved							
52	Vendor specific							
59								
60	(MSB) CRC							
63								(LSB)

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

The FUNCTION field shall be set to 01h.

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 (i.e., 0Eh). A RESPONSE LENGTH field set to 00h indicates there are 14 additional dwords (i.e., 56 additional bytes) before the CRC field in the response frame.

A SAS-1.1 FORMAT bit set to one indicates that bytes 40 through 59 are as defined in this standard. A SAS-1.1 FORMAT bit set to zero indicates that bytes 40 through 59 are vendor-specific as defined in the original version of this standard.

ASCII data fields (e.g., the VENDOR IDENTIFICATION field, the PRODUCT IDENTIFICATION field, and PRODUCT REVISION LEVEL field, and the COMPONENT VENDOR IDENTIFICATION field) shall contain only graphic codes (i.e., code values 20h through 7Eh). Left-aligned fields shall place any unused bytes at the end of the field (i.e., at the highest offset) and the unused bytes shall be filled with space characters (i.e., 20h).

The VENDOR IDENTIFICATION field contains eight bytes of ASCII data identifying the vendor of the subsystem (e.g., the board or enclosure) containing the component. The data shall be left-aligned within the field. The vendor identification string shall be one assigned by INCITS for use in the standard INQUIRY data VENDOR IDENTIFICATION field. A list of assigned vendor identification strings is in SPC-3 and on the T10 web site (http://www.t10.org).

The PRODUCT IDENTIFICATION field contains sixteen bytes of ASCII data identifying the type of the subsystem (e.g., the board or enclosure model number) containing the component, as defined by the vendor of the subsystem. The data shall be left-aligned within the field. The PRODUCT IDENTIFICATION field should be changed whenever the subsystem design changes in a way noticeable to a user (e.g., a different stock-keeping unit (SKU)).

The PRODUCT REVISION LEVEL field contains four bytes of ASCII data identifying the revision level of the subsystem (e.g., the board or enclosure) containing the component, as defined by the vendor of the subsystem. The data shall be left-aligned within the field. The PRODUCT REVISION LEVEL field should be changed whenever the subsystem design changes (e.g., any component change, even including resistor values).

All components on a subsystem should have the same values for their VENDOR IDENTIFICATION fields, PRODUCT IDENTIFICATION fields, and PRODUCT REVISION LEVEL fields.

NOTE 4 - Application clients may use the VENDOR IDENTIFICATION field and PRODUCT IDENTIFICATION field to identify the subsystem (e.g., for a user interface). Application clients may use the VENDOR IDENTIFICATION field, PRODUCT IDENTIFICATION field, PRODUCT REVISION LEVEL field to perform workarounds for problems in a specific revision of a subsystem.

The COMPONENT VENDOR IDENTIFICATION field contains eight bytes of ASCII data identifying the vendor of the component (e.g., the expander device) containing the SMP target port. The data shall be left-aligned within the field. The component vendor identification string shall be one assigned by INCITS for use in the standard INQUIRY data VENDOR IDENTIFICATION field. A list of assigned vendor identification strings is in SPC-3 and on the T10 web site (http://www.t10.org).

The COMPONENT ID field contains a 16-bit identifier identifying the type of the component (e.g., the expander device model number) containing the SMP target port, as defined by the vendor of the component. The COMPONENT ID field should be changed whenever the component's programming interface (e.g., the SMP target port definition) changes.

The COMPONENT REVISION LEVEL field contains an 8-bit identifier identifying the revision level of the component (e.g., the expander device) containing the SMP target port, as defined by the vendor of the component. The COMPONENT REVISION LEVEL field should be changed whenever the component changes but its programming interface does not change.

NOTE 5 - Application clients may use the COMPONENT VENDOR IDENTIFICATION field and the COMPONENT ID field to interpret vendor-specific information (e.g., vendor-specific SMP functions) correctly for that component. Application clients may use the COMPONENT VENDOR IDENTIFICATION field, the COMPONENT ID field, and the COMPONENT REVISION LEVEL field to perform workarounds for problems in a specific revision of a component.

The vendor-specific bytes are defined by the vendor of the subsystem (e.g., the board or enclosure) containing the component.

The CRC field is defined in 10.4.3.2.

10.4.3.4 DISCOVER function

The DISCOVER function returns the physical link configuration information for the specified phy. This SMP function provides information from the IDENTIFY address frame received by the phy and additional phy-specific information. This SMP function shall be implemented by all SMP target ports.

Table 10 defines the request format.

Table 10 — DISCOVER request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (10h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)							
15	CRC							(LSB)

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

The FUNCTION field shall be set to 10h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 2). A REQUEST LENGTH field set to 00h indicates there are 2 additional dwords (i.e., 8 additional bytes) before the CRC field in the request frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for the link configuration information being requested.

The CRC field is defined in 10.4.3.1.

Table 11 defines the response format.

