

Draft Minutes  
MC/DTE Interface Working Group  
T10/02-059r0  
January 15, 2002 – Houston, TX  
9:02 AM – 4:00 PM

The meeting was called to order by the facilitator, Paul Suhler, at 9:02.

### 1. Introductions

Paul Suhler led the group in introductions. The following people were in attendance:

Michael Banther	HP	michael_banther@hp.com
Joe Breher	Exabyte Corp	joebre@exabyte.com
Kevin Butt	IBM	kdbutt@us.ibm.com
Sergio Encarnacao	Overland Data	sencarnacao@overlanddata.com
Paul Entzel	Quantum	paul.entzel@quantum.com
Halvard Eriksen	Tandberg Data	halvard.eriksen@tandberg.com
Kevin Gingerich	Texas Instruments	k-gingerich@ti.com
Susan Gray	Quantum	susan.gray@quantum.com
Robert Griswold	Crossroads	rgriswold@crossroads.com
Steve Hellwege	Quantum	steve.hellwege@quantumatl.com
Lee Jesionowski	IBM	ljesion@us.ibm.com
Bjarte Myrold	Tandberg Data	bjarte.myrold@tandberg.com
Erich Oetting	StorageTek	erich_oetting@stortek.com
Paul Suhler	Seagate	paul.a.suhler@seagate.com
Andy Tran	Compaq	andy.tran4@compaq.com
Rod Wideman	ADIC	rod.wideman@adic.com

### 2. Approval of previous minutes

The minutes of the previous meeting held November 6, 2001 were presented (T10/01-349r0). These were approved as written

### 3. Approval of this agenda

The proposed agenda was approved, with addition of "where do we go from here" under 6.- Unscheduled Business.

### 4. Review of Action Items

Handled out of order; see below.

## 5. Patent Policy

Paul Entzel pointed out that the group should adopt a stated patent policy. This yielded an action item for Paul Suhler to find out whether T10 rules require individual patent applications to be disclosed.

Bob Griswold stated that T10 patent policy does not require disclosure of individual patents.

## 6. Scheduled Business

### 6.1 Call for secretary

Paul Suhler gave a call for volunteers for the role of secretary for the group. Joe Breher volunteered for this position. His appointment was approved without dissent.

### 6.2 Project Proposal Overview

Paul Suhler presented "Project Overview.pdf". The salient points are summarized as:

#### Project Proposals

Automation/Drive Interface – Commands (ADC) 02-008r0

Automation/Drive Interface – Transport Protocol (ADT) 02-009r0

Automation/Drive Interface – Physical Layer (ADP) 02-010r0

#### Proposed Milestones

1 List of proposed inclusions Mar 2002

2 First draft May 2002

3 Last new technical input Sep 2002

4 Document ready for first letter ballot Nov 2002

5 Forward to NCITS for first public review Mar 2003

Paul asked if the group felt that these milestones were achievable. There was no evidence presented to indicate that they were not.

### 6.3 Seagate Encapsulated SCSI Protocol

Paul Suhler presented 02-002r0.pdf - "Viper 200 LTO Tape Drive Library Interface Specification: Encapsulated SCSI Protocol".

Discussed Vendor Specific Inquiry Vital Product Data Pages – Drive Status Page. This page contains information on the loading state of the drive. Michael Banther initiated discussion of specificity of returned info for drive loading states. It was pointed out that we may or may not be able to specify a loading state machine model that will work for all drives. Lee Jesionowski indicated that he would like to separate the concept of "state" from the concept of "last action" in the drive loading state machine model.

Paul Suhler moved the discussion to 5.6, Vendor Specific Mode Pages, Table 12 - Interface Control Mode Page. This page contains serial port mode variables and addressing information for the drive's automation interface. Joe Breher opened question of whether Port \* Selection ID is sufficient for iSCSI. Bob Griswold pointed out that names in iSCSI are text[256]. David Hafner is reputed to be working on a mapping of iSCSI names to smaller identifiers.

