

Document: X3T10/95-___r0

Date: April 18, 1998

To: X3T10 Committee Membership

From: Edward A. Gardner

Subject: FCP-2 use of Disconnect-Reconnect Mode Page

Document 95-208r3 was accepted at the last X3T10 meeting (Bedford, NH), splitting the description of the current disconnect-reconnect mode page between SPC and SIP to allow protocol specific use of that mode page. This document is a first pass at a proposal for FCP-2 use of that mode page. This is a work in progress intended to elicit discussion, not a completed proposal.

Remove any existing discussion of disconnect-reconnect mode page parameters from FCP, then add the following 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 FCP protocol and their interpretation.

[editorial note: FCP-2 is supposed to define SAM compliant service interfaces, with application clients and service delivery subsystems and the like. I am assuming that FCP-2 uses similar terminology to SIP, e.g. that "Target Role Agent" is a correct term. The terminology in this section may have to be updated if FCP-2 changes the meaning of or introduces new terms.]

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 Fibre Channel Protocol. 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 FCP. 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.

An FCP interconnect tenancy corresponds to an FCAL loop tenancy. A tenancy begins when a device opens the loop (issues an OPN primitive) and ends when the loop is closed (both devices have issued CLS primitives). An FCP Target Role Agent requests an interconnect tenancy by arbitrating for the loop. It concludes an FCP interconnect tenancy by issuing a CLS primitive. Interconnect tenancies are only meaningful for loop configurations.

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

The bus inactivity limit field is in transmission word increments. This field indicates the maximum time that the target may withhold sending CLS without sending information frames or permitting their reception. Specifically a target device shall send CLS if the following would persist for more than the bus inactivity limit number of transmission words:

1. The target is unable to send an information frame, whether due to it having no frames to send or it having zero available BB credit.
2. Either the initiator has zero available BB credit (the target has begun receiving the "last" frame) and the target is unable to send an R_RDY, or the initiator has sent a CLS or TTC primitive. *[note: The TTC primitive is the Transfer Tenancy Control primitive being*

considered for inclusion in FCAL-2. This sentence assumes that the TTC proposal is adopted by X3T11.]

The bus inactivity limit is ignored in non-loop configurations.

[note: Should a bus inactivity limit value of zero indicate a zero inactivity limit (i.e., send CLS immediately), or no limit (i.e. indefinitely large)?]

The disconnect time limit and connect time limit fields are in 128 transmission word increments (i.e. 512 bytes). A value of zero indicates that the respective time limit does not apply. The disconnect and connect time limits are described in SPC. The disconnect and connect time limits are ignored in non-loop configurations.

The maximum burst size field indicates the maximum amount of data that the device server may transfer using a single Fibre Channel sequence. This parameter does not affect how much data is transferred in an interconnect tenancy. The maximum burst size value encoding is described in SPC.

The enable modify data pointers (EMDP) bit indicates whether or not the target may re-order multiple Fibre Channel sequences. If the EMDP bit is zero, the target shall generate continuously increasing DATA_RO values within each FCP IO operation. If the EMDP bit is one, the target may generate consecutive DATA_RO values that are not continuously increasing within a single FCP IO operation. Note that this bit does not affect the order of frames within a sequence.

The FARd, FAWrt, and FASat bits indicate whether the target shall use the access fairness algorithm when arbitrating for the loop. An FA bit of one indicates that the target shall use the access fairness algorithm. An FA bit of zero indicates that the target may choose to not use the access fairness algorithm. The FARd bit controls arbitration when the target wishes to send one or more FCP_DATA frames to an initiator. The FAWrt bit controls arbitration when the target wishes to send one or more FCP_XFER_RDY frames to an initiator. The FASat bit controls arbitration when the target wishes to send one or more FCP_RSP frames to an initiator or FCP_CMND frames to another target. If the target intends to send multiple frame types, it may choose to not use the access fairness algorithm if any one applicable FA bit is zero. The FA bits are ignored in non-loop configurations.

A disconnect immediate (DImm) bit of zero indicates that the Target Role Agent may send FCP_DATA or FCP_XFER_RDY frames for a command during the same interconnect tenancy in which it receives the command. A disconnect immediate (DIMM) bit of one indicates that the Target Role Agent shall not send FCP_DATA or FCP_XFER_RDY frames for a command during the same interconnect tenancy in which it receives the command. The disconnect immediate (DIMM) bit is ignored in non-loop configurations.

The data transfer disconnect control (DTDC) field is not appropriate to FCP devices. If DTDC 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 first burst size field indicates the maximum amount of data that an initiator may send following a command, without waiting for an FCP_XFER_RDY. The target shall configure its buffers to accept up to this much data following any write-type command.