Document:	X3T10/95-208r1	Date:	April 18, 1998
To:	X3T10 Committee Membership		
From:	Edward A. Gardner		
Subject:	Protocol Specific use of Disconnect-Reconnect Mode Page		

At the Harrisburg X3T10 working group meeting I asked how we should deal with protocol specific use of mode page parameters, specifically those in the Disconnect-reconnect page. The concensus at that meeting was that SPC should describe the mode page per se and the individual protocol documents should describe how the parameters apply to the specific protocol. The following is my proposal for how the current Disconnect-reconnect mode page description should be divided between the SPC and SIP documents. I've also incorporated a few new parameters, such as Giles initial burst size proposal. A separate document will propose additions to FCP-2 to describe how these parameters apply to that protocol.

1. Proposed Text for SPC (with revision marks)

Replace section 8.3.2 with the following. All changes from SPC revision 7 are shown with change bars and cross-outs for deleted text or underlines for new text. For clarity, this text with most revision marks removed is repeated in the next section of this document.

8.3.2 Disconnect-reconnect page

The disconnect-reconnect page (see table 83) provides the application client the means to tune the performance of the service delivery subsystem. The name for this mode page (disconnect-reconnect) comes from the SCSI-2 parallel bus. A SCSI-3 device based on any of the protocols (SIP, SBP, FCP, SSP, etc.) may use appropriate parameters in the disconnect-reconnect mode page. The parameters appropriate to each protocol and their interpretation for that protocol are specified in the individual protocol documents.

If a parameter that is not appropriate to the specific protocol implemented by the SCSI-3 device is non-zero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

Table 83 - Disconnect-reconnect page

Byte \ Bit	7	6	5	4	3	2	1	0
0	PS	Reserved		Page code (02h)				
1				Page length (0Eh)				
2				Buffer full ratio				
3				Buffer empty ratio				
4	(MSB)			Bus inactivity limit				
5								(LSB)
6	(MSB)			Disconnect time limit				
7								(LSB)
8	(MSB)			Connect time limit				
9								(LSB)
10	(MSB)			Maximum burst size				
11								(LSB)
12	EMDP	PARd	PAWrt	PAStat	DImm		DTDC	
13				Reserved				
14	(MSB)			Initial burst size				
15								(LSB)

The buffer full ratio field indicates to the device server, on read operations, how full the buffer should be prior to attempting to move data on the interconnect. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in clause 5.2.

The buffer empty ratio field indicates to the device server, on write operations, how empty the buffer should be prior to attempting to move data on the interconnect. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in clause 5.2.

The bus inactivity limit field indicates the maximum time that the target is permitted to lockout other uses of the interconnect without actually moving data. If the bus inactivity limit would be exceeded the device server shall attempt to interrupt the data moving operation, within the restrictions placed on it by the application client and the applicable SCSI-3 protocol. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no bus inactivity limit.

The disconnect time limit field indicates the minimum time that the target shall wait between attempts to move data on the interconnect. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no disconnect time limit.

The connect time limit field indicates the maximum time that the target is allowed to use the interconnect during a single data moving operation. If the connect time limit is exceeded the device server shall attempt to interrupt the data moving operation, within the restrictions placed on it by the application client and the applicable SCSI-3 protocol. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no connect time limit.

The maximum burst size field indicates the maximum amount of data that the device server shall transfer during a single data moving operation. This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per data moving operation.

The enable modify data pointers (EMDP) bit indicates whether or not the initiator allows the data transfer to be re-ordered by the target. If the EMDP bit is zero, the target shall not re-order the data transfer. If the EMDP bit is one, the target is allowed to re-order the data transfer.

The PARd, PAWrt, and PAStat bits indicate whether the target may use a higher priority arbitration method when performing read data transfers, write data transfers, and status or control message transfer respectively. A PA bit of zero indicates that the target shall use the normal or default arbitration method for the protocol. A PA bit of one indicates that the target may use an optional higher priority arbitration method.

A disconnect immediate (DImm) bit of zero indicates that the target may begin execution of a command (e.g. perform a data transfer) during the same use of the interconnect in which it receives the command. Whether or not the target does so may be based on the target's internal algorithms, the rules of the applicable SCSI-3 protocol, and settings of the other parameters in this mode page. A disconnect immediate (DImm) bit of one indicates that the target shall not begin execution of a command during the same use of the interconnect that it receives the command.

