

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 8 December 2006
 Subject: 06-322r3 SAS-2 Response to abandon-class OPEN_REJECT

Revision history

Revision 0 (7 July 2006) First revision

Revision 1 (11 August 2006) Incorporated comments from July SAS protocol WG - made some corrections and added some new rules concerning OPEN_REJECT (CONNECTION RATE NOT SUPPORTED) handling for STP/SATA bridges.

Revision 2 (30 October 2006) Incorporated comments from Vicky Duerk, Intel, and the September SAS protocol WG.

Revision 3 (8 December 2006) Incorporated comments from November SAS protocol WG.

Related documents

sas2r07 - Serial Attached SCSI - 2 (SAS-2) revision 7

Serial ATA 2.5 (not publicly available; see www.sata-io.org for details)

Overview

1. If an STP target port receives an abandon-class OPEN_REJECT such as OPEN_REJECT (PROTOCOL NOT SUPPORTED), it needs to stop trying to open that initiator port and abort any outstanding commands from that STP initiator port. If the STP target port is in an STP/SATA bridge (in an expander), the way it aborts those commands is by performing a link reset sequence on the phy attached to the SATA device. If an STP initiator port receives an abandon-class OPEN_REJECT, it needs to consider all commands to that target as completed with an error.

This is the same as if an I_T nexus loss occurs, so the I_T nexus loss definition is expanded to including receiving an abandon-class OPEN_REJECT.

2. OPEN_REJECT (CONNECTION RATE NOT SUPPORTED) presents additional problems. It is labeled as abandon-class, but the request is not really abandoned unless it requested a 1.5 Gbps connection rate; if it was higher than that, then it is reattempted with a lower connection rate. This is clarified throughout.

3. The CONNECTION RATE field selection rules for both SSP and STP target ports are adjusted to prohibit infinite retries at connection rates that are not working (literally allowed by the current rules). STP target ports are allowed the SSP-style option of remembering rates per command; SSP target ports are allowed the STP-style option of remembering the last known good connection rate with the initiator.

4. "SATA rate matching."

If a SATA physical link is faster than the SAS connection rate (which could be the cause of an OPEN_REJECT (CONNECTION RATE NOT SUPPORTED), then the STP/SATA bridge is going to be overrun by the incoming data from the SATA drive.

The existing rule that the SATA physical link rate shall be set the same as the highest possible connection rate on the potential pathways to the STP initiator port doesn't help the error case and doesn't help if multiple initiators are present with different physical link rates. This rule is proposed to be constrained to single-initiator cases or eliminated altogether (assuming "SATA rate matching" described below works).

The only defense the STP/SATA bridge can offer is to send SATA_HOLD; each time it releases SATA_HOLD for n dwords, up to 21+n more incoming dwords are unleashed. SATA 1.0/1.0a HT and DT state machines originally stated that HOLD was not allowed in either direction for non-Data FISes (e.g., register, DMA Setup, Set Device Bits, PIO Setup) but SATA 2.5 seems to have dropped that restriction. So, SATA_HOLD seems like a viable solution. Although this will create extra HOLD/HOLDA turnarounds on the SATA physical link, it is only consuming bandwidth that is not achievable

The STP/SATA bridge still must handle incoming primitives without losing any important information.

The outgoing direction (SAS to SATA) is easier; the STP/SATA bridge just rate matches, using either ALIGN (0) (but they must be transmitted in pairs) or SATA_HOLD.

Suggested changes

3.1.33 connection rate: The effective rate of dwords through the pathway between a SAS initiator phy and a SAS target phy, established through the connection request.

3.1.92 I_T nexus: When referring to SAS ports (see 3.1.189), aA nexus that exists between a SAS initiator port and a SAS target port. When referring to SCSI ports (see 3.1.207)(e.g., SSP ports), a relationship between a SCSI initiator port and a SCSI target port. See SAM-4.