Table 11 — DISCOVER response (part 1 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (10h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (0Ch)							

Table 11 — DISCOVER response (part 2 of 2)

Byte\Bit	7	6	5	4	3	2	1	0
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	Reserved	ATTACHED DEVICE TYPE			Reserved			
13	Reserved				NEGOTIATED PHYSICAL LINK RATE			
14	Reserved				ATTACHED SSP INITIATOR	ATTACHED STP INITIATOR	ATTACHED SMP INITIATOR	ATTACHED SATA HOST
15	ATTACHED SATA PORT SELECTOR	Reserved			ATTACHED SSP TARGET	ATTACHED STP TARGET	ATTACHED SMP TARGET	ATTACHED SATA DEVICE
16	SAS ADDRESS							
23								
24	ATTACHED SAS ADDRESS							
31								
32	ATTACHED PHY IDENTIFIER							
33	Reserved							
39								
40	PROGRAMMED MINIMUM PHYSICAL LINK RATE				HARDWARE MINIMUM PHYSICAL LINK RATE			
41	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				HARDWARE MAXIMUM PHYSICAL LINK RATE			
42	PHY CHANGE COUNT							
43	VIRTUAL PHY	Reserved			PARTIAL PATHWAY TIMEOUT VALUE			
44	Reserved				ROUTING ATTRIBUTE			
45	Reserved	CONNECTOR TYPE						
46	CONNECTOR ELEMENT INDEX							
47	CONNECTOR PHYSICAL LINK							
48	Reserved							
49								
50	Vendor specific							
51								
52	(MSB)	CRC						
55								(LSB)

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

The FUNCTION field shall be set to 10h.

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 (i.e., 12). A RESPONSE LENGTH field set to 00h indicates there are 12 additional dwords (i.e., 48 additional bytes) before the CRC field in the response frame.

The PHY IDENTIFIER field indicates the phy for which physical configuration link information is being returned.

The ATTACHED DEVICE TYPE field indicates the DEVICE TYPE value received during the link reset sequence and is defined in table 12.

Table 12 — ATTACHED DEVICE TYPE **field**

Code	Description
000b	No device attached
001b	End device
010b	Edge expander device
011b	Fanout expander device
All others	Reserved

The ATTACHED DEVICE TYPE field shall only be set to a value other than 000b after:

a) the identification sequence is complete if a SAS device or expander device is attached; or
b) the initial Register - Device to Host FIS has been received if a SATA phy is attached.

The NEGOTIATED PHYSICAL LINK RATE field is defined in table 13 and indicates the physical link rate negotiated during the link reset sequence. The negotiated physical link rate may be less than the programmed minimum physical link rate or greater than the programmed maximum physical link rate if the programmed physical link rates have been changed since the last link reset sequence.

Table 13 — NEGOTIATED PHYSICAL LINK RATE **field**

Code	Name	Description
0h	UNKNOWN	Phy is enabled; unknown physical link rate. [a]
1h	DISABLED	Phy is disabled.
2h	PHY_ RESET_ PROBLEM	Phy is enabled; the phy obtained dword synchronization for at least one physical link rate during the SAS speed negotiation sequence (see 6.7.4.2), but the SAS speed negotiation sequence failed (i.e., the last speed negotiation window, using a physical link rate expected to succeed, failed). These failures may be logged in the SMP REPORT PHY ERROR LOG function (see 10.4.3.5) and/or the Protocol-Specific Port log page (see 10.2.8.1).
3h	SPINUP_ HOLD	Phy is enabled; detected a SATA device and entered the SATA spinup hold state. The LINK RESET and HARD RESET operations in the SMP PHY CONTROL function (see 10.4.3.10) may be used to release the phy. This field shall be updated to this value at SATA spinup hold time (see 6.8.7 and 6.10)(i.e., after the COMSAS Detect Timeout timer expires during the SATA OOB sequence) if SATA spinup hold is supported.
4h	PORT_ SELECTOR	Phy is enabled; detected a SATA port selector. The physical link rate has not been negotiated since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The SATA spinup hold state has not been entered since the last time the phy's SP state machine entered the SP0:OOB_COMINIT state. The value in this field may change to 3h, 8h, or 9h if attached to the active phy of the SATA port selector. Presence of a SATA port selector is indicated by the ATTACHED SATA PORT SELECTOR bit.
8h	G1	Phy is enabled; 1,5 Gbps physical link rate. This field shall be updated to this value after the speed negotiation sequence completes.
9h	G2	Phy is enabled; 3,0 Gbps physical link rate. This field shall be updated to this value after the speed negotiation sequence completes.
All others	Reserved.	

[a] This code may be used by an application client in its local data structures to indicate an unknown negotiated physical link rate (e.g., before the discover process has queried the phy).

Table 14 describes the ATTACHED SATA PORT SELECTOR bit and the ATTACHED SATA DEVICE bit.

Table 14 — 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).

[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).
[b] Whenever the ATTACHED SATA PORT SELECTOR bit changes, the phy shall generate a BROADCAST(CHANGE) notification.
[c] For the purposes of the ATTACHED SATA DEVICE bit, the SATA port selector is not considered a SATA device.
[d] The ATTACHED SATA DEVICE bit shall be updated at SATA spin-up hold time (see 6.8.7 and 6.10).

An ATTACHED SATA HOST bit set to one indicates a SATA host port is attached. An ATTACHED SATA HOST bit set to zero indicates a SATA host port is not attached.

NOTE 6 - Support for SATA hosts is outside the scope of this standard.

If a SAS phy reset sequence occurs (see 6.7.4)(i.e., one or more of the ATTACHED SSP INITIATOR PORT bit, ATTACHED STP INITIATOR PORT bit, the ATTACHED SMP INITIATOR PORT bit, the ATTACHED SSP TARGET PORT bit, the ATTACHED STP TARGET PORT bit, and/or the ATTACHED SMP TARGET PORT bit is set to one), then the ATTACHED SATA PORT SELECTOR bit, the ATTACHED SATA DEVICE bit, and the ATTACHED SATA HOST bit shall each be set to zero.

