

TO: T10 Membership
FROM: Paul A. Suhler, Certance
DATE: 17 October 2003
SUBJECT: T10/03-319r0, ADT Transfer Ready IU / Response IU Race Condition

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 Request 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 IU with the ACCEPT bit set to one
- SCSI Response
- SCSI Transfer Ready
- SCSI Data
- VHF Data

A negotiation exchange is opened by transmission of an ACK for a Port Login IU with ACCEPT set to zero and closed by transmission of an ACK for a Port Login IU with ACCEPT set to one.

A SCSI exchange is opened by transmission or reception of an ACK for a SCSI Request IU and closed by transmission or reception of an ACK for a SCSI Response IU.

A VHF Data exchange is opened by transmission or reception of an ACK for a Request for VHF Data IU and closed by transmission or reception of an ACK for a VHF Data IU.

More precise wording for the above is given below.

3. Proposed Changes

Insert the following new subclause after subclause 6.5.8:

6.5.9 Link service exchange lifetime

Link service exchanges may be negotiation exchanges, port logout exchanges, pause exchanges, or NOP exchanges. Port logout IUs, Pause IUs, and NOP IUs are sent in single-IU exchanges and the lifetime of the exchange is only for that single IU.

A negotiation exchanges becomes active in a port when the port transmits a Port Login IU with the ACCEPT bit set to zero and the port receives an ACK IU for it. A negotiation exchange becomes inactive in a port after (i) the port receives a Port Login IU with the ACCEPT bit set to one and transmits an ACK IU for it; or (ii) the port transmits a Port Login IU with the ACCEPT bit set to one and receives an ACK IU for it.

If a port receives a Port Login IU with the ACCEPT bit set to one in an inactive exchange, it shall transmit an ACK IU and discard the Port Login IU.

Insert the following new subclause after subclause 7.1.5:

7.1.6 SCSI exchange lifetime

A SCSI encapsulation exchange becomes active in an initiator port after the port transmits a SCSI Request IU and receives an ACK IU for it. A SCSI encapsulation exchange becomes active in a target port after the port receives a valid SCSI Request IU and transmits an ACK IU for it.

A SCSI encapsulation exchange becomes inactive in an initiator port after (i) the port receives a SCSI Response IU for that exchange and transmits an ACK IU for it; (ii) the port transmits a SCSI Request IU containing a ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET, or TARGET RESET task management request affecting the exchange, receives a SCSI Response IU with a RESPONSE CODE of for the task management request, and transmits an ACK IU for it; or (iii) an I_T nexus loss occurs.

A SCSI encapsulation exchange becomes inactive in a target port after (i) the port transmits a SCSI Response IU for that exchange and receives an ACK IU for it; (ii) the port receives a SCSI Request IU containing a ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET, or TARGET RESET task management request affecting the exchange, transmits a SCSI Response IU with a RESPONSE CODE of for the task management request, and receives an ACK IU for it; (iii) a hard reset occurs; or (iv) an I_T nexus loss occurs.

If a port receives a SCSI Response IU, SCSI Transfer Ready IU, or SCSI Data IU in an inactive exchange, it shall transmit an ACK IU and discard the SCSI IU.

Insert the following new subclause after subclause 7.2.6:

7.2.7 ADC fast access exchange lifetime

ADC fast access exchanges may be either VHF data exchanges or AER control exchanges. AER IUs are sent in single-IU exchanges and the lifetime of the exchange is only for that single IU.

An ADC VHF data exchange becomes active in an initiator port after the port transmits a Request for VHF Data IU and receives an ACK IU for it. An ADC fast access exchange becomes active in a target port after the port receives a Request for VHF Data IU and transmits an ACK IU for it.

An ADC VHF data exchange becomes inactive in an initiator port after the port receives a VHF Data IU for the exchange and transmits an ACK IU for it. An ADC fast access exchange becomes inactive in a target port after the port transmits a VHF Data IU for the exchange and receives an ACK IU for it. An ADC fast access exchange becomes inactive in either an initiator or target port after (i) the port receives a Port Login IU and transmits an ACK IU for it; or (ii) the port receives a Port Logout IU and transmits an ACK IU for it.

If a port receives an ADC VHF Data IU in an inactive exchange, it shall transmit an ACK IU for it and discard the ADC VHF Data IU.

An ADC AER control exchange becomes active in an automation port after the port transmits an AER Control IU and receives an ACK IU for it. An ADC AER control exchange becomes active in a DTD port after the port receives an AER Control IU and transmits an ACK IU for it.

An ADC AER control exchange becomes inactive in an automation port after the port receives an AER Control IU and transmits an ACK IU for it. An ADC AER control exchange becomes inactive in a target port after the port transmits an AER Control IU and receives an ACK IU for it.

If an automation port receives an AER Control IU in an inactive exchange, it shall transmit an ACK IU for it and discard the AER Control IU.