3.1.93 I_T nexus loss: When referring to SAS ports (see 3.1.189), aA condition where a SAS port determines that another SAS port is no longer available. See 4.5. When referring to SCSI ports (see 3.1.207)(e.g., SSP ports), a condition resulting from the events defined by SAM-4 in which the SCSI device performs the I_T nexus loss operations described in SAM-4, SPC-4, and the appropriate command standards. See SAM-4.

3.1.94 I_T_L nexus: A nexus that exists between a SCSI initiator port, a SCSI target port, and a logical unit. This relationship extends the prior I_T nexus. See SAM-4.

3.1.95 I_T_L_Q nexus: A nexus between a SCSI initiator port, a SCSI target port, a logical unit, and a task. This relationship extends the prior I_T nexus or I_T_L nexus. See SAM-4.

3.1.137 nexus: When referring to SAS ports (see 3.1.189), aA relationship between a SAS initiator port and a SAS target port. When referring to SCSI ports (see 3.1.207)(e.g., SSP ports), a relationship between a SCSI initiator port and a SCSI target port that may extend to a logical unit and a task. See SAM-4.

3.1.147 pathway blocked count: The number of times the port has retried this connection request due to receiving OPEN_REJECT (PATHWAY BLOCKED), OPEN_REJECT (RESERVED STOP 0), or OPEN_REJECT (RESERVED STOP 1).

3.1.206 SCSI initiator port: A SCSI initiator device object that acts as the connection between SCSI application clients and a service delivery subsystem through which indications and responses are routed. See SAM-4.

3.1.207 SCSI port: A SCSI initiator port, SCSI target port, or SCSI target/initiator port. See SAM-4.

3.1.209 SCSI target port: A SCSI target device object that contains a task router and acts as the connection between SCSI device servers and task managers and a service delivery subsystem through which requests and confirmations are routed. See SAM-4.

4.1.12 Connections

...

The connection rate is the effective rate of dwords through the pathway between a SAS initiator phy and a SAS target phy, established through the connection request. Every logical phy shall support a 1,5 Gbps connection rate regardless of its logical link rate.

...

4.5 I_T nexus loss

When a SAS port receives OPEN_REJECT (NO DESTINATION), OPEN_REJECT (PATHWAY BLOCKED), OPEN_REJECT (RESERVED INITIALIZE 0), OPEN_REJECT (RESERVED INITIALIZE 1), OPEN_REJECT (RESERVED STOP 0), OPEN_REJECT (RESERVED STOP 1), 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 expires; or
- c) the I_T nexus loss timer, if any, expires (see 8.2.2.1, 10.2.7.2, and 10.4.3.14).

An I_T nexus loss occurs in a SAS port when:

- a) the I_T nexus loss timer expires; or
- b) the SAS port receives an abandon-class OPEN_REJECT (see table 90 in 7.2.5.12).

I_T nexus loss is handled by the port layer state machines (see 8.2.2.3). In some cases, the I_T nexus loss timer is overridden for connection requests through self-configuring expander devices as described in 4.7.1.

If ~~the I_T nexus loss timer expires~~ an I_T nexus loss occurs 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. If an I_T nexus loss occurs in an SSP initiator port, a SCSI application client should send an I_T NEXUS RESET task management function to the SSP target port when it reestablishes communication.

If ~~the I_T nexus loss timer expires~~ an I_T nexus loss occurs in an STP initiator port, then the STP initiator port shall send a Transport Event Notification (Nexus Loss, [Device]) indication to the portATA application client (i.e., create a nexus loss event)(see ATA8-AAM). The ATA application client shall consider any commands for the lost STP target port to be completed with an error. ~~This serves as a nexus loss event (see ATA8-AAM).~~

If ~~the I_T nexus loss timer expires~~ an I_T nexus loss occurs in an STP target port, then the ~~port~~ the ATA device server shall abort any-all outstanding commands for the lost STP initiator port. 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 so the ATA device server in the SATA device aborts all outstanding commands.

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 due to I_T nexus loss timer expiration, 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).

4.6.4 Expander connection router (ECR)

The ECR routes messages between pairs of expander logical phys as configured by the ECM. Enough routing resources shall be provided to support at least one connection.