The ATTACHED SSP INITIATOR PORT bit indicates the value of the SSP INITIATOR PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED STP INITIATOR PORT bit indicates the value of the STP INITIATOR PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED SMP INITIATOR PORT bit indicates the value of the SMP INITIATOR PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED SSP TARGET PORT bit indicates the value of the SSP TARGET PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED STP TARGET PORT bit indicates the value of the STP TARGET PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED SMP TARGET PORT bit indicates the value of the SMP TARGET PORT field received in the IDENTIFY address frame (see 7.8.2) during the identification sequence.

The ATTACHED SSP INITIATOR PORT bit, ATTACHED STP INITIATOR PORT bit, ATTACHED SMP INITIATOR PORT bit, ATTACHED SSP TARGET PORT bit, ATTACHED STP TARGET PORT bit, and ATTACHED SMP TARGET PORT bit shall be updated at the end of the identification sequence.

If a SATA phy reset sequence occurs (see 6.7.3)(i.e., the ATTACHED SATA PORT SELECTOR bit is set to one, the ATTACHED SATA DEVICE bit is set to one, or the ATTACHED SATA HOST bit is set to one), then the ATTACHED SSP INITIATOR PORT bit, ATTACHED STP INITIATOR PORT bit, ATTACHED SMP INITIATOR PORT bit, ATTACHED SSP TARGET PORT bit, ATTACHED STP TARGET PORT bit, and ATTACHED SMP TARGET PORT bit shall each be set to zero.

The SAS ADDRESS field contains the value of the SAS ADDRESS field transmitted in the IDENTIFY address frame during the identification sequence. If the phy is an expander phy, the SAS ADDRESS field contains the SAS address of the expander device (see 4.2.4). If the phy is a SAS phy, the SAS ADDRESS field contains the SAS address of the SAS port (see 4.2.6).

The ATTACHED SAS ADDRESS field contains the value of the SAS ADDRESS field received in the IDENTIFY address frame during the identification sequence. If the attached port is an expander port, the ATTACHED SAS ADDRESS field contains the SAS address of the attached expander device (see 4.2.4). If the attached port is a SAS port, the ATTACHED SAS ADDRESS field contains SAS address of the attached SAS port (see 4.2.6). If the attached port is a SATA device port, the ATTACHED SAS ADDRESS field contains the SAS address of the STP/SATA bridge (see 4.6.2).

The ATTACHED SAS ADDRESS field shall be updated:

a) after the identification sequence completes, if a SAS phy or expander phy is attached; or
b) after the COMSAS Detect Timeout timer expires (see 6.8.3.9), if a SATA phy is attached.

An STP initiator port should not make a connection request to the attached SAS address until the ATTACHED DEVICE TYPE field is set to a value other than 000b.

The ATTACHED PHY IDENTIFIER field contains a phy identifier for the attached phy:

a) If the attached phy is a SAS phy or an expander phy, the ATTACHED PHY IDENTIFIER field contains the value of the PHY IDENTIFIER field received in the IDENTIFY address frame during the identification sequence:
   A) If the attached phy is a SAS phy, the ATTACHED PHY IDENTIFIER field contains the phy identifier of the attached SAS phy in the attached SAS device;
   B) If the attached phy is an expander phy, the ATTACHED PHY IDENTIFIER field contains the phy identifier (see 4.2.7) of the attached expander phy in the attached expander device; and
b) If the attached phy is a SATA device phy, the ATTACHED PHY IDENTIFIER field contains 00h;
c) If the attached phy is a SATA port selector phy and the expander device is able to determine the port of the SATA port selector to which it is attached, the ATTACHED PHY IDENTIFIER field contains 00h or 01h; and
d) If the attached phy is a SATA port selector phy and the expander device is not able to determine the port of the SATA port selector to which it is attached, the ATTACHED PHY IDENTIFIER field contains 00h.

The ATTACHED PHY IDENTIFIER field shall be updated:

a) after the identification sequence completes, if a SAS phy or expander phy is attached; or
b) after the COMSAS Detect Timeout timer expires (see 6.8.3.9), if a SATA phy is attached.

The PROGRAMMED MINIMUM PHYSICAL LINK RATE field indicates the minimum physical link rate set by the PHY CONTROL function (see 10.4.3.10). The values are defined in table 15. The default value shall be the value of the HARDWARE MINIMUM PHYSICAL LINK RATE field.

The HARDWARE MINIMUM PHYSICAL LINK RATE field indicates the minimum physical link rate supported by the phy. The values are defined in table 16.

The PROGRAMMED MAXIMUM PHYSICAL LINK RATE field indicates the maximum physical link rate set by the PHY CONTROL function (see 10.4.3.10). The values are defined in table 15. The default value shall be the value of the HARDWARE MAXIMUM PHYSICAL LINK RATE field.

Table 15 — PROGRAMMED MINIMUM PHYSICAL LINK RATE **and** PROGRAMMED MAXIMUM PHYSICAL LINK **rate fields**

Code	Description
0h	Not programmable
8h	1,5 Gbps
9h	3,0 Gbps
All others	Reserved

The HARDWARE MAXIMUM PHYSICAL LINK RATE field indicates the maximum physical link rate supported by the phy. The values are defined in table 16.

Table 16 — HARDWARE MINIMUM PHYSICAL LINK RATE **and** HARDWARE MAXIMUM PHYSICAL LINK RATE **fields**

Code	Description
8h	1,5 Gbps
9h	3,0 Gbps
All others	Reserved

The PHY CHANGE COUNT field counts the number of 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) for any reason described in 7.11 originating from the expander phy other than forwarding a BROADCAST (CHANGE).