Paul Suhler walked the group through Protocol Exchange types. In regards to the Data In Transaction, Paul mentioned that he somewhat regrets the requirement for a "transfer ready" indication from the destination device, as it is an unnecessary complication.

Michael Banther asked about SCSI bus performance within the scope of commands forwarded by the drive to the library. Paul acknowledged that this could be a problem, however, he was unaware of any field problems with this in Seagate's shipping product.

Joe Breher asked about reset implementation. Specifically, the drive can issue Reset Task Management Function to the library. Paul indicated that this ability can be disabled.

Application layer - Link layer protocol was discussed by Paul Suhler. Bus access is handled through collision detection. There is no arbitration for bus ownership.

Lee Jesionowski opened a discussion of timeouts. Specifically, should they be relaxed, or efficient. The Seagate specification is built around tight timeouts, IBM's is built around loose simple timeouts, assuming low error rate. Erich Oetting argued for loose timeouts. Michael Banther asked whether we are specifying command layer timeouts, or just link layer. Consensus was we need to specify both, but the current discussion should be constrained to link layer timeouts.

A discussion began around the topic of collision detection and multidrop issues. TCP/IP was brought up as a protocol having dealt with these issues already. Joe Breher asked whether many of these topics may have been already resolved by GPP or SBP. The discussion then vectored into other possibilities - I2C, USB, Ethernet, CAN, and the number of gates required to implement these protocols.

Paul Suhler indicated that Seagate's protocol originally considered supporting an address in the header to support multidrop, but no longer does.

Paul Suhler presented section 7.3, Command Payload. Paul indicated it borrows concepts from FCP. Lee Jesionowski asked the purpose of Data Length, as the CDB already contains the data. Joe Breher pointed out that this removes requirement that the link layer understand the various possible CDB formats.

Joe Breher asked to keep whatever protocol we forward as compliant as possible with the concepts embodied in SAM and SAM-2. Erich Oetting asked about the drive forwarding an FCP frame. Paul Entzel pointed out that what would be forwarded would be more in line with an FC sequence - not a frame.

Erich Oetting asked about baud rate changes. Specifically, when is it allowed change? Paul Suhler indicated that the communication employs automatic detection of baud rate. Once communications are established, a 'send mode select' operation is sent over the bus to perform the baud rate change.

Joe Breher posed a question to the group at large about the maximum length to allow for between the library controller and the drive. Erich Oetting indicated up to 20 feet, maybe longer.

Michael Banther indicated that some libraries multiplex multiple drives onto a single bus. This affects timeouts (among other things). It was deemed desirable that we not force automation vendors into not being able to use this architecture.

#### 6.4 Quantum Packetized Serial Protocol

Paul Entzel presented T10/02-036r0. This document indicates that Quantum has a patent on a portion of his proposal covering Baud rate negotiation. This document outlines Quantum's policy in regards to this patent, as applied to the ADI project.

Paul Entzel presented T10/02-013r0, "Packetized Serial Protocol Specification", as an implementation of a transport layer for RS-422 based communications. Paul indicated that his presentation of this document will focus on features of this protocol presentation that he has not seen in the other proposals on today's agenda.

Paul Entzel indicated that this proposal employs a 16-bit CRC, rather than a checksum. The stated reason for this was that checksum is not powerful enough for some physical layers.

Paul Entzel indicated that this proposal employs guaranteed unique characters for SOF (start of frame) and EOF (end of frame), which lends itself to hardware acceleration for generating an interrupt upon reception of frame - not upon reception of byte. Lee Jesionowski indicated that his proposal to be discussed later allows for this as well.

Paul Entzel pointed out that in this protocol, the frame header specifies which device is addressed by that frame.

Paul Entzel pointed out that this protocol readily allows for the transport of existing diagnostic interfaces, as well as other legacy protocols.