When forwarding dwords during a connection from a source phy with a higher physical link rate to a destination phy with a lower physical link rate, rate matching (see 7.13) ensures the dwords are at a connection rate equal to or less than the lower physical link rate. However, the ECR may be requested to forward more dwords than the destination phy is able to accept if:

- a) an invalid dword occurs during a deletable primitive;
- b) an invalid dword occurs during a CLOSE; or
- c) multiple invalid dwords occur during a BREAK.

The ECR may discard dwords if needed and count them as receive elasticity buffer overflows (see 4.11).

When forwarding dwords from a SATA physical link with a higher physical link rate to a SAS physical link with a lower physical link rate, the SATA host port in the STP/SATA bridge shall throttle incoming FISes with SATA_HOLD (see 7.17.3).

NOTE 1 - If SATA_HOLD is deasserted for n dwords, the SATA device sends up to (21+n) more data dwords.

SATA_HOLD does not affect primitives (see ATA/ATAPI-7 V3). The STP/SATA bridge may expand or contract repeated and continued primitives without changing their functional meaning.

When forwarding dwords from a SAS logical link with a lower logical link rate to a SATA physical link with a higher physical link rate, the SATA host port in the STP/SATA bridge shall perform a process similar to rate matching (see 7.13) by inserting ALIGN (0) and/or SATA_HOLD on the SATA physical link whenever it underflows.

NOTE 2 - SATA requires that ALIGN (0) be sent in pairs (see ATA/ATAPI-7 V3).

4.8.3 Route table optimization

...

If the management application client supports route table optimization, then the management application client should provide a vendor-specific method for initiating a check of the resulting expander route tables. The check should be performed under the following situations:

- a) when an **OPEN_REJECT (NO DESTINATION)** is received for a connection request to a SAS address-
[L T nexus loss occurs for a destination port](#) that is expected to be **accepted** present;
- b) when a discover process has been completed;
- c) when another SMP initiator port is discovered in the SAS domain; or
- d) when a self-configuring expander device is discovered in the SAS domain.

4.11 Phy event information

...

Table 1 — PHY EVENT INFORMATION SOURCE field

Code	Name	Type ^a	Description
...
21h	Transmitted abandon-class OPEN_REJECT abandon count	WC	Number of times the phy received an OPEN address frame and transmitted an abandon-class OPEN_REJECT (see 7.2.5.12). In expander devices, forwarded OPEN_REJECTs shall not be counted.
22h	Received abandon-class OPEN_REJECT abandon count	WC	Number of times the phy originated an OPEN address frame and received an abandon-class OPEN_REJECT (see 7.2.5.12). In expander devices, OPEN_REJECTs in response to forwarded OPEN address frames shall not be counted.
23h	Transmitted retry-class OPEN_REJECT retry count	WC	Number of times the phy received an OPEN address frame and transmitted a retry-class OPEN_REJECT (see 7.2.5.12). In expander devices, forwarded OPEN_REJECTs shall not be counted.
24h	Received retry-class OPEN_REJECT retry count	WC	Number of times the phy originated an OPEN address frame and received a retry-class OPEN_REJECT (see 7.2.5.12). In expander devices, OPEN_REJECTs in response to forwarded OPEN address frames shall not be counted.
^a The Type column indicates the source type: a) WC = wrapping counter b) PVD = peak value detector c) N/A = not applicable ^b This standard also defines a saturating counter that counts this event (see 10.2.8.1 and 10.4.3.7).			

7.2.5.12 OPEN_REJECT

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

All of the OPEN_REJECT versions defined in table 2 shall result in the originating port abandoning the connection request.