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 7 - 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.

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 PARTIAL PATHWAY TIMEOUT VALUE field indicates the partial pathway timeout value in microseconds (see 7.12.4.5) set by the PHY CONTROL function (see 10.4.3.10).

NOTE 8 - The recommended default value for PARTIAL PATHWAY TIMEOUT VALUE is 7 µs. ~~The partial pathway timeout value may be set by the PHY CONTROL function (see 10.4.3.10).~~

Editor's Note 9: above change makes the wording more consistent with other fields

---

The ROUTING ATTRIBUTE field indicates the routing attribute supported by the phy (see 4.6.7.1) and is defined in table 17.

Table 17 — ROUTING ATTRIBUTE **field**

Code	Name	Description
0h	Direct routing attribute	Direct routing method for attached end devices. Attached expander devices are not supported on this phy.
1h	Subtractive routing attribute	Either: a) subtractive routing method for attached expander devices; or b) direct routing method for attached end devices.
2h	Table routing attribute	Either: a) table routing method for attached expander devices; or b) direct routing method for attached end devices.
All others	Reserved	

The ROUTING ATTRIBUTE field shall not change based on the attached device type.

The CONNECTOR TYPE field indicates the type of connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2). A CONNECTOR TYPE field set to 00h indicates no connector information is available and that the CONNECTOR ELEMENT INDEX field and the CONNECTOR PHYSICAL LINK fields are invalid and shall be ignored.

The CONNECTOR ELEMENT INDEX indicates the element index of the SAS Connector element representing the connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2).

The CONNECTOR PHYSICAL LINK field indicates the physical link in the connector used to access the phy, as reported by the enclosure services process for the enclosure (see the SAS Connector element in SES-2).

The CRC field is defined in 10.4.3.2.

10.4.3.5 REPORT PHY ERROR LOG function

The REPORT PHY ERROR LOG function returns error logging information about the specified phy. This SMP function may be implemented by any SMP target port.

Table 18 defines the request format.

Table 18 — REPORT PHY ERROR LOG request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (11h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)			CRC				
15								(LSB)

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

The FUNCTION field shall be set to 11h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 2). A REQUEST LENGTH field set to 00h indicates there are 2 additional dwords (i.e., 8 additional bytes) before the CRC field in the request frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which information shall be reported.

The CRC field is defined in 10.4.3.1.

Table 19 defines the response format.

Table 19 — REPORT PHY ERROR LOG response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (11h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (06h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)	INVALID DWORD COUNT						
15								(LSB)
16	(MSB)	RUNNING DISPARITY ERROR COUNT						
19								(LSB)
20	(MSB)	LOSS OF DWORD SYNCHRONIZATION COUNT						
23								(LSB)
24	(MSB)	PHY RESET PROBLEM COUNT						
27								(LSB)
28	(MSB)	CRC						
31								(LSB)

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

The FUNCTION field shall be set to 11h.

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 (i.e., 6). A RESPONSE LENGTH field set to 00h indicates there are 6 additional dwords (i.e., 24 additional bytes) before the CRC field in the response frame.

The PHY IDENTIFIER field indicates the phy (see 4.2.7) for which information is being reported.

The INVALID DWORD COUNT field indicates the number of invalid dwords (see 3.1.100) that have been received outside of phy reset sequences (i.e., between when the SP state machine (see 6.8) sends a Phy Layer Ready (SAS) confirmation or Phy Layer Ready (SATA) confirmation and when it sends a Phy Layer Not Ready confirmation to the link layer). The count shall stop at the maximum value. The INVALID DWORD COUNT field is set to a vendor-specific value after power on.

The RUNNING DISPARITY ERROR COUNT field indicates the number of dwords containing running disparity errors (see 6.2) that have been received outside of phy reset sequences. The count shall stop at the maximum value. The RUNNING DISPARITY ERROR COUNT field is set to a vendor-specific value after power on.

The LOSS OF DWORD SYNCHRONIZATION COUNT field indicates the number of times the phy has restarted the link reset sequence because it lost dword synchronization (see 6.9) (i.e., the SP state machine transitioned from SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready to SP0:OOB_COMINIT (see 6.8)). The count shall stop at the maximum value. The LOSS OF DWORD SYNCHRONIZATION COUNT field is set to a vendor-specific value after power on.

The PHY RESET PROBLEM COUNT field indicates the number of times the phy did not obtain dword synchronization during the final SAS speed negotiation window (see 6.7.4.2). The count shall stop at the maximum value. The PHY RESET PROBLEM COUNT field is set to a vendor-specific value after power on.

The CRC field is defined in 10.4.3.2.

10.4.3.6 REPORT PHY SATA function

The REPORT PHY SATA function returns information about the SATA state for a specified phy. This SMP function shall be implemented by SMP target ports that share SAS addresses with STP target ports and by SMP target ports in expander devices with STP/SATA bridges. This SMP function shall not be implemented by any other type of SMP target port.

Table 203 defines the request format.

Table 20 — REPORT PHY SATA request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (12h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)							
15	CRC							(LSB)

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

The FUNCTION field shall be set to 12h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 2). A REQUEST LENGTH field set to 00h indicates there are 2 additional dwords (i.e., 8 additional bytes) before the CRC field in the request frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) for which information shall be reported.

The CRC field is defined in 10.4.3.1.

Table 204 defines the response format.

