

## **T10/07-391 revision 1**

Date: October 02, 2007

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: SAS-2: Limiting SAS Target response to OPEN\_REJECT (RETRY)

### **1 Overview**

In a SAS topology with a large imbalance between the number of initiators and targets (e.g., one or two initiators with 50 or more targets) with a heavy workload it is not unusual for the minority SAS device to become bandwidth limited. When this happens the bandwidth limited SAS device will start issuing OPEN\_REJECT (RETRY) to most open attempts. In response to the OPEN\_REJECT (RETRY) many SAS devices will immediately issue another OPEN. This causes a flood of OPEN/OPEN\_REJECT (RETRY)s which accomplishes nothing except to add congestion to an already congested topology.

This proposal brings the DISCONNECT TIME LIMIT field of the Disconnect-Reconnect mode page into SAS-2.

Revision 1 - Moves the timer to the Shared Port Control mode page for SSP and the CONFIGURE GENERAL function for SMP and STP. Adds the timer description into the port layer specifically the PL\_OC state machine.

### **2 Proposed SAS-2 changes**

#### **8.1 PL (port layer) state machines**

##### **8.1.1 PL state machines overview**

The PL (port layer) consists of state machines that run in parallel and perform the following functions:

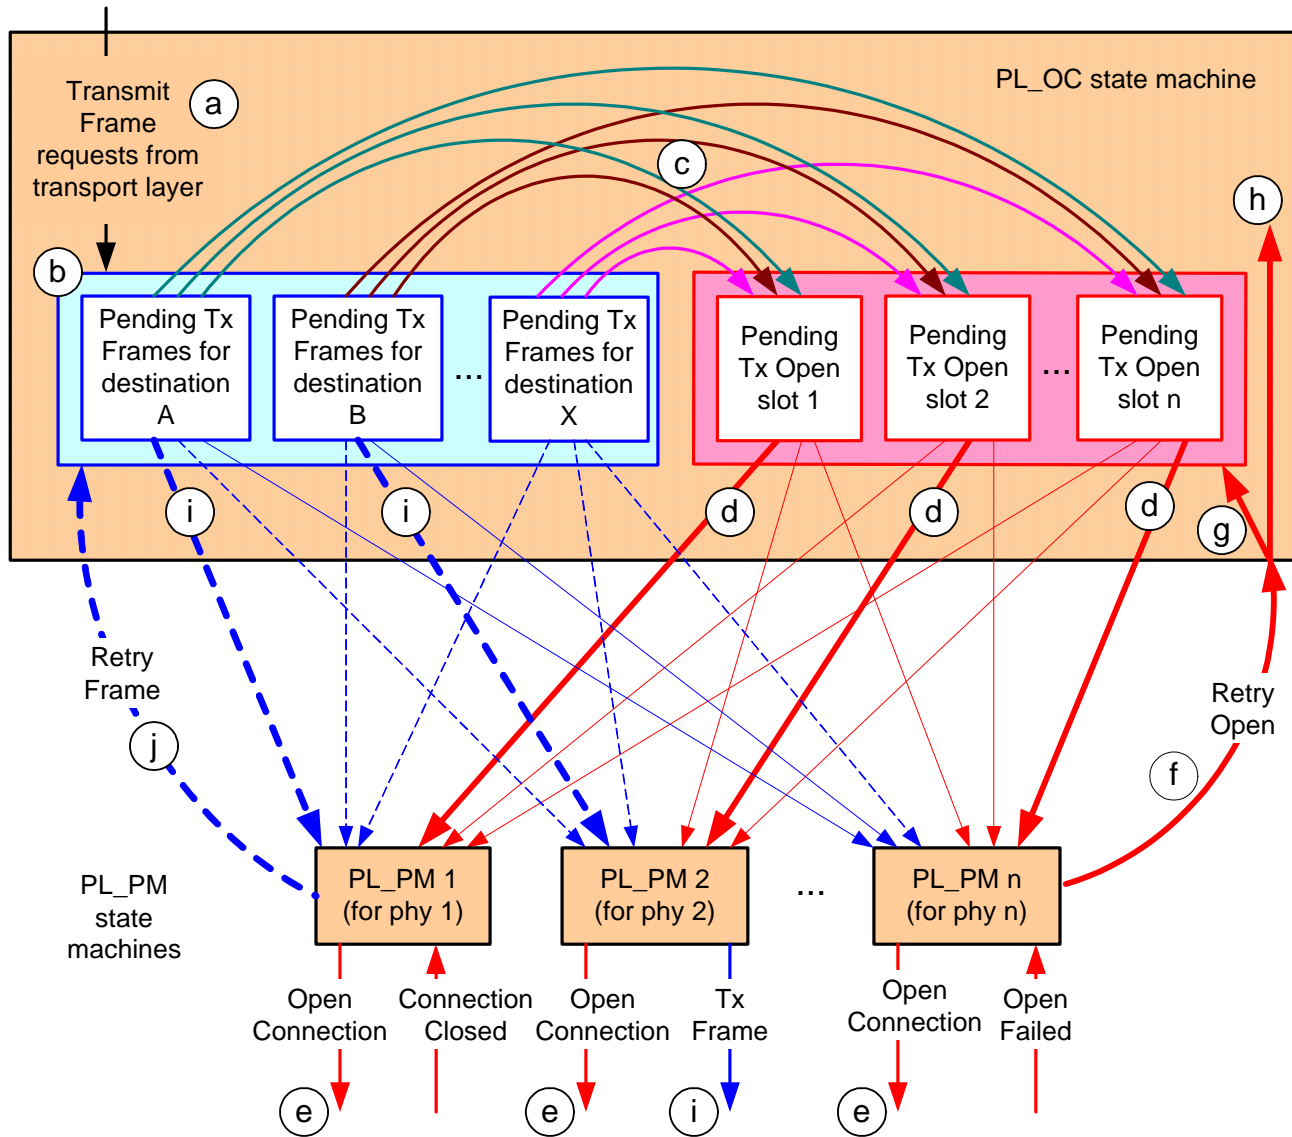
- a) receive requests from the SSP, SMP, and STP transport layer state machines for connection management (e.g., requests to open or close connections) and frame transmission;
- b) send requests to the SAS link layer state machines for connection management and frame transmission;
- c) receive confirmation from the SAS link layer state machines; and
- d) send confirmations to the SSP, SMP, and STP transport layer state machines.