Table 2 — Abandon-class OPEN_REJECT ~~abandon~~-primitives

Primitive	Originator	Description
OPEN_REJECT (BAD DESTINATION)	Expander phy	A connection request arrives through an expander phy using the direct routing or table routing method and the expander device determines the connection request would have to be routed to the same expander port as the expander port through which the connection request arrived (e.g., the destination SAS address equals the source SAS address), and the expander device has not chosen to return OPEN_REJECT (NO DESTINATION) (see 7.12.4.3 and 7.12.4.4).
OPEN_REJECT (CONNECTION RATE NOT SUPPORTED)	Any phy	The requested connection rate is not supported on some physical link on the pathway between the source phy and destination phy. When a SAS initiator phy is directly attached to a SAS target phy, the requested connection rate is not supported by the destination phy. <u>If the connection rate is 1,5 Gbps, this shall be considered an abandon-class OPEN_REJECT.</u> <u>If the connection rate is greater than 1,5 Gbps, the connection request may be shall be modified and reattempted as described in 7.12.2.2-7.12.2.3.</u>
OPEN_REJECT (PROTOCOL NOT SUPPORTED)	Destination phy	Phy with destination SAS address exists but the destination phy does not support the requested initiator/target role, protocol, initiator connection tag, or features (i.e., the values in the INITIATOR PORT bit, the PROTOCOL field, the INITIATOR CONNECTION TAG field, and/or the FEATURES field in the OPEN address frame are not supported).
OPEN_REJECT (RESERVED ABANDON 1)	Unknown	Reserved. Process the same as OPEN_REJECT (WRONG DESTINATION).
OPEN_REJECT (RESERVED ABANDON 2)		
OPEN_REJECT (RESERVED ABANDON 3)		
OPEN_REJECT (STP RESOURCES BUSY)	Destination phy	STP target port with destination SAS address exists but the STP target port has an affiliation with another STP initiator port or all of the available task file registers have been allocated to other STP initiator ports (see 7.17.5). Process the same as OPEN_REJECT (WRONG DESTINATION) for non-STP connection requests.
OPEN_REJECT (WRONG DESTINATION)	Destination phy	The destination SAS address does not match the SAS address of the SAS port to which the connection request was delivered.
OPEN_REJECT (ZONE VIOLATION)	Zoning expander phy	The connection request is from a zone group that does not have permission to access the zone group that contains the destination phy.

Editor's Note 1: move ZONE VIOLATION to the bottom of all references to it throughout the standard.

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

Table 3 — Retry-class OPEN_REJECT ~~retry~~ primitives

Primitive	Originator	Description
OPEN_REJECT (NO DESTINATION) ^a	Expander phy	Either: a) No such destination phy; b) the expander device determines the connection request would have to be routed to the same expander port as the expander port through which the connection request arrived (e.g., the destination SAS address equals the source SAS address) and the expander device has not chosen to return OPEN_REJECT (BAD DESTINATION) (see 7.12.4.3 and 7.12.4.4); or c) the SAS address is valid for an STP target port in an STP/SATA bridge, but the initial Register - Device to Host FIS has not been successfully received (see 10.4.3.8).
OPEN_REJECT (PATHWAY BLOCKED) ^b	Expander phy	An expander device determined the pathway was blocked by higher priority connection requests.
OPEN_REJECT (RESERVED CONTINUE 0) ^c	Unknown	Reserved. Process the same as OPEN_REJECT (RETRY).
OPEN_REJECT (RESERVED CONTINUE 1) ^c		
OPEN_REJECT (RESERVED INITIALIZE 0) ^a	Unknown	Reserved. Process the same as OPEN_REJECT (NO DESTINATION).
OPEN_REJECT (RESERVED INITIALIZE 1) ^a		
OPEN_REJECT (RESERVED STOP 0) ^b	Unknown	Reserved. Process the same as OPEN_REJECT (PATHWAY BLOCKED).
OPEN_REJECT (RESERVED STOP 1) ^b		
OPEN_REJECT (RETRY) ^c	Destination phy	Phy with destination SAS address exists but is not able to accept connections.
^a If the I_T Nexus Loss timer is already running, it continues running; if it is not already running, it is initialized and started. Stop retrying the connection request if the I_T Nexus Loss timer expires. ^b If the I_T Nexus Loss timer is already running, it continues running. Stop retrying the connection request if the I_T Nexus Loss timer expires. ^c If the I_T Nexus Loss timer (see 8.2.2) is already running, it is stopped.		