Table 21 — REPORT PHY SATA response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (12h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (0Eh)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	Reserved							
11	Reserved					STP I_T NEXUS LOSS OCCURRED	AFFILIATIONS SUPPORTED	AFFILIATION VALID
12	Reserved							
15								
16	STP SAS ADDRESS							
23								
24	REGISTER DEVICE TO HOST FIS							
43								
44	Reserved							
47								
48	AFFILIATED STP INITIATOR SAS ADDRESS							
55								
56	STP I_T NEXUS LOSS SAS ADDRESS							
63								
~~56~~ 64	(MSB) CRC							
~~59~~ 67								(LSB)

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

The FUNCTION field shall be set to 12h.

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 (i.e., 15). A RESPONSE LENGTH field set to 00h indicates there are 13 additional dwords (i.e., 52 additional bytes) before the CRC field in the response frame.

The PHY IDENTIFIER field indicates the phy (see 4.2.7) for which information is being reported.

An STP I_T NEXUS LOSS OCCURRED bit set to one indicates that the STP target port encountered an I_T nexus loss for the STP initiator port whose SAS address is indicated in the STP I_T NEXUS LOSS SAS ADDRESS field. An STP I_T NEXUS LOSS OCCURRED bit set to zero indicates that either an I_T nexus loss has not occurred or the STP target port has successfully established a connection with the indicated STP initiator port.

An AFFILIATIONS SUPPORTED bit set to one indicates that affiliations are supported by the STP target port containing the specified phy. An AFFILIATIONS SUPPORTED bit set to zero indicates that affiliations are not supported by the STP target port containing the specified phy.

~~The~~An AFFILIATION VALID bit ~~shall be~~ set to one ~~when~~indicates that the AFFILIATED STP INITIATOR SAS ADDRESS field is valid and the STP target port containing the specified phy is maintaining an affiliation (see 7.17.5).

~~The~~An AFFILIATION VALID bit ~~shall be~~ set to zero ~~when~~indicates that no affiliation is being maintained.

The STP SAS ADDRESS field contains the SAS address (see 4.2.2) of the STP target port that contains the specified phy.

The REGISTER DEVICE TO HOST FIS field contains the contents of the initial Register - Device to Host FIS. For an STP/SATA bridge, this is delivered by the attached SATA device after a link reset sequence (see ATA/ATAPI-7 V3 and SATAII-EXT). For a native STP target port in an end device, this is directly provided.

The FIS contents shall be stored with little-endian byte ordering (i.e., the first byte, byte 24, contains the FIS Type).

For an STP/SATA bridge, the first byte of the field (i.e., the FIS Type) shall be initialized to zero on power on and whenever the phy has restarted the link reset sequence after losing dword synchronization (see 6.9)(i.e., the SP state machine transitioned from SP15:SAS_PHY_Ready or SP22:SATA_PHY_Ready to SP0:OOB_COMINIT (see 6.8)) to indicate the field is invalid and the attached SATA device has not delivered a Register – Device to Host FIS. The first byte of the field shall be set to 34h when the attached SATA device has delivered the initial Register – Device to Host FIS. The remaining contents of the REGISTER DEVICE TO HOST FIS field shall remain constant until a link reset sequence causes the attached SATA device to deliver another initial Register – Device to Host FIS.

An STP/SATA bridge that receives a connection request for a SATA device that has not successfully delivered the initial Register – Device to Host FIS shall return an OPEN_REJECT (NO DESTINATION).

NOTE 9 - If there is a problem receiving the expected initial Register - Device to Host FIS, the STP/SATA bridge should use SATA_R_ERR to retry until it succeeds. In the DISCOVER response, the ATTACHED SATA DEVICE bit is set to one and the ATTACHED SAS ADDRESS field is valid, but the ATTACHED DEVICE TYPE field is set to 000b (i.e., no device attached) during this time.

The AFFILIATED STP INITIATOR SAS ADDRESS field contains the SAS address (see 4.2.2) of the last STP initiator port that ~~currently has~~had an affiliation with the STP target port that contains the specified phy.

The STP I_T NEXUS LOSS SAS ADDRESS field contains the SAS address (see 4.2.2) of the last STP initiator port for which the STP target port experienced an I_T nexus loss (see 4.5).

Editor's Note 10: The SAS ADDRESS fields maintain their values even after the corresponding VALID bits are set to zero; they are not zeroed out immediately. This seems to match SAS-1.1 implementations for the AFFILIATED STP INITIATOR SAS ADDRESS field.

The CRC field is defined in 10.4.3.2.

10.4.3.7 REPORT ROUTE INFORMATION function

The REPORT ROUTE INFORMATION function returns an expander route entry from the expander route table within an expander device. This SMP function shall be supported by SMP target ports in expander devices if the EXPANDER ROUTE INDEXES field is non-zero in the REPORT GENERAL function. This SMP function may be used as a diagnostic tool to resolve topology issues.

Table 22 defines the request format.

Table 22 — REPORT ROUTE INFORMATION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (13h)							
2	Reserved							
3	REQUEST LENGTH (02h)							
4	Reserved							
5								
6	(MSB)	EXPANDER ROUTE INDEX						
7								(LSB)
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11								
12	(MSB)	CRC						
15								(LSB)

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

The FUNCTION field shall be set to 13h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 2). A REQUEST LENGTH field set to 00h indicates there are 3 additional dwords (i.e., 8 additional bytes) before the CRC field in the request frame.

The EXPANDER ROUTE INDEX field specifies the expander route index for the expander route entry being requested (see 4.6.7.3).

The PHY IDENTIFIER field specifies the phy for which the expander route entry is being requested.

