

Document: X3T10/95-208r3 Date: September 14, 1995
 To: X3T10 Committee Membership
 From: Edward A. Gardner
 Subject: Protocol Specific use of Disconnect-Reconnect Mode Page

1. Proposed Text for SPC

Replace section 8.3.2 with the following.

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 may be specified in the individual protocol documents.

The device server communicates the parameter values in this mode page to the service delivery subsystem (e.g., to its Target Role Agent). Similarly the application client may also communicate parameter values to the service delivery subsystem (e.g., those controlling behavior of its Initiator Role Agent). This communication is internal to the initiator or target device and is outside the scope of SCSI-3.

If a parameter that is not appropriate for 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	7	6	5	4	3	2	1	0
Byte								
0	PS	Reserved		Page code (02h)				
1				Page length (0Eh)				
2				Buffer full ratio				
3				Buffer empty ratio				
4	(MSB)			Bus inactivity limit				(LSB)
5								
6	(MSB)			Disconnect time limit				(LSB)
7								
8	(MSB)			Connect time limit				(LSB)
9								
10	(MSB)			Maximum burst size				(LSB)
11								
12	EMDP	FARd	FAWrt	FAStat	DIImm		DTDC	
13				Reserved				
14	(MSB)			First burst size				(LSB)
15								

An interconnect tenancy is a period of time during which a target device owns or may access the interconnect. For example, on arbitrated interconnects, a tenancy typically begins when a device successfully arbitrates for the interconnect and ends when the device releases the interconnect for use by other devices. Data and other information transfers take place during interconnect tenancies. The precise definition of an interconnect tenancy is specified in the individual protocol documents.

The buffer full ratio field indicates to the device server, on read operations, how full the buffer should be prior to requesting an interconnect tenancy. 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 requesting an interconnect tenancy. 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 full and buffer empty ratios are numerators of a fractional multiplier that has 256 as its denominator. A value of zero indicates that the device server determines when to request an interconnect tenancy consistent with the disconnect time limit parameter. These parameters are advisory to the target.

NOTE 55 As an example, consider a device server with ten 512-byte buffers and a specified buffer full ratio of 3Fh. The formula is: $\text{INTEGER}((\text{ratio}/256)*\text{number of buffers})$. Thus $\text{INTEGER}((3\text{Fh}/256)*10) = 2$. On read operations, the device server should request an interconnect tenancy whenever two or more buffers are full.

The bus inactivity limit field indicates the maximum time that the target is permitted to maintain an interconnect tenancy without data or information transfer. If the bus inactivity limit would be exceeded the device server shall conclude the interconnect tenancy, within the restrictions placed on it by 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 interconnect tenancies. 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 duration of a single interconnect tenancy. If the connect time limit would be exceeded the device server shall conclude the interconnect tenancy, within the restrictions placed on it by 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 maximum amount of data that the device server may 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.). The relationship (if any) between data moving operations and interconnect tenancies is specified in the individual protocol documents. 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 FARd, FAWrt, and FAStat bits indicate whether the target should use fair or unfair (e.g. priority) arbitration when requesting an interconnect tenancy for a read data transfer, write data transfer, and status or control message transfer respectively. An FA bit of one indicates that the target should use fair arbitration. An FA bit of zero indicates that the target should use unfair (e.g. priority) arbitration.

A disconnect immediate (DImm) bit of zero indicates that the target may transfer data for a command during the same interconnect tenancy in which it receives the command. Whether or not the target does so may depend upon 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 transfer data for a command during the same interconnect tenancy in which it receives the command.

The data transfer disconnect control (DTDC) field (see table 84) defines further restrictions on when multiple interconnect tenancies are permitted.

Table 84 - Data transfer disconnect control

DTDC	Description
000b	Data transfer disconnect control is not used. Interconnect tenancies are controlled by the other fields in this page.
001b	A target shall transfer all data for a command within a single interconnect tenancy.
010b	Reserved
011b	A target shall transfer all data for a command and complete the command within a single interconnect tenancy.
100b- 111b	Reserved

The first burst size field indicates the maximum amount of data that a target may transfer for a command during the same interconnect tenancy in which it receives the 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 there is no first burst size limit.

2. 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 parameters in that mode page that are appropriate for the SIP protocol and their interpretation.

The application client passes parameter values controlling the service delivery subsystem to the target device server by means of this mode page. The target device server in turn communicates the parameter values to its Target Role Agent to control use of the SCSI-3 Parallel Interface.

Parameter values are communicated from the target device server to its Target Role Agent using an internal interface (within the target device) that is outside the scope of SCSI-3.

Only the parameters discussed below are appropriate to SIP. If any parameter other than those discussed below 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.

A SIP interconnect tenancy begins with a successful selection or reselection and ends with the subsequent Bus Free. A SIP Target Role Agent requests an interconnect tenancy by issuing a Selection Service request. It concludes a SIP interconnect tenancy by issuing a Bus Free Service request.

The buffer full and buffer empty ratios are interpreted as described in SPC.

The bus inactivity limit, disconnect time limit and connect time limit fields are in 100 us increments. A value of zero indicates that the respective time limit does not apply. The bus inactivity limit field indicates the maximum time that the target device server may maintain a SIP interconnect tenancy without an active Command, Data Out, Data In, Status, Message Out, or Message In service request. A service request is active from the time the request is issued until the corresponding confirmation is returned. The disconnect and connect time limits are described in SPC.

The maximum burst size field indicates the maximum amount of data that the device server may transfer during a single SIP interconnect tenancy. A single SIP data moving operation corresponds to the data transferred during a single SIP interconnect tenancy. The maximum burst size only applies if the initiator has granted the disconnect privilege (see n.n.n). The maximum burst size value encoding is described in SPC.

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 Role Agent may request Data In or Data Out Services following a Command Service without an intervening Bus Free Service request (i.e., during the same SIP interconnect tenancy). A disconnect immediate (DImm) bit of one indicates that the Target Role Agent shall issue a Bus Free Service request between a Command

Service request and subsequent Data In or Data Out Service requests. The disconnect immediate (DImm) bit only applies when the initiator has granted the disconnect privilege (see n.n.n).

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	Once it has issued a Data In or Data Out service request, the Target Role Agent shall not issue a Bus Free Service request (shall not disconnect) until the last Data In or Data Out service request has been issued for the command (all the data transferred). The connect time limit and bus inactivity limit are ignored during the data transfer.
010b	Reserved
001b	Once it has issued a Data In or Data Out service request, the Target Role Agent shall not issue a Bus Free Service request (shall not disconnect) until the device server issues the Send SCSI command service response (the command is complete). The connect time limit and bus inactivity limit are ignored once data transfer has begun.
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.