

To: T10 Technical Committee
 From: Rob Elliott, HP (elliott@hp.com)
 Date: 25 October 2002
 Subject: 02-443r0 SAS Handling Link Rate Not Supported

Revision History

Revision 0 (25 October 2002) First revision

Related Documents

sas-r02b - Serial Attached SCSI revision 2b

Overview

Pak-lung Seto of Intel raised two questions on the T10 reflector about issues not addressed in sas-r02b.

1. What CONNECTION RATE (aka CONNECTION LINK RATE) does the target use in its OPENS?
2. What does a target do when it receives OPEN_REJECT (LINK RATE NOT SUPPORTED)?

Suggested Changes

1. What CONNECTION RATE should the target use in its OPENS?

Some options (I recommend option a))(for 7.7.3 OPEN address frame or 7.12.2.1 Opening a connection):

- a) When requesting a connection to an initiator port, a target port shall set the CONNECTION RATE field to the value received in the last connection request from the initiator port unless it has received an OPEN_REJECT (LINK RATE NOT SUPPORTED) (see 7.12.2.2 for handling of OPEN_REJECT (LINK RATE NOT SUPPORTED)).
 -) Discussion: requires the value be stored only in an I_T context, not with I_T_L_Q context.
- b) When requesting a connection to an initiator port, a target port shall set the CONNECTION RATE field to the value received in connection requests from the initiator port unless it has received an OPEN_REJECT (LINK RATE NOT SUPPORTED)(see 7.12.2.2 for handling of OPEN_REJECT (LINK RATE NOT SUPPORTED)). An initiator port shall use the same CONNECTION RATE field value for all connection requests to the same target port, and shall only change the CONNECTION RATE field value when it has no commands outstanding to that target port. Targets are not required to check consistency of the CONNECTION RATE field in different connection requests from the same initiator port.
 -) Discussion: this is like the INITIATOR CONNECTION TAG; the target can remember it any way it wants, in an I_T or I_T_L_Q context. Problem: If the initiator gets an OPEN_REJECT (LINK RATE NOT SUPPORTED) and has to change its rate while commands are outstanding, it will violate the rule above and confuse the target. The INITIATOR CONNECTION TAG doesn't have this problem, since there is no reason to change it.
- c) When requesting a connection to an initiator port, a 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 (LINK RATE NOT SUPPORTED)(see 7.12.2.2 for handling of OPEN_REJECT (LINK RATE NOT SUPPORTED)).
 -) Discussion: requires the value be stored in an I_T_L_Q context. If the target has commands marked for 1.5 Gbps and 3.0 Gbps, it would have to close a 3 Gbps connection to handle transfers for the 1.5 Gbps commands.

2. What does a device do when it receives OPEN_REJECT (LINK RATE NOT SUPPORTED)? This should not happen in normal operation, but could occur if there is a wide link with multiple link rates (already an unusual case) and all the fastest physical links go offline.

I suggest this rule for targets (for 7.12.2.2 Connection request responses):

If a target port receives an OPEN_REJECT (LINK RATE NOT SUPPORTED) in response to one of its connection requests that has the CONNECTION RATE field set to greater than 1.5 Gbps, it shall retry the connection request with the CONNECTION RATE field set to 1.5 Gbps.

Discussion: The target already handles OPEN REJECT (RETRY). This is just a variation on that - retrying and changing the rate to 1.5 Gbps. This way the target only downshifts when needed, but it does keep going. This follows the parallel SCSI model of "fallback to asynchronous" and continue to operate (albiet at a slowe rate) rather than giving up completely.

For initiators, an unexpected OPEN_REJECT(LINK RATE NOT SUPPORTED) should not happen, since a CHANGE should arrive before that response arrives. However, the initiator may still see OPEN_REJECT(LINK RATE NOT SUPPORTED) for a connection request in flight when CHANGE arrives. After new discovery, the initiator should repeat its connection request with a proper connection rate. This adversely affects option b) in issue #1.

For 7.7.3 OPEN address frame:

When requesting a connection to a target port, an initiator port shall set the CONNECTION RATE field to a value supported by the pathway (i.e., if there is a wide link in the pathway, at least one physical link in the wide link in the pathway supports the rate; if there is a narrow link in the pathway, that physical link supports the rate).