The CRC field is defined in 10.4.3.1.

Table 23 defines the response format.

Table 23 — REPORT ROUTE INFORMATION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (13h)							
2	FUNCTION RESULT							
3	RESPONSE LENGTH (09h)							
4	Reserved							
5								
6	(MSB)	EXPANDER ROUTE INDEX						
7								(LSB)
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11								
12	EXPANDER ROUTE ENTRY DISABLED	Reserved						
13	Reserved							
15								
16	ROUTED SAS ADDRESS							
23								
24	Reserved							
39								
40	(MSB)	CRC						
43								(LSB)

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

The FUNCTION field shall be set to 13h.

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 (i.e., 9). A RESPONSE LENGTH field set to 00h indicates there are 9 additional dwords (i.e., 36 additional bytes) before the CRC field in the response frame.

The EXPANDER ROUTE INDEX field contains the expander route index for the expander route entry being returned (see 4.6.7.3).

The PHY IDENTIFIER field contains the phy identifier for the expander route entry being returned.

The EXPANDER ROUTE ENTRY DISABLED bit indicates whether the ECM shall use the expander route entry to route connection requests (see 4.6.7.3). If the EXPANDER ROUTE ENTRY DISABLED bit is set to zero, then the ECM shall use the expander route entry to route connection requests. If the EXPANDER ROUTE ENTRY DISABLED bit is set to one, the ECM shall not use the expander route entry to route connection requests.

The ROUTED SAS ADDRESS field contains the SAS address in the expander route entry (see 4.6.7.3).

The CRC field is defined in 10.4.3.2.

10.4.3.8 CONFIGURE ROUTE INFORMATION function

The CONFIGURE ROUTE INFORMATION function sets an expander route entry within the expander route table of a configurable expander device. This SMP function shall be supported by SMP target ports in expander devices if the CONFIGURABLE ROUTE TABLE field is set to one in the REPORT GENERAL response data. Other SMP target ports shall not support this SMP function.

Table 24 defines the request format.

Table 24 — CONFIGURE ROUTE INFORMATION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (90h)							
2	Reserved							
3	REQUEST LENGTH (09h)							
4	Reserved							
5								
6	(MSB)							
7	EXPANDER ROUTE INDEX							(LSB)
8	Reserved							
9	PHY IDENTIFIER							
10	Reserved							
11								
12	DISABLE EXPANDER ROUTE ENTRY	Reserved						
13	Reserved							
15								
16	ROUTED SAS ADDRESS							
23								
24	Reserved							
39								
40	(MSB)							
43	CRC							(LSB)

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

The FUNCTION field shall be set to 90h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 9). A REQUEST LENGTH field set to 00h indicates there are 9 additional dwords (i.e., 36 additional bytes) before the CRC field in the request frame.

The EXPANDER ROUTE INDEX field specifies the expander route index for the expander route entry being configured (see 4.6.7.3).

The PHY IDENTIFIER field specifies the phy for which the expander route entry is being configured (see 4.6.7.3).

The DISABLE EXPANDER ROUTE ENTRY bit specifies whether the ECM shall use the expander route entry to route connection requests (see 4.6.7.3). If the DISABLE EXPANDER ROUTE ENTRY bit is set to zero, then the ECM shall use the expander route entry to route connection requests. If the DISABLE EXPANDER ROUTE ENTRY bit is set to one, the ECM shall not use the expander route entry to route connection requests.

The ROUTED SAS ADDRESS field contains the SAS address for the expander route entry being configured (see 4.6.7.3).

The CRC field is defined in 10.4.3.1.

Table 25 defines the response format.

Table 25 — CONFIGURE ROUTE INFORMATION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (90h)							
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 90h.

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 (i.e., 0).

The CRC field is defined in 10.4.3.2.

10.4.3.9 PHY CONTROL function

The PHY CONTROL function requests actions by the specified phy. This SMP function may be implemented by any SMP target port.

Table 209 defines the request format.

Table 26 — PHY CONTROL request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (91h)							
2	Reserved							
3	REQUEST LENGTH (09h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	PHY OPERATION							
11	Reserved							UPDATE PARTIAL PATHWAY TIMEOUT VALUE
12	Reserved							
31								
32	PROGRAMMED MINIMUM PHYSICAL LINK RATE				Reserved			
33	PROGRAMMED MAXIMUM PHYSICAL LINK RATE				Reserved			
34	Reserved							
35								
36	Reserved				PARTIAL PATHWAY TIMEOUT VALUE			
37	Reserved							
39								
40	(MSB) CRC							
43								(LSB)

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

The FUNCTION field shall be set to 91h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 9). A REQUEST LENGTH field set to 00h indicates there are 9 additional dwords (i.e., 36 additional bytes) before the CRC field in the request frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) to which the PHY CONTROL request applies.

Table 27 defines the PHY OPERATION field.

Table 27 — PHY OPERATION field (part 1 of 2)

Code	Operation	Description
00h	NOP	No operation.
01h	LINK RESET	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) requirements related to this phy operation in an expander device. 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). The SMP response shall be returned without waiting for the link reset to complete.
02h	HARD RESET	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) requirements related to this phy operation in an expander device. If the specified phy is a virtual phy, perform an internal reset and enable the specified phy. Any affiliation (see 7.17.5) shall be cleared. The SMP response shall be returned without waiting for the hard reset to complete.
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) requirements related to this phy operation in an expander device.
04h	Reserved	
05h	CLEAR ERROR LOG	Clear the error log counters (see 10.4.3.5) for the specified phy.