NOTE 3 - Some SAS logical phys compliant with earlier versions of this standard also transmit OPEN_REJECT (RETRY) if they receive an OPEN address frame while their SL_CC state machines are in the SL_CC5:BreakWait state (see 7.14.4.7).

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

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

See 7.12 for details on connection requests.

7.8.3 OPEN address frame

Table 115 defines the OPEN address frame format used for connection requests.

...

The CONNECTION RATE field specifies the connection rate (see 4.1.11) being requested between the source and destination, and is defined in table 112.

Table 112 — CONNECTION RATE field

Code	Description
8h	1,5 Gbps
9h	3 Gbps
Ah	6 Gbps
Bh - Fh	Reserved for future connection rates
All others	Reserved

~~When requesting a connection to a SAS target port, a SAS initiator port shall should set the CONNECTION RATE field to a value supported by at least one potential pathway.~~

~~When requesting an SSP connection to an SSP initiator port, an SSP target port shall set the CONNECTION RATE field to the connection rate in effect when the command was received unless it has received an OPEN_REJECT (CONNECTION RATE NOT SUPPORTED). See 7.12.2.27.12.2.3 for details on handling OPEN_REJECT (CONNECTION RATE NOT SUPPORTED).~~

A SAS initiator port shall set the initial CONNECTION RATE field to:

- a) the highest supported connection rate supported by a potential pathway as determined during the discover process (e.g., based on the logical link rates of each logical link reported in the SMP DISCOVER responses); or
- b) the logical link rate of the logical phy used to transmit the OPEN address frame.

If a SAS initiator port expects a connection rate to work but the connection request results in OPEN_REJECT (CONNECTION RATE NOT SUPPORTED), the discover process information is no longer current and should be run again.

A SAS target port shall set the initial CONNECTION RATE field to:

- a) the last known good connection rate established with the SAS initiator port; or
- b) for the first frame it intends to transmit in the connection, the connection rate that was used by the SAS initiator port to deliver the command or task management function for that frame.

Each time a connection request with a connection rate greater than 1,5 Gbps results in OPEN_REJECT (CONNECTION RATE NOT SUPPORTED), the SAS port shall reattempt the connection request with a lower connection rate (e.g., drop from 6 Gbps to 3 Gbps or 1,5 Gbps) and send the same frames in the resulting connection that it intended to send at the initial connection rate.

~~The SAS target port should send frames in a connection to the SAS initiator port regardless of whether the saved connection rate for that command matches the current connection rate; the SAS target port should not close the connection just to reopen the connection at the saved connection rate.~~

~~When requesting an STP connection to an STP initiator port, an STP target port shall set the CONNECTION RATE field to the last value received in a connection request from the STP initiator port unless the STP target port has received an OPEN_REJECT (CONNECTION RATE NOT SUPPORTED).~~

...

The PATHWAY BLOCKED COUNT field specifies the number of times the port has retried this connection request due to receiving OPEN_REJECT (PATHWAY BLOCKED), [OPEN_REJECT \(RESERVED STOP 0\)](#), or [OPEN_REJECT \(RESERVED STOP 1\)](#). The port shall not increment the PATHWAY BLOCKED COUNT value past FFh. If the port changes connection requests, it shall use a PATHWAY BLOCKED COUNT [offset set to 00h](#).

7.12.2 Opening a connection

7.12.2.1 Connection request

The OPEN address frame (see 7.8.3) is used to open a connection from a source port to a destination port using one source phy and one destination phy.

To make a connection request, the source port shall transmit an OPEN address frame through an available phy. The source phy shall transmit idle dwords after the OPEN address frame until it receives a response or aborts the connection request with BREAK.

After transmitting an OPEN address frame, the source phy shall initialize and start a 1 ms Open Timeout timer. Whenever an AIP is received, the source phy shall reinitialize and restart the Open Timeout timer. Source phys are not required to enforce a limit on the number of AIPs received before aborting the connection request. When any connection response is received, the source phy shall reinitialize the Open Timeout timer. If the Open Timeout timer expires before a connection response is received, the source phy shall transmit BREAK to abort the connection request (see 7.12.6).