The port layer state machines are as follows:

- a) PL\_OC (port layer overall control) state machines (see 8.1.2); and
- b) PL\_PM (port layer phy manager) state machines (see 8.2.3).

There is one PL\_OC state machine per port (see 4.1.4). There is one PL\_PM state machine for each phy contained in the port. Phys are assigned to ports by the management application layer. More than one port in a SAS device may have the same SAS address.

Figure 1 shows examples of the port layer state machines and their interaction with the transport and link layers.



**Figure 1 — Port layer examples**

The following is a description of the example processes in figure 1. These example processes do not describe all of the possible condition or actions.

- Transmit Frame requests are received by the PL\_OC state machine;
- the PL\_OC state machine converts Transmit Frame requests into pending Tx Frame messages associated with the destination SAS address;
- the PL\_OC state machine generates a pending Tx Open message for a pending Tx Frame message when there is a pending Tx Open slot available (i.e., the number of pending Tx Open messages is less than or equal to the number of phys);
- the PL\_OC state machine sends a pending Tx Open message as a Tx Open message to a PL\_PM state machine when a PL\_PM machine is available; a slot is then available for a new pending Tx Open message;
- when a PL\_PM state machine receives a Tx Open message, the PL\_PM state machine attempts to establish a connection with the destination SAS address through the link layer;
- if a PL\_PM state machine is unable to establish a connection with the destination SAS address, then the PL\_PM state machine sends a Retry Open message to the PL\_OC state machine;

- g) [after a reject to open limit timer has expired, if any](#), the PL\_OC state machine converts a Retry Open message to a pending Tx Open message if there is a pending Tx Open slot available;

---

---

Editor's Note 1: The current item g should be split into two entries at the ;