Table 27 — PHY OPERATION **field** (part 2 of 2)

Code	Operation	Description
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.
07h	TRANSMIT SATA PORT SELECTION SIGNAL	This function shall only be supported by phys in an expander device. 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 (CHANGE) requirements related to this phy operation in an expander device. Any affiliation (see 7.17.5) shall be cleared. 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. 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.
All others	Reserved	

If the PHY IDENTIFIER field specifies the phy which is being used for the SMP connection and a phy operation of LINK RESET, HARD RESET, or DISABLE is requested, the SMP target port shall not perform the requested operation and shall return a function result of SMP FUNCTION FAILED in the response frame.

An UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit set to one specifies that the PARTIAL PATHWAY TIMEOUT VALUE field shall be honored. An UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit set to zero specifies that the PARTIAL PATHWAY TIMEOUT VALUE field shall be ignored.

The PROGRAMMED MINIMUM PHYSICAL LINK RATE field specifies the minimum physical link rate the phy shall support during a link reset sequence (see 4.4.1). Table 28 defines the values for this field. If this field is changed along with a phy operation of LINK RESET or HARD RESET, that phy operation shall utilize the new value for this field. This value is reported in the DISCOVER function (see 10.4.3.5).

The PROGRAMMED MAXIMUM PHYSICAL LINK RATE field specifies the maximum physical link rates the phy shall support during a link reset sequence (see 4.4.1).Table 28 defines the values for this field. If this field is changed along with a phy operation of LINK RESET or HARD RESET, that phy operation shall utilize the new value for this field. This value is reported in the DISCOVER function (see 10.4.3.5).

Table 28 — PROGRAMMED MINIMUM PHYSICAL LINK RATE **and** PROGRAMMED MAXIMUM PHYSICAL LINK RATE **fields**

Code	Description
0h	Do not change current value
8h	1,5 Gbps
9h	3,0 Gbps
All others	Reserved

If the PROGRAMMED MINIMUM PHYSICAL LINK RATE field or the PROGRAMMED MAXIMUM PHYSICAL LINK RATE field is set to an unsupported or reserved value, or the PROGRAMMED MINIMUM PHYSICAL LINK RATE field and PROGRAMMED MAXIMUM PHYSICAL LINK RATE field are set to an invalid combination of values (e.g., the minimum is greater than the maximum), the SMP target port shall not change either of their values and may return a function result of SMP FUNCTION FAILED in the response frame. If it returns a function result of SMP FUNCTION FAILED, it shall not perform the requested phy operation.

The PARTIAL PATHWAY TIMEOUT VALUE field specifies the amount of time in microseconds the expander phy shall wait after receiving an Arbitrating (Blocked On Partial) confirmation from the ECM before requesting that the ECM resolve pathway blockage (see 7.12.4.6). A PARTIAL PATHWAY TIMEOUT VALUE field value of zero (i.e., 0 μs) specifies that partial pathway resolution shall be requested by the expander phy immediately upon reception of an Arbitrating (Blocked On Partial) confirmation from the ECM. The PARTIAL PATHWAY TIMEOUT VALUE field is only honored when the UPDATE PARTIAL PATHWAY TIMEOUT VALUE bit is set to one. This value is reported in the DISCOVER function (see 10.4.3.5).

The CRC field is defined in 10.4.3.1.

Table 29 defines the response format.

Table 29 — PHY CONTROL response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (91h)							
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 91h.

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 (i.e., 0).

The CRC field is defined in 10.4.3.2.

10.4.3.10 PHY TEST FUNCTION function

The PHY TEST FUNCTION function requests actions by the specified phy. This SMP function may be implemented by any SMP target port.

Table 30 defines the request format.

Table 30 — PHY TEST FUNCTION request

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (40h)							
1	FUNCTION (92h)							
2	Reserved							
3	REQUEST LENGTH (09h)							
4	Reserved							
8								
9	PHY IDENTIFIER							
10	PHY TEST FUNCTION							
11	PHY TEST PATTERN							
12	Reserved							
14								
15	Reserved				PHY TEST PATTERN PHYSICAL LINK RATE			
16	Reserved							
39								
40	(MSB) CRC							
43								(LSB)

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

The FUNCTION field shall be set to 92h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 10). A REQUEST LENGTH field set to 00h indicates there are 9 additional dwords (i.e., 36 additional bytes) ,before the CRC field in the request frame.

The PHY IDENTIFIER field specifies the phy (see 4.2.7) to which the PHY TEST PATTERN request applies.

If the PHY IDENTIFIER field specifies the phy which is being used for the SMP connection, the SMP target port shall not perform the requested operation and shall return a function result of SMP FUNCTION FAILED in the response frame.

The PHY TEST FUNCTION field specifies the phy test function to be performed, and is defined in table 31. If the PHY TEST FUNCTION field specifies a phy test function that is not supported by the phy, the SMP target port shall return a function result of UNKNOWN PHY TEST FUNCTION in the response frame.