The OPEN address frame flows through expander devices onto intermediate physical links. If an expander device on the pathway is unable to forward the connection request, it returns OPEN_REJECT (see 7.12.4). If the OPEN address frame reaches the destination, it returns either OPEN_ACCEPT or OPEN_REJECT unless the OPEN address frame passed an OPEN address frame from the destination with higher arbitration priority (see 7.12.3). Rate matching shall be used on any physical links in the pathway with negotiated physical link rates that are faster than the requested connection rate (see 7.13).

7.12.2.2 Results of a connection request

After a phy transmits an OPEN address frame, it shall expect one or more of the results listed in table 113.

Table 113 — Connection Results of a connection request

Result	Description
Receive AIP	Arbitration in progress. When an expander device is trying to open a connection to the selected destination port, it returns an AIP to the source phy. The source phy shall reinitialize and restart its Open Timeout timer each time it receives an AIP. AIP is sent by an expander device while it is internally arbitrating for access to an expander port.
Receive OPEN_ACCEPT	Connection request accepted. OPEN_ACCEPT is transmitted by the destination phy.
Receive OPEN_REJECT	Connection request rejected. OPEN_REJECT is transmitted by the destination phy or by an expander device in the partial pathway. The different versions are described in 7.2.5.12. See 4.5 for I_T nexus loss handling. See 7.8.3 for handling of OPEN_REJECT (CONNECTION RATE NOT SUPPORTED) for connection rates greater than 1.5 Gbps.
Receive OPEN address frame	If AIP has been previously detected, this indicates an overriding connection request. If AIP has not yet been detected, this indicates two connection requests crossing on the physical link. Arbitration fairness determines which one wins (see 7.12.3).
Receive BREAK	The destination phy or an expander device in the partial pathway may reply with BREAK indicating the connection is not being established.
Open Timeout timer expires	The source phy shall abort the connection request by transmitting BREAK (see 7.12.6). See 4.5 for I_T nexus loss handling.

~~After an OPEN_REJECT (CONNECTION RATE NOT SUPPORTED) has been received by a SAS target port, the SAS target device shall set the connection rate for future requests for that I_T_L_Q nexus to:~~

- ~~a) the last value received in a connection request from the SAS initiator port;~~
- ~~b) 1,5-Gbps; or~~
- ~~c) the connection rate in effect when the command was received.~~

7.13 Rate matching

Each successful connection request contains the connection rate (see 4.1.12) of the pathway.

Each logical phy in the pathway shall insert deletable primitives between dwords if its logical link rate is faster than the connection rate as described in table 114.

Table 114 — Rate matching deletable primitive insertion requirements

Logical link rate	Connection rate	Requirement
1,5 Gbps	1,5 Gbps	None
3 Gbps	1,5 Gbps	One deletable primitive within every 2 dwords that are not physical link rate tolerance management deletable primitives (i.e., every overlapping window of 2 dwords)(e.g., a repeating pattern of a deletable primitive followed by a dword, or a repeating pattern of a dword followed by an deletable primitive)
	3 Gbps	None
6 Gbps	1,5 Gbps	Three deletable primitivse within every 4 dwords that are not physical link rate tolerance management deletable primitives (i.e., 3 in every overlapping window of 4 dwords)
	3 Gbps	One deletable primitive within every 2 dwords that are not physical link rate tolerance management deletable primitives (i.e., every overlapping window of 2 dwords)(e.g., a repeating pattern of a deletable primitive followed by a dword, or a repeating pattern of a dword followed by an deletable primitive)
	6 Gbps	None

Deletable primitives inserted for rate matching are in addition to deletable primitives inserted for physical link rate tolerance management (see 7.3). See Annex H for a summary of their combined requirements.

Figure 1 shows an example of rate matching between a 3 Gbps source phy and a 3 Gbps destination phy, with an intermediate 1,5 Gbps physical link in between them.

Figure 1 — Rate matching example [no change]

A phy originating dwords shall start rate matching at the selected connection rate starting with the first dword that is not a deletable primitive inserted for physical link rate tolerance management following:

- a) transmitting the EOAF for an OPEN address frame; or
- b) transmitting an OPEN_ACCEPT.