---

---

- h) if the PL\_OC state machine converts a Retry Open message into a pending Tx Open message, then the pathway blocked count and arbitration wait time context from the Retry Open message are applied to the pending Tx Open message;
- i) if the PL\_OC state machine does not convert a Retry Open to a pending Tx Open frame, then the PL\_OC discards the Retry Open message. The PL\_OC state machine may create a new Tx Open message for the same pending Tx Frame at a later time. If the PL\_OC state machine discards a Retry Open message, then the pathway blocked count and arbitration wait time context from the Retry Open message are also discarded;
- j) after a PL\_PM state machine establishes a connection with a destination SAS address, the PL\_OC state machine sends pending Tx Frame messages for the destination to the PL\_PM state machine as Tx Frame messages;
- k) if a PL\_PM state machine is unable to send a Tx Frame message to the link layer as a Tx Frame request (e.g., due to a credit timeout), then the PL\_PM state machine sends a Retry Frame message to the PL\_OC state machine, and the PL\_OC state machine converts the Retry Frame message into a pending Tx Frame message; and
- l) if the PL\_PM state machine is able to send a Tx Frame message as a Tx Frame request to the link layer, then the PL\_PM state machine sends a Transmission Status confirmation to the transport layer.

The Transmission Status confirmation from either the PL\_OC state machine or a PL\_PM state machine shall include the following as arguments:

- a) tag;
- b) destination SAS address; and
- c) source SAS address.

## 8.1.2 PL\_OC (port layer overall control) state machine

### 8.1.2.1 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; [and](#)
- j) [may have Reject To Open Limit timers.](#)

This state machine consists of the following states:

- a) PL\_OC1:Idle (see 8.1.2.2) (initial state); and
- b) PL\_OC2:Overall\_Control (see 8.1.2.3).

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: <ul style="list-style-type: none"> <li>a) for SSP target ports, the value in the I_T NEXUS LOSS TIME field in the Protocol-Specific Port mode page (see 10.2.7.4);</li> <li>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 10.2.7.4);</li> <li>c) for STP target ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function (see 10.4.3.18);</li> <li>d) for STP initiator ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function (see 10.4.3.18) for the STP target port with that destination SAS address; or</li> <li>e) for SMP initiator ports, the value in the STP SMP I_T NEXUS LOSS TIME field in the SMP CONFIGURE GENERAL function (see 10.4.3.18).</li> </ul>
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.
<a href="#">Open to reject limit timer</a>	<a href="#">One per Retry Open message</a>	<a href="#">Depending on the protocol used by the port:</a> <ul style="list-style-type: none"> <li>a) <a href="#">for SSP target ports, the value in the OPEN TO REJECT LIMIT field in the Shared Port Control mode page (see 10.2.7.6);</a></li> <li>b) <a href="#">for SSP initiator ports, the value in the OPEN TO REJECT LIMIT field in the Shared Port Control mode page for the SSP target port with that destination SAS address (see 10.2.7.6);</a></li> <li>c) <a href="#">for STP target ports, the value in the STP SMP OPEN TO REJECT LIMIT field in the SMP CONFIGURE GENERAL function (see 10.4.3.18);</a></li> <li>d) <a href="#">for STP initiator ports, the value in the STP SMP OPEN TO REJECT LIMIT field in the SMP CONFIGURE GENERAL function (see 10.4.3.18) for the STP target port with that destination SAS address; or</a></li> <li>e) <a href="#">for SMP initiator ports, the value in the STP SMP OPEN TO REJECT LIMIT field in the SMP CONFIGURE GENERAL function (see 10.4.3.18).</a></li> </ul>

Figure 2 shows the PL\_OC state machine.

