

To: T10 Technical Committee
From: Rob Elliott, HP (elliott@hp.com)
Date: 26 October 2006
Subject: 06-466r0 SAS-2 OPEN_REJECT RETRY during self-configuration changes

Revision history

Revision 0 (26 October 2006) First revision

Related documents

sas2r06 - Serial Attached SCSI - 2 (SAS-2) revision 6
06-384r1 - OPEN_REJECT (RETRY) during zoning changes (Rob Elliott, HP)

Overview

After an initiator performs the discover process and discovers a target, it may start opening connections to that target. If a self-configuring expander along the pathway has not yet put the target's SAS address into its downstream routing tables, the initiator will get OPEN_REJECT (NO DESTINATION). This is not a problem; the initiator also discovers the expanders and should have noticed that one of them has its CONFIGURING bit set to one, indicating that it should extend the I_T nexus loss time indefinitely.

If the downstream route tables are all complete, the initiator could open the target successfully, send a read command (like INQUIRY), and close the connection. However, when the target tries to open the initiator, there is no guarantee that all the expanders have finished self-configuration in the upstream direction, so the target could get OPEN_REJECT (NO DESTINATION). Since targets don't perform the discover process, it doesn't know an expander still has its CONFIGURING bit set to one and that it should extend its I_T nexus loss time. If self-configuration takes too long, it will declare an I_T nexus loss (unless the initiator happens to open it in the meantime).

Even on the initiator side, reliance on the CONFIGURING bit is imperfect. The CONFIGURING bit could change in the time between the initiator discovering the expander and the initiator attempting to open the target. BROADCAST (CHANGE)s don't propagate immediately, and the discover process in the expander doesn't set the CONFIGURING bit to one until it detects that a change is needed.

Zoning expanders return OPEN_REJECT (RETRY) rather than OPEN_REJECT (ZONE VIOLATION) while zoning configuration changes are in progress. Self-configuring expanders could do the same, and return OPEN_REJECT (RETRY) rather than OPEN_REJECT (NO DESTINATION) if self-configuring changes are in progress.

Editor's Note 1: Should this take effect immediately as soon as BROADCAST (CHANGE) is received and the discover process starts, or continue to wait until the discover process determines that a change is needed?

Editor's Note 2: How about OPEN_REJECT (BAD DESTINATION)?

After expander Y and Z are attached, assume X and Y complete self-configuration quickly while Z is slow:

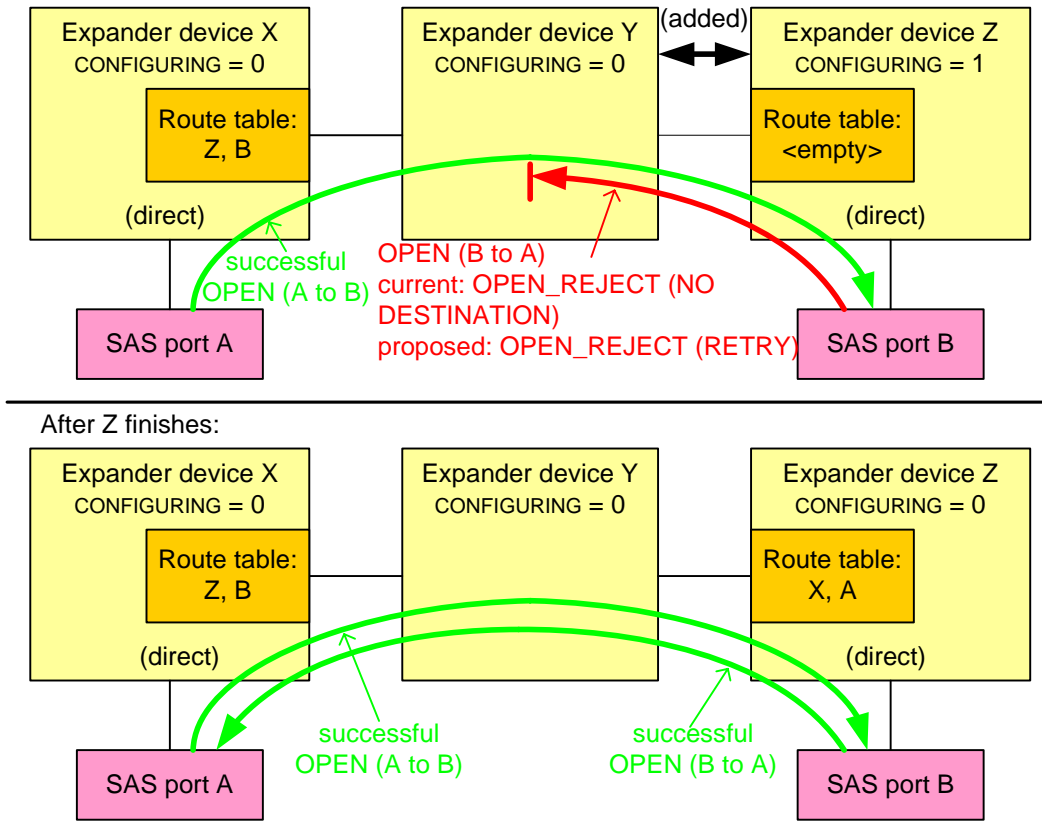


Figure 1 — Configuring not complete example

Suggested changes

4.6.6.3 ECM interface

...

Table 1 describes the confirmations from the ECM to an expander phy. These confirmations are sent in confirmation of a Request Path request.