An expander phy forwarding dwords shall not insert deletable primitives for rate matching based on counting dwords transmitted. It shall insert deletable primitives whenever it underflows.

The source phy transmits idle dwords including deletable primitives at the selected connection rate while waiting for the connection response. This enables each expander device to start forwarding dwords from the source phy to the destination phy after forwarding an OPEN_ACCEPT.

A phy shall stop inserting deletable primitives for rate matching after:

- a) transmitting the first dword in a CLOSE;
- b) transmitting the first dword in a BREAK;
- c) transmitting the first dword in a BREAK_REPLY;
- d) receiving an OPEN_REJECT for a connection request; or
- e) losing arbitration to a received OPEN address frame.

~~If an expander phy attached to a SATA phy is using a physical link rate greater than the maximum connection rate supported by the pathway from an STP initiator port, a management application client should use the SMP_PHY_CONTROL function (see 10.4.3.18) to set the PROGRAMMED_MAXIMUM_PHYSICAL_LINK_RATE field of the expander phy to the maximum connection rate supported by the pathway from that STP initiator port.~~

7.16 SSP link layer

7.16.1 Opening an SSP connection

An SSP phy that accepts an OPEN address frame shall transmit at least one RRDY in that connection within 1 ms of transmitting an OPEN_ACCEPT. If the SSP phy is not able to grant credit, it shall respond with OPEN_REJECT (RETRY) and not accept the connection request.

7.17 STP link layer

7.17.6 Opening an STP connection

If no STP connection exists when the SATA host port in an STP/SATA bridge receives a SATA_X_RDY from the attached SATA device, the STP target port in the STP/SATA bridge shall establish an STP connection to the appropriate STP initiator port before it transmits a SATA_R_RDY to the SATA device.

Wide STP initiator ports shall not request more than one connection at a time to an STP target port. Wide STP target ports shall not request more than one connection at a time to an STP initiator port.

While a wide STP target port is waiting for a response to a connection request or has established a connection to an STP initiator port, it shall:

- a) reject incoming connection requests from that STP initiator port with OPEN_REJECT (RETRY); and
- b) if affiliations are supported, reject incoming connection requests from other STP initiator ports with OPEN_REJECT (STP RESOURCES BUSY).

While a wide STP initiator port is waiting for a response to a connection request to an STP target port, it shall not reject an incoming connection request from that STP target port because of its outgoing connection request. It may reject incoming connection requests for other reasons (see 7.2.5.12).

If a wide STP initiator port receives an incoming connection request from an STP target port while it has a connection established with that STP target port, it shall reject the request with OPEN_REJECT (RETRY).

The first dword that an STP phy sends inside an STP connection after OPEN_ACCEPT that is not an ALIGN or NOTIFY shall be an STP primitive (e.g., SATA_SYNC).

7.18 SMP link layer

7.18.3 Opening an SMP connection

An SMP target port shall not attempt to establish an SMP connection.

8.2.2.2.1 PL_OC1:Idle state description

...

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.

8.2.2.3.4 PL_OC2:Overall_Control state unable to establish a connection

...

If this state receives an Unable To Connect (Connection Rate Not Supported), Unable To Connect (Protocol Not Supported), Unable To Connect (Zone Violation), Unable To Connect (Reserved Abandon 1), Unable To

Connect (Reserved Abandon 2), Unable To Connect (Reserved Abandon 3), 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.

10.2.7.2.2 Protocol-Specific Port mode page - short format

...

The I_T NEXUS LOSS TIME field contains the time that the SSP target port shall retry connection requests to an SSP initiator port that are rejected with [certain](#) responses indicating [that](#) the SSP initiator port may no longer be present (see 8.2.2) before recognizing an I_T nexus loss (see 4.5). Table 179 defines the values of the I_T NEXUS LOSS TIME field. If this mode page is not implemented, the I_T nexus loss time is vendor specific. This value is enforced by the port layer (see 8.2.2).