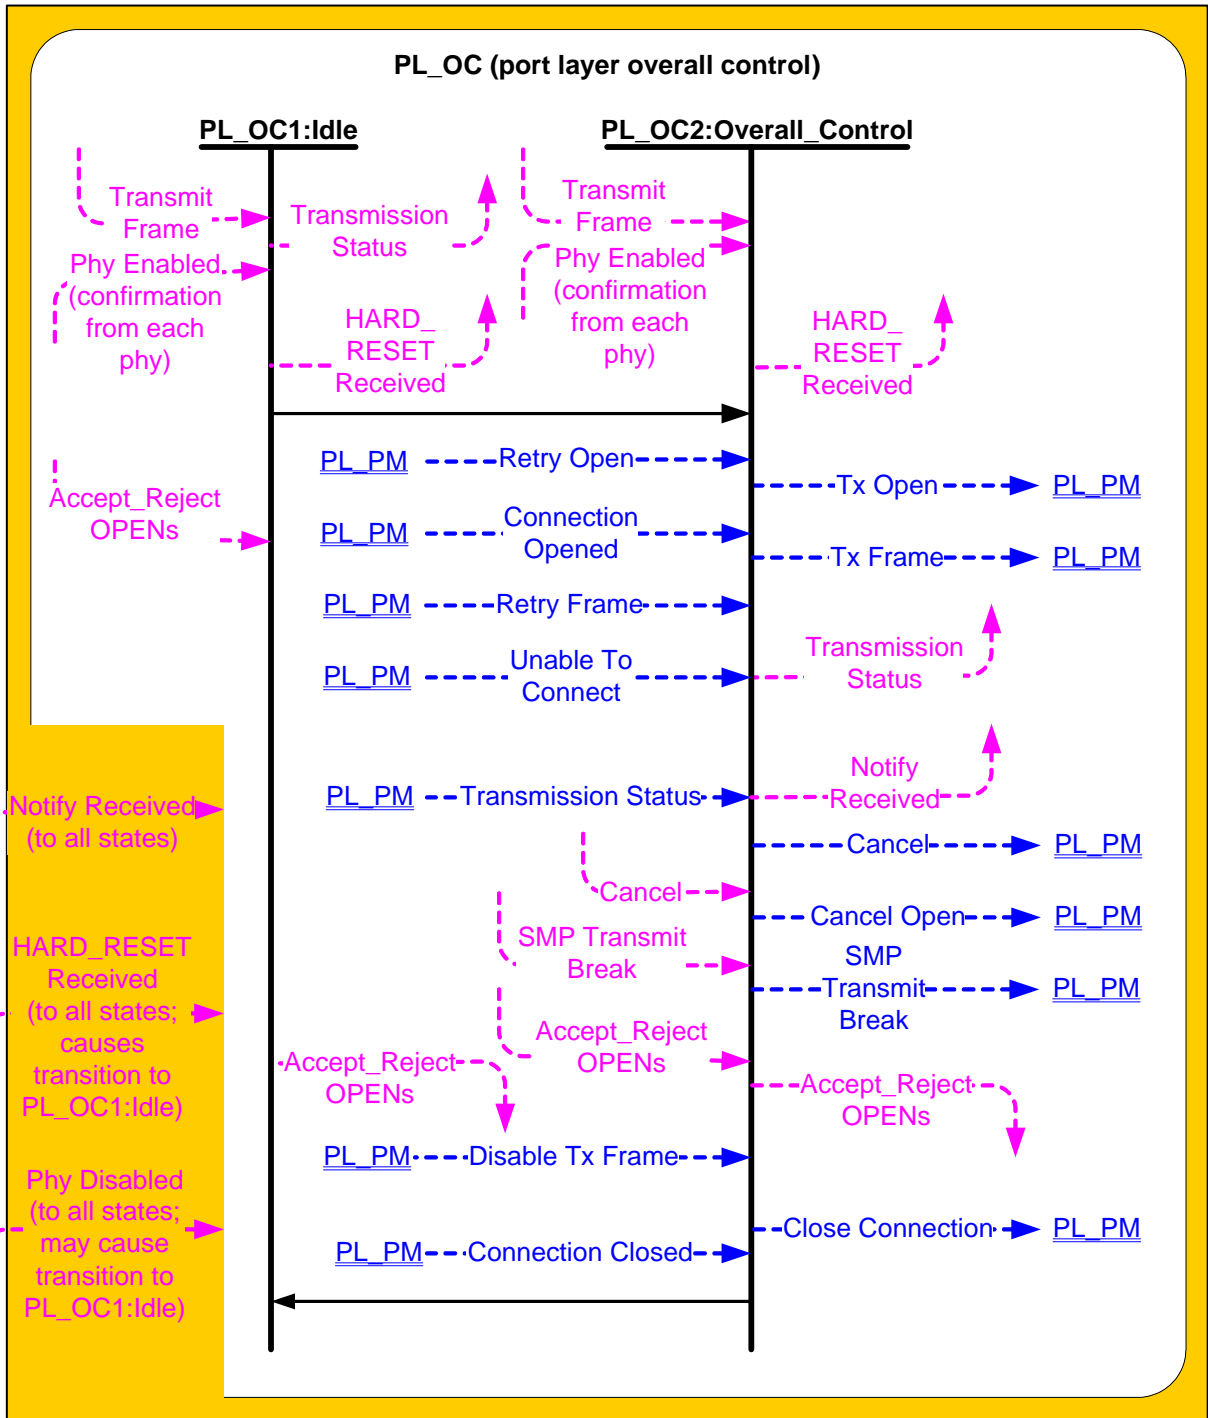


Figure 2 — PL\_OC (port layer overall control) state machine

### 8.1.2.2 PL\_OC1:Idle state

#### 8.1.2.2.1 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 Notify Received (Power Loss Expected) confirmation, then this state shall send a Notify Received (Power Loss Expected) confirmation to the transport layer.

If this state receives an Accept\_Reject Opens request, then this state shall send an Accept\_Reject Opens request to all link layers in the port.

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.

#### **8.1.2.2.2 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.

#### **8.1.2.3 PL\_OC2:Overall\_Control state**

##### **8.1.2.3.1 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, 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 Protocol-Specific Port mode page (see 10.2.7.4) 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 an SSP initiator port; or;
- b) the protocol is STP and the port is an STP initiator port.

When this state creates an I\_T Nexus Loss timer it shall:

- 1) initialize the I\_T Nexus Loss timer as specified in table 1 (see 8.1.2.1); and
- 2) not start the I\_T Nexus Loss timer.

If this state machine is in an SSP initiator port, 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 10.2.7.4) as the initial value for its I\_T Nexus Loss timer.

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.

If this state receives a Notify Received (Power Loss Expected) confirmation, then this state shall:

- a) discard all pending Tx Frame messages, if any;
- b) delete all I\_T Nexus Loss timers, if any;
- c) send a Close Connection message to all the PL\_PM state machines;
- d) send a Cancel Open message to all the PL\_PM state machines; and
- e) send a Notify Received (Power Loss Expected) confirmation to the transport layer.