Table 1 — ECM to expander phy confirmations (part 1 of 2)

Message	Description
Arbitrating (Normal)	Confirmation that the ECM has received the Request Path request.
Arbitrating (Waiting On Partial)	Confirmation that the ECM has determined that: <ul style="list-style-type: none"> a) there is a destination port capable of routing to the requested destination SAS address; b) at least one phy within the destination port supports the requested connection rate; c) each of the phys within the destination port is returning a Phy Status (Partial Pathway) or Phy Status (Blocked Partial Pathway) response; and d) at least one of the phys within the destination port is returning a Phy Status (Partial Pathway) response.
Arbitrating (Blocked On Partial)	Confirmation that the ECM has determined that: <ul style="list-style-type: none"> a) there is a destination port capable of routing to the requested destination SAS address; b) at least one phy within the destination port supports the requested connection rate; and c) each of the phys within the destination port is returning a Phy Status (Blocked Partial Pathway) response.
Arbitrating (Waiting On Connection)	Confirmation that the ECM has determined that the connection request is blocked due to one of the following reasons: <ul style="list-style-type: none"> a) the connection request is blocked by an active connection; or b) there are insufficient routing resources within the expander to complete the connection request. A connection request shall be considered blocked by an active connection when: <ul style="list-style-type: none"> a) there is a destination port capable of routing to the requested destination SAS address; b) at least one phy within the destination port supports the requested connection rate; c) each of the phys within the destination port is returning a Phy Status (Partial Pathway), Phy Status (Blocked Partial Pathway), or Phy Status (Connection) response; and d) at least one of the phys within the destination port is returning a Phy Status (Connection) response.
Arb Won	Confirmation that an expander phy has won path arbitration.
Arb Lost	Confirmation that an expander phy has lost path arbitration.

Table 1 — ECM to expander phy confirmations (part 2 of 2)

Message	Description
Arb Reject (No Destination)	Confirmation that the expander device is not configuring (see 4.7.3) and: <ol style="list-style-type: none"> there is no operational expander phy capable of routing to the requested destination SAS address; or the requested destination SAS address maps back to the requesting port (see 7.12.4.3).
Arb Reject (Bad Destination)	Confirmation that: <ol style="list-style-type: none"> the requested destination SAS address maps back to the requesting port; the requesting port is using the direct routing method or the table routing method; and the ECM has not chosen to return Arb Reject (No Destination) (see 7.12.4.3).
Arb Reject (Connection Rate Not Supported)	Confirmation that the ECM has determined that there is a destination port capable of routing to the requested destination SAS address but no phys within the destination port are configured to support the requested connection rate.
Arb Reject (Zone Violation)	Confirmation that the ECM has determined that there is a zoning violation (see 4.9.3).
Arb Reject (Pathway Blocked)	Confirmation that the ECM has determined that the requesting expander phy shall back off according to SAS pathway recovery rules.
Arb Reject (Retry)	Confirmation that: <ol style="list-style-type: none"> the ECM has determined that there is a zoning violation (see 4.9.3) while the zoning expander device is locked; or the expander device is configuring (see 4.7.3) and the ECM has determined that: <ol style="list-style-type: none"> there is no operational expander phy capable of routing to the requested destination SAS address; or the requested destination SAS address maps back to the requesting port (see 7.12.4.3).

...

4.7 Discover process

4.7.2 Discover process overview

...

~~If~~**While** a self-configuring expander device's CONFIGURING bit is set to one in the SMP REPORT GENERAL response, connection requests for destination ports two or more levels beyond the self-configuring expander device ~~may result in responses such as OPEN_REJECT (NO DESTINATION) (see 7.2.5.12). If a management application client (e.g., in an end device) detects a self-configuring expander device with the configuring bit set to one, the SAS port shall stop the I_T Nexus Loss timer for I_T nexuses involving such destination ports as long as the CONFIGURING bit is set to one (see 8.2.2.1) that would have returned OPEN_REJECT (NO DESTINATION) return OPEN_REJECT (RETRY) instead (see 4.7.3).~~

...

4.7.3 Discover process in a self-configuring expander device

The management application client of a self-configuring expander device shall configure:

- the expander routing table in that expander device; and

- b) the expander routing table in each externally configurable expander device in the SAS domain that is not located behind another self-configuring expander device.

When a self-configuring expander device receives a Broadcast (Change) the management application client shall start the discover process using the expander port that received the Broadcast (Change). If a change to the expander route table is identified then the management device server shall set its CONFIGURING bit to one in the SMP REPORT GENERAL response (see 10.4.3.3).

The management application client shall set the CONFIGURING bit to zero when the discover process is complete. When the CONFIGURING bit changes from one to zero:

- a) a zoning expander device with zoning enabled shall originate a Broadcast (Change) on each expander port that has access to the expander port through which the discover process was performed based on the zone permission table; and
- b) an expander device with zoning disabled shall originate a Broadcast (Change) on each expander port other than the one through which the discover process was performed.

When a Broadcast (Change) has been received by a self-configuring expander device, all previously valid SAS addresses shall continue to be routable until they are determined to be no longer valid. When a change has been identified, all unaffected SAS addresses shall continue to be routable.

[While the CONFIGURING bit is set to one, the expander device shall return OPEN_REJECT \(RETRY\) for any connection requests that would have resulted in OPEN_REJECT \(NO DESTINATION\).](#)

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 — OPEN_REJECT abandon primitives

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

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

Table 3 — OPEN_REJECT retry primitives

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

Editor's Note 3: "would" is not great standardese, but copying the three lengthy NO DESTINATION

[reasons is not a good alternative](#)