Table 31 — PHY TEST FUNCTION **field**

Code	Name	Description
00h	STOP	If the selected phy is performing a phy test function, then the selected phy shall stop performing the phy test function and originate a link reset sequence. If the selected phy is not performing a phy test function, then this function has no effect on the selected phy.
01h	TRANSMIT_ PATTERN	If the selected phy is not performing a phy test function, the selected phy shall be set to transmit the phy test pattern specified by the PHY TEST PATTERN field at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field and set to ignore its receiver. If the selected phy receives data while transmitting the pattern, then the selected phy shall ignore the received data. If the selected phy is performing a phy test function, the SMP target port shall return a function result of PHY TEST FUNCTION IN PROGRESS in the response frame.
02h - EFh	Reserved	
F0h - FFh	Vendor specific	

If the PHY TEST FUNCTION field is set to 01h, the PHY TEST PATTERN field specifies the phy test pattern to be performed, and the same as that defined in table 180 for the Protocol-Specific diagnostic page (see 10.2.9.1). The phy test pattern shall be sent at the physical link rate specified by the PHY TEST PATTERN PHYSICAL LINK RATE field.

The PHY TEST PATTERN PHYSICAL LINK RATE field specifies the physical link rate at which the phy test function, if any, shall be performed. Table 32 defines the values for this field.

Table 32 — PHY TEST PATTERN PHYSICAL LINK RATE **field**

Code	Description
8h	1,5 Gbps
9h	3,0 Gbps
All others	Reserved

The CRC field is defined in 10.4.3.1.

Table 33 defines the response format.

Table 33 — PHY TEST FUNCTION response

Byte\Bit	7	6	5	4	3	2	1	0
0	SMP FRAME TYPE (41h)							
1	FUNCTION (92h)							
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 92h.

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 (i.e., 0).

The CRC field is defined in 10.4.3.2.

10.4.3.10 CONFIGURE GENERAL function

The CONFIGURE GENERAL function requests actions by the device containing the SMP target port. This SMP function may be implemented by any SMP target port.

Table 34 defines the request format.

Table 34 — 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 (03h)							
4	Reserved							
7								
8	Reserved					UPDATE STP SMP I_T NEXUS LOSS TIME	UPDATE STP MAXIMUM CONNECT TIME LIMIT	UPDATE STP BUS INACTIVITY TIME LIMIT
9	Reserved							
11	(MSB)	STP BUS INACTIVITY TIME LIMIT						
11								(LSB)
12	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						
13								(LSB)
14	(MSB)	STP SMP I_T NEXUS LOSS TIME						
15								(LSB)
16	(MSB)	CRC						
19								(LSB)

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

The FUNCTION field shall be set to 80h.

The REQUEST LENGTH field contains the number of dwords that follow, not including the CRC field (i.e., 3).

An UPDATE STP BUS INACTIVITY TIME LIMIT bit set to one specifies that the STP BUS INACTIVITY TIME LIMIT field shall be honored. An UPDATE STP BUS INACTIVITY TIME LIMIT bit set to zero specifies that the STP BUS INACTIVITY TIME LIMIT field shall be ignored.

An UPDATE STP MAXIMUM CONNECT TIME LIMIT bit set to one specifies that the STP MAXIMUM CONNECT TIME LIMIT field shall be honored. An UPDATE STP MAXIMUM CONNECT TIME LIMIT bit set to zero specifies that the STP MAXIMUM CONNECT TIME LIMIT field shall be ignored.

An UPDATE STP SMP I_T NEXUS LOSS TIME bit set to one specifies that the STP SMP I_T NEXUS LOSS TIME field shall be honored. An UPDATE STP SMP I_T NEXUS LOSS TIME bit set to zero specifies that the STP SMP I_T NEXUS LOSS TIME field shall be ignored.

The STP BUS INACTIVITY TIME LIMIT field contains the maximum period that an STP target port is permitted to maintain a connection (see 4.1.10) without transferring a frame to the STP initiator port. This value shall be the number of 100 μs increments between frames that the STP target port transmits during a connection.

When this number is exceeded, the STP target port shall close the connection. A value of zero in this field specifies that there is no bus inactivity time limit. This value is reported in the REPORT GENERAL function (see 10.4.3.x). The bus inactivity time limit is enforced by the port layer (see 8.2.3).

The STP MAXIMUM CONNECT TIME LIMIT field contains the maximum duration of a connection (see 4.1.10). This value shall be the number of 100 µs increments that an STP target port transmits during a connection after which the STP target port shall connection at the next opportunity (e.g., a value of one in this field means that the time is less than or equal to 100 µs and a value of two in this field means that the time is less than or equal to 200 µs). If an STP target port is transferring a frame when the maximum connection time limit is exceeded, the STP target port shall complete transfer of the frame before closing the connection. A value of zero in this field specifies that there is no maximum connection time limit. This value is reported in the REPORT GENERAL function (see 10.4.3.x). The maximum connection time limit is enforced by the port layer (see 8.2.3).

The STP SMP I_T NEXUS LOSS TIME field contains the time that an STP target port shall retry connection requests to an STP initiator port that are rejected with responses indicating the STP initiator port may no longer be present (see 8.2.2) before recognizing an I_T nexus loss (see 4.5). Table 35 defines the values of the STP SMP I_T NEXUS LOSS TIME field. This value is enforced by the port layer (see 8.2.2).

Table 35 — STP SMP I_T NEXUS LOSS TIME field

Code	Description
0000h	Vendor-specific amount of time.
0001h to FFFEh	Time in milliseconds.
FFFFh	The STP target port shall never recognize an I_T nexus loss (i.e., it shall retry the connection requests forever).

NOTE 10 - The default value of the STP SMP I_T NEXUS LOSS TIME field should be non-zero. It is recommended that this value be 2 000 ms.

NOTE 11 - An STP initiator port should retry connection requests for the time indicated by the STP SMP I_T NEXUS LOSS field in the SMP REPORT GENERAL function for the STP target port to which it is trying to establish a connection (see 4.5).