#### **8.1.2.3.2 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, the I\_T nexus loss time is not set to FFFFh, and the CONFIGURING bit is set to zero in the REPORT GENERAL response (see 10.4.3.3) for each expander device between this port and the destination port that is two or more levels away from this port, then this state shall start the timer;
- b) 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;
- c) if the I\_T Nexus Loss timer is running, then this state shall not stop the timer; and
- d) 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:

- 1) create an Arbitration Wait Time timer for the pending Tx Open message;
- 2) set the Arbitration Wait Time timer for the pending Tx Open message to the arbitration wait time argument from the Retry Open message; and
- 3) 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.

If a pending Tx Open message was created as the result this state receiving a Retry Open (Retry) message and:

- a) the protocol for the connection is SSP, the Shared Port Control mode page is implemented, and OPEN TO REJECT LIMIT field in the Shared Port Control mode page (see 10.2.7.6) is not set to zero;
- b) the protocol for the connection is SMP and the SMP STP OPEN TO REJECT LIMIT field is not set to zero in the SMP REPORT GENERAL response (see 10.4.3.3); or
- c) the protocol for the connection is STP and the SMP STP OPEN TO REJECT LIMIT field is not set to zero in the SMP REPORT GENERAL response (see 10.4.3.3).

then, this state shall:

- 1) create a Reject To Open Limit timer;
- 2) initialize the Reject To Open Limit timer as specified in table 1;
- 3) start the Reject To Open Limit timer; and
- 4) wait at least until the Reject To Open Limit timer expires before sending a Tx Open message.

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.