The data transfer disconnect control (DTDC) field (see table 84) defines further restrictions on when a disconnect is permitted.

Table 84 - Data transfer disconnect control

DTDC	Description
000b	Data transfer disconnect control is not used. Disconnect is controlled by the other fields in this page.
001b	A target shall transfer all data for a command within a single use of the interconnect.
001b	A target shall transfer all data for a command and complete the command within a single use of the interconnect.
100b-111b	Reserved

The initial burst size field indicates the maximum amount of data that the initiator may transfer in conjunction with a command. This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates that initiators shall not transfer data in conjunction with commands.

2. Proposed Text for SPC (without revision marks)

The following is the same text shown in the previous section, but with most revision marks removed for clarity. Intentional functional changes from the current specification are flagged with revision marks, all other revision marks have been removed.

8.3.2 Disconnect-reconnect page

The disconnect-reconnect page (see table 83) provides the application client the means to tune the performance of the service delivery subsystem. The name for this mode page (disconnect-reconnect) comes from the SCSI-2 parallel bus. A SCSI-3 device based on any of the protocols (SIP, SBP, FCP, SSP, etc.) may use appropriate parameters in the disconnect-reconnect mode page. The parameters appropriate to each protocol and their interpretation for that protocol are specified in the individual protocol documents.

If a parameter that is not appropriate to the specific protocol implemented by the SCSI-3 device is non-zero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

Table 83 - Disconnect-reconnect page

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Reserved		Page code (02h)				
1				Page length (0Eh)				
2				Buffer full ratio				
3				Buffer empty ratio				
4	(MSB)							
5				Bus inactivity limit				(LSB)
6	(MSB)							
7				Disconnect time limit				(LSB)
8	(MSB)							
9				Connect time limit				(LSB)
10	(MSB)							
11				Maximum burst size				(LSB)
12	EMDP	PARd	PAWrt	PAStat	DImm		DTDC	
13				Reserved				
14	(MSB)							
15				Initial burst size				(LSB)

The buffer full ratio field indicates to the device server, on read operations, how full the buffer should be prior to attempting to move data on the interconnect. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in clause 5.2.

The buffer empty ratio field indicates to the device server, on write operations, how empty the buffer should be prior to attempting to move data on the interconnect. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in clause 5.2.

The bus inactivity limit field indicates the maximum time that the target is permitted to lockout other uses of the interconnect without actually moving data. If the bus inactivity limit would be exceeded

the device server shall attempt to interrupt the data moving operation, within the restrictions placed on it by the application client and the applicable SCSI-3 protocol. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no bus inactivity limit.

The disconnect time limit field indicates the minimum time that the target shall wait between attempts to move data on the interconnect. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no disconnect time limit.

The connect time limit field indicates the maximum time that the target is allowed to use the interconnect during a single data moving operation. If the connect time limit is exceeded the device server shall attempt to interrupt the data moving operation, within the restrictions placed on it by the application client and the applicable SCSI-3 protocol. This value may be rounded as defined in clause 5.2. A value of zero indicates that there is no connect time limit.

The burst size field indicates the amount of data that the device server shall transfer during a single data moving operation. This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per data moving operation.

The enable modify data pointers (EMDP) bit indicates whether or not the initiator allows the data transfer to be re-ordered by the target. If the EMDP bit is zero, the target shall not re-order the data transfer. If the EMDP bit is one, the target is allowed to re-order the data transfer.

The PARd, PAWrt, and PAStat bits indicate whether the target may use a higher priority arbitration method when performing read data transfers, write data transfers, and status or control message transfer respectively. A PA bit of zero indicates that the target shall use the normal or default arbitration method for the protocol. A PA bit of one indicates that the target may use an optional higher priority arbitration method.

A disconnect immediate (DImm) bit of zero indicates that the target may begin execution of a command (e.g. perform a data transfer) during the same use of the interconnect in which it receives the command. Whether or not the target does so may be based on the target's internal algorithms, the rules of the applicable SCSI-3 protocol, and settings of the other parameters in this mode page. A disconnect immediate (DImm) bit of one indicates that the target shall not begin execution of a command during the same use of the interconnect that it receives the command.

