

TO: T10 Membership
FROM: Paul A. Suhler, Certance
DATE: 18 November 2003
SUBJECT: T10/03-319r3, ADT Transfer Ready IU / Response IU Race Condition

0. Revisions

Rev. 3:

Incorporated changes from 17 November working group teleconference (T10/04-002r0). Also made the following changes:

- In Negotiation and Fast Access exchanges, removed the requirement to send or receive an ACK IU for the exchange to end.

Rev. 2:

Incorporated changes from 3 – 4 November 2003 teleconference. See minutes, T10/03-378r0, discussion item e. Exceptions are:

- Several of us have concluded that SCSI task lifetimes (as per SAM-2/3) do not exactly match ADT SCSI exchange lifetimes. Thus, I left the SCSI exchange lifetime definition in terms of IUs, although transmission or reception of an ACK for the final IU is no longer required. See below.

Rev. 1:

Incorporated changes from 20 October 2003 teleconference. See minutes, T10/03-341r0, discussion item e. Differences from changes suggested in the minutes are:

- Rather than defining I_T Nexus Loss, specified the occurrence of Port Login or Port Logout.
- Defined negotiation exchange lifetimes in terms of transitions among negotiation states, which are already defined in terms of IU transmission and reception.

Rev. 0:

Initial proposal

1. Summary

We have identified a possible race condition in a data-in command. The problem can arise when the initiator port requests more data in a command than the target port wishes to send. Furthermore, the last data-in IU sent by the target port satisfies the amount of data requested by the most recent Transfer Ready IU sent by the initiator port:

Initiator Port	Target Port
SCSI Command IU ----->	
buffer allocation length = 9k	
first burst length = 4 k	
	(target has only 8k to send)
	<----- SCSI Data IU (2k)
	<----- SCSI Data IU (2k)
Xfer Rdy IU (4k)----->	
	<----- SCSI Data IU (2k)
	<----- SCSI Data IU (2k)
Xfer Rdy IU (1k) ----> RACE <-----	SCSI Response IU (GOOD Status)

As shown, the initiator port does not know that the target port will send more data, so it sends a Transfer Ready IU. The target port, having no more data-in to send, chooses to send a Response IU. The race can actually occur anywhere on the physical medium or in the protocol stacks of the initiator or the target.

This problem can be solved by specifying a solution to a more general problem, i.e., what to do with IUs that aren't a part of an existing exchange (and which don't start a new exchange).

2. Proposed Solution

A port receiving any of the following IUs outside of an open exchange will transmit an ACK IU and discard the IU:

- Port Login with the ACCEPT bit set to one
- SCSI Response
- SCSI Transfer Ready
- SCSI Data
- VHF Data
- AER Control (DTD port only)

Definitions of exchange lifetimes are given below.

3. Proposed Changes

If the editor wishes to move some or all of the following text into other subclauses than those suggested below, I would regard that as an editorial change. Changes since r2 are in [blue](#).

3.1 Add the following definition:

3.1.x simple exchange: An exchange consisting of two information units, an IU with a type other than Acknowledgement and the corresponding Acknowledgement IU.

3.2 Make the following changes to Table 15, NAK IU status code values:

This assumes that the previous codes [05h](#) – 08h were removed as per the 3 – 4 November meeting.

Status Code	Description
05h	INVALID EXCHANGE ID
07h – 2Fh	Reserved

3.3 Insert the following new subclause after subclause 6.5.8:

6.5.9 Link service exchange lifetime

6.5.9.1 Link service exchange types

Link service exchanges may be negotiation exchanges, port logout exchanges, pause exchanges, or NOP exchanges.

6.5.9.2 Simple link service exchange lifetime

Port logout IUs, Pause IUs, and NOP IUs are sent in simple exchanges. A simple exchange begins in the sending port with the transmission of the IU and ends with the reception of the

corresponding Acknowledgement IU. A simple exchange begins in the receiving port with the reception of a valid IU and ends with the transmission of the Acknowledgement IU.

6.5.9.3 Negotiation exchange lifetime

In a port initiating a negotiation exchange, the exchange begins when the port transmits a Port Login IU with the ACCEPT bit set to zero in a nonexistent exchange. In a port not initiating a negotiation exchange, the exchange begins when the port receives a Port Login IU with the ACCEPT bit set to zero in a nonexistent exchange. A negotiation exchange ends in a port when either:

- a) The port has sent a Login IU with the ACCEPT bit set to one and receives a Login IU with the ACCEPT bit set to one; or
- b) The port has received a Login IU with the ACCEPT bit set to one and sends a Login IU with the ACCEPT bit set to one.
- c) The port has received a Login IU with a difference exchange ID, indicating that negotiation has restarted.