### 10.2.7.6 Shared Port Control mode page

The Shared Port Control mode page contains parameters that affect SSP target port operation. If the mode page is implemented by one logical unit in a SCSI target device, it shall be implemented by all logical units in the SCSI target device that support the MODE SELECT or MODE SENSE commands.

The mode page policy (see SPC-4) for this mode page shall be shared.

Table 2 defines the format of this mode page.

**Table 2 — Shared Port Control mode page**

Byte\Bit	7	6	5	4	3	2	1	0
0	PS	SPF (1b)	PAGE CODE (19h)					
1	SUBPAGE CODE (02h)							
2	(MSB)	PAGE LENGTH (000Ch)						(LSB)
3								
4	Reserved							
5	Reserved				PROTOCOL IDENTIFIER (6h)			
6	(MSB)	POWER LOSS TIMEOUT						(LSB)
7								
8								
9	<a href="#">REJECT TO OPEN LIMIT</a>							
10								
15	Reserved							

The PARAMETERS SAVEABLE (PS) bit is defined in SPC-4.

The SUBPAGE FORMAT (SPF) bit shall be set to one to access this mode page.

The PAGE CODE field shall be set to 19h.

The SUBPAGE CODE field shall be set to 02h.

The PAGE LENGTH field shall be set to the number of bytes in the page after the PAGE LENGTH field (i.e., 000Ch).

The PROTOCOL IDENTIFIER field shall be set to 6h indicating this is a SAS SSP specific mode page.

The POWER LOSS TIMEOUT field contains the maximum time, in one millisecond increments, that a target port shall respond to connection requests with OPEN\_REJECT (RETRY) after receiving NOTIFY (POWER LOSS EXPECTED) (see 7.2.5.11.3). A POWER LOSS TIMEOUT field set to 0000h is vendor-specific. The power loss timeout shall be restarted on each NOTIFY (POWER LOSS EXPECTED) that is received.

The REJECT TO OPEN LIMIT field specifies the minimum time in 100 us increments that the target port shall wait between a receiving an OPEN\_REJECT (RETRY) or an OPEN\_REJECT (PATHWAY BLOCK) retry class of OPEN\_REJECT primitives (see table 113xx) and issuing an OPEN address frame on the same I\_T nexus. This value may be rounded as defined in SPC-4. A value of 0000h specifies that the time limit between receiving an OPEN\_REJECT (RETRY) or an OPEN\_REJECT (PATHWAY BLOCK) retry class of OPEN\_REJECT primitives (see table 113xx) and issuing an OPEN address frame on the same I\_T nexus is vendor specific. This time limit is enforced by the port layer (see 8.2.3).

#### **10.4.3.18 CONFIGURE GENERAL function**

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

Table 3 defines the request format.

**Table 3 — 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 (04h)							
4	(MSB)	EXPECTED EXPANDER CHANGE COUNT						(LSB)
5								
6	Reserved							
7								
8	Reserved		<a href="#">UPDATE SMP STP REJECT TO OPEN LIMIT</a>	UPDATE INITIAL TIME TO REDUCED FUNCTIONALITY	UPDATE STP SMP I_T NEXUS LOSS TIME	UPDATE STP MAXIMUM CONNECT TIME LIMIT	UPDATE STP BUS INACTIVITY TIME LIMIT	
9	Reserved							
10	(MSB)	STP BUS INACTIVITY TIME LIMIT						(LSB)
11								
12	(MSB)	STP MAXIMUM CONNECT TIME LIMIT						(LSB)
13								
14	(MSB)	STP SMP I_T NEXUS LOSS TIME						(LSB)
15								
16	INITIAL TIME TO REDUCED FUNCTIONALITY							
17								
18	<a href="#">SMP STP REJECT TO OPEN LIMIT</a>							
19	Reserved							
20	(MSB)	CRC						(LSB)
23								

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

The FUNCTION field shall be set to 80h.

The REQUEST LENGTH field shall be set to 04h.