The data transfer disconnect control (DTDC) field (see table 84) defines further restrictions on when a disconnect is permitted.

Table 84 - Data transfer disconnect control

DTDC	Description
000b	Data transfer disconnect control is not used. Disconnect is controlled by the other fields in this page.
001b	A target shall transfer all data for a command within a single use of the interconnect.
010b	Reserved
001b	A target shall transfer all data for a command and complete the command within a single use of the interconnect.
100b-111b	Reserved

The initial burst size field indicates the maximum amount of data that the initiator may transfer in conjunction with a command. This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates that initiators shall not transfer data in conjunction with commands.

3. Proposed Text for SIP

Add the following text to SIP as a new section.

n.n.n Use of Disconnect-reconnect page parameters

The disconnect-reconnect page (defined in SPC) provides the application client the means to tune the performance of the service delivery subsystem. The following describes the applicability and interpretation of the parameters in that mode page to the SIP protocol.

If any parameter other than those discussed in this section is non-zero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

The buffer full and buffer empty ratios are numerators of a fractional multiplier that has 256 as its denominator. A value of zero indicates that the target determines when to initiate reselection consistent with the disconnect time limit parameter. These parameters are advisory to the target.

NOTE nn As an example, consider a device server with ten 512-byte buffers and a specified buffer full ratio of 3Fh. The formula is: INTEGER((ratio/256)*number of buffers). Thus INTEGER((3Fh/256)*10) = 2. On read operations, the device server should attempt to reselect the initiator whenever two or more buffers are full.

The bus inactivity limit field indicates the maximum time in 100 us increments that the target is permitted to assert the BSY signal without a REQ/ACK handshake. If the bus inactivity limit is exceeded the target shall attempt to disconnect if the initiator has granted the disconnect privilege (see n.n.n) and it is not restricted by DTDC. A value of zero indicates that there is no bus inactivity limit.

The disconnect time limit field indicates the minimum time in 100 us increments that the target shall wait after releasing the SCSI bus before attempting reselection. A value of zero indicates that there is no disconnect time limit.

The connect time limit field indicates the maximum time in 100 ms increments that the target is allowed to use the SCSI bus before disconnecting, if the initiator has granted the disconnect privilege (see n.n.n) and it is not restricted by DTDC. A value of zero indicates that there is no connect time limit.

The maximum burst size field indicates the maximum amount of data that the target shall transfer during a data phase before disconnecting if the initiator has granted the disconnect privilege. This value is expressed in increments of 512 bytes (e.g. a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per connection.

The enable modify data pointers (EMDP) bit indicates whether or not the initiator allows the Modify Data Pointers message to be sent by the target. If the EMDP bit is a zero, the target shall not issue the Modify Data Pointers message. If the EMDP bit is a one, the target is allowed to issue Modify Data Pointers messages.

If the EMDP bit is a one and the initiator responds to a Modify Data Pointer message with a Message Reject, then the target shall go to the STATUS phase and send the initiator a CHECK CONDITION. The sense key shall be set to ABORTED COMMAND and the sense code shall be set to INVALID MESSAGE ERROR.

A disconnect immediate (DImm) bit of zero indicates that the target may disconnect after command phase if it chooses to do so based on its internal algorithms, the setting of the DiscPriv bit in the IDENTIFY message and the settings of the other parameters in this mode page.

A disconnect immediate (DImm) bit of one indicates that the target shall attempt to disconnect immediately after every command phase for those connections in which disconnections are allowed.

The data transfer disconnect control (DTDC) field (see table xx) defines further restrictions on when a disconnect is permitted.

Table xx - Data transfer disconnect control

DTDC	Description
000b	Data transfer disconnect control is not used. Disconnect is controlled by the other fields in this page.
001b	A target shall not attempt to disconnect once the data transfer of a command has started until all data the command is to transfer has been transferred. The connect time limit and bus inactivity limit are ignored during the data transfer.
010b	Reserved
011b	A target shall not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once data transfer has started.
100b-111b	Reserved

If DTDC is non-zero and the maximum burst size is non-zero, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.