If a port receives a Port Login IU with the ACCEPT bit set to one in a nonexistent exchange, it shall transmit a NAK IU with a status of INVALID EXCHANGE ID and discard the Port Login IU.

3.4 Insert the following new subclause after subclause 7.1.5:

7.1.6 SCSI task exchange lifetime

A SCSI task exchange begins in an initiator port after the port transmits a SCSI Command IU or a SCSI Task Management IU. A SCSI task exchange begins in a target port after the port receives a SCSI Command IU or a SCSI Task Management IU.

A SCSI encapsulation exchange ends in an initiator port after:

- 1) the port receives a SCSI Response IU for that exchange;
- 2) the port transmits a SCSI Task Management IU containing a task management request aborting the exchange and receives a SCSI Response IU with a RESPONSE CODE of COMMAND OR TASK MANAGEMENT FUNCTION COMPLETE for the task management request; or
- 3) a nexus loss occurs.

A SCSI encapsulation exchange ends in a target port after:

- 1) the port transmits a SCSI Response IU for that exchange;
- 2) the port receives a SCSI Task Management IU containing a task management request aborting the exchange and transmits a SCSI Response IU with a RESPONSE CODE of COMMAND OR TASK MANAGEMENT FUNCTION COMPLETE for the task management request;
- 3) a hard reset occurs; or
- 4) a nexus loss occurs.

If a port receives a SCSI Response IU, SCSI Transfer Ready IU, or SCSI Data IU in a nonexistent exchange, it shall transmit a NAK IU with a status of INVALID EXCHANGE ID and discard the SCSI IU.

3.5 Insert the following new subclause after subclause 7.2.6:

7.2.7 Fast Access exchange lifetime

7.2.7.1 Fast Access exchange types

Fast Access exchanges may be either VHF Data exchanges, AER Control exchanges, or AER Data exchanges.

7.2.7.2 VHF Data exchange lifetime

A VHF data exchange begins in an [ADT port in an automation device](#) after the port transmits a Request for VHF Data IU. The exchange begins in an [ADT port in a DTD](#) after the port receives a Request for VHF Data IU.

A VHF data exchange ends in an [ADT port in an automation device](#) after the port receives a VHF Data IU for the exchange. The exchange ends in an [ADT port in a DTD](#) after the port transmits a VHF Data IU for the exchange.

If an [ADT port in an automation device](#) receives a VHF Data IU in a nonexistent exchange, it shall transmit a NAK IU with a status of INVALID EXCHANGE ID and discard the VHF Data IU. If an [ADT port in a DTD](#) receives a VHF Data IU, it shall transmit a NAK IU with a status of UNSUPPORTED FRAME TYPE FOR SELECTED PROTOCOL and discard the VHF Data IU.

7.2.7.3 AER Control exchange lifetime

An AER control exchange begins in an [ADT port in an automation device](#) after the port transmits an AER Control IU. The exchange begins in an [ADT port in a DTD](#) after the port receives an AER Control IU.

An AER control exchange ends in an [ADT port in an automation device](#) after the port receives an AER Control IU. The exchange ends in an [ADT port in a DTD](#) after the port transmits an AER Control IU.

If an [ADT port in an automation device](#) receives an AER Control IU in a nonexistent exchange, it shall transmit a NAK IU with a status of UNSUPPORTED FRAME TYPE FOR SELECTED PROTOCOL and discard the AER Control IU.

7.2.7.4 AER Data exchange lifetime

AER data exchanges are simple exchanges. The exchange begins in an [ADT port in a DTD](#) with the transmission of the IU and ends with the reception of the corresponding Acknowledgement IU. A simple exchange begins in an [ADT port in an automation device](#) with the reception of a valid IU and ends with the transmission of the corresponding Acknowledgement IU.

If an [ADT port in a DTD](#) receives an AER Data IU, it shall transmit a NAK IU with a status of UNSUPPORTED FRAME TYPE FOR SELECTED PROTOCOL and discard the AER Data IU.

3.6 Make the following global change:

Change “ADC Fast Access” to “Fast Access” with the words capitalized.