If the management device server is not in an expander device or the EXPECTED EXPANDER CHANGE COUNT field is set to 0000h, the EXPECTED EXPANDER CHANGE COUNT field shall be ignored. If the management device server is in an expander device and the EXPECTED EXPANDER CHANGE COUNT field is not set to 0000h, then:

- a) if the EXPECTED EXPANDER CHANGE COUNT field contains the current expander change count (i.e., the value of the EXPANDER CHANGE COUNT field that would be returned by an SMP REPORT GENERAL response at this time), the management device server shall process the function; and
- b) If the EXPECTED EXPANDER CHANGE COUNT field does not contain the current expander change count, the management device server shall return a function result of INVALID EXPANDER CHANGE COUNT in the response frame (see table 315 in 10.4.3.2).

An UPDATE SMP STP REJECT TO OPEN LIMIT bit set to one specifies that the SMP STP REJECT TO OPEN LIMIT field shall be honored. An UPDATE SMP STP REJECT TO OPEN LIMIT bit set to zero specifies that the SMP STP REJECT TO OPEN LIMIT field shall be ignored.

An UPDATE INITIAL TIME TO REDUCED FUNCTIONALITY bit set to one specifies that the INITIAL TIME TO REDUCED FUNCTIONALITY field shall be honored. An UPDATE INITIAL TIME TO REDUCED FUNCTIONALITY bit set to zero specifies that the INITIAL TIME TO REDUCED FUNCTIONALITY field shall be ignored.

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 specifies the maximum time in 100 μs increments that an STP target port is permitted to maintain a connection (see 4.1.12) while transmitting and receiving SATA\_SYNC. When this time is exceeded, the STP target port shall close the connection. A value of 0000h in this field specifies that there is no bus inactivity time limit. This value is reported in the STP BUS INACTIVITY TIME LIMIT field in the SMP REPORT GENERAL response (see 10.4.3.3). The bus inactivity time limit is enforced by the port layer (see 8.2.3).

The STP MAXIMUM CONNECT TIME LIMIT field specifies the maximum duration of a connection (see 4.1.12) in 100 μs increments (e.g., a value of 0001h in this field means that the time is less than or equal to 100 μs and a value of 0002h in this field means that the time is less than or equal to 200 μs). When this time is exceeded, the STP target port shall close the connection at the next opportunity. If the 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 0000h in this field specifies that there is no maximum connection time limit. This value is reported in the STP MAXIMUM CONNECT TIME LIMIT field in the SMP REPORT GENERAL response (see 10.4.3.3). The maximum connection time limit is enforced by the port layer (see 8.2.3).

The STP SMP I\_T NEXUS LOSS TIME field specifies the time that an STP target port or SMP initiator port shall retry connection requests that are rejected with responses indicating the destination port may no longer be present (see 8.2.2) before recognizing an I\_T nexus loss (see 4.5). Table 4 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 4 — STP SMP I\_T NEXUS LOSS TIME field**

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

NOTE 2 - 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 3 - An STP initiator port should retry connection requests for the time indicated by the STP SMP I\_T NEXUS LOSS TIME field in the SMP REPORT GENERAL response for the STP target port to which it is trying to establish a connection.

The INITIAL TIME TO REDUCED FUNCTIONALITY field specifies the minimum period of time, in 100 ms increments, that an expander device shall wait from originating a Broadcast (Expander) to reducing functionality (see 4.6.8). This value is reported in the INITIAL TIME TO REDUCED FUNCTIONALITY field in the SMP REPORT GENERAL response (see 10.4.3.3).

The SMP STP REJECT TO OPEN LIMIT field specifies the minimum time in 100 us increments that an STP target port or an SMP initiator port shall wait between a receiving an OPEN\_REJECT (RETRY) or an OPEN\_REJECT (PATHWAY BLOCK) retry class of OPEN\_REJECT primitives (see table 113xx) and issuing an OPEN address frame on the same I\_T nexus. This value may be rounded as defined in SPC-4. A value of 0000h specifies that the time limit between receiving an OPEN\_REJECT (RETRY) or an OPEN\_REJECT (PATHWAY BLOCK) retry class of OPEN\_REJECT primitives (see table 113xx) and issuing an OPEN address frame on the same I\_T nexus is vendor specific. This time limit is enforced by the port layer (see 8.2.3).

The CRC field is defined in 10.4.3.1.