Paul Entzel indicated that Quantum's current implementation does not allow for queued commands, or processing of a new 'info command' while currently processing an 'action' command. Instead, the tape drive queues up commands and issues one at a time to the loader (attached media changer). However, this could be worked around by sending all 'info' commands to one address, and all 'action' commands to another address. Michael Banther asked if the address could be thought of as a virtual channel. Paul acknowledged that this is a useful abstraction. Erich Oetting likened the address to a TCP Port ID. Paul elaborated that the address was originally thought of as 'task' in a multitasking system.

## 6.5 IBM Library/Drive Interface

Lee Jesionowski presented T10/02-021r0, "IBM Library/ Drive Interface (LDI)".

With respect to the physical layer, Lee asked if the stop bit configuration should be controlled by a hardware jumper or a 'mode select' type implementation. The group came to no clear conclusion on this issue. Other signals implemented in IBM's physical layer include a drive present signal, which the library can test to determine whether a drive is present, and a wrap tool signal, which the drive uses to test RS-422 functionality in manufacturing.

Lee then presented a comparison of various features of each proposed protocol (IBM's, Quantum's and Seagate's).

Lee indicated that Seagate's flow control was presented as byte level due to the fact that the number of bytes in the frame is specified in the X\_RDY.

Lee suggested writing a profile against another already existing specification as a means of expediting the development of the ADI project. Joe Breher asked if SBP may form a suitable basis for the ADI project. Erich Oetting suggested that SBP is tied tightly to the FireWire physical layer. Joe Breher then asked if GPP would suffice. Paul Suhler indicated that he had investigated GPP for this purpose, but found it unnecessarily complicated.

The point was made that of the three layers, the physical layer will probably be the last thing to change. Michael Banther thereby advocated getting this pinned down first, as it has a hardware life cycle, so we better get started on it right away.

The discussion turned to our standards development philosophy. Rough consensus was reached that work should be done to expand Lee Jesionowski's comparison of the various proposals into a true requirements specification.

It was suggested that the group create a 'profile' derived from SSC-2. This profile would define a non-proper subset of SSC-2.

The group wondered what would the proper home for ADI as an ongoing development? Paul Suhler indicated that optical media changer vendors do not seem to be interested in this development, as evidenced by their lack of participation after being invited. Candidates for a 'home' included SSC-x, CAP, and SMC-x. It was asserted that covering the ADI business at the CAP meeting would have downside of difficulty of obtaining time at the CAP meeting. No conclusion was reached.

## 6.6 Putting a stake in the ground

Paul Suhler opened a discussion regarding "What is the next step?"

Erich Oetting asked what the statuses of the project proposals were. Bob Griswold indicated that the three project proposal documents are prepared, and that there is time allocated at the upcoming T10 Plenary meeting to discuss and vote on these proposals.

Bob Griswold presented the project proposal T10/02-008r0, "Automation/Drive Interface – Commands (ADC)". Kevin Butt asked if we are improperly restricting the direction of communication from the library to the drive. Paul Suhler pointed out that elsewhere in the proposal, the drive as 'initiator' is covered. Upon

this clarification, the group was satisfied that the proposal addressed the necessary points. Bob then presented the proposals T10/02-009r0, "Automation/Drive Interface – Transport Protocol (ADT)", and T10/02-010r0, "Automation/Drive Interface – Physical Layer (ADP)" as they differ from the ADC proposal. The group was satisfied that these proposals addressed the necessary points. An action item was given to Bob Griswold to present these proposals to the plenary.

## 7. Review of old action items

- a) Bob Griswold will produce a Project Proposal to submit to T10 for a standard for this interface -- complete
- b) Bob Griswold will prepare a shell document from which to develop the standard -- complete
- c) Everyone interested should critique each document and submit a written report to T10 – incomplete, but closed by affirmation
- d) Any company that wishes may submit their automation interface for consideration, by way of a T10 numbered proposal -- complete
- e) Paul Suhler will re-submit his proposal with the "Seagate Confidential" footer removed -- complete
- f) Each document that was presented today is to be submitted to T10 as a numbered proposal --- complete

Erich Oetting was given a new action item to bring a proposal for automation/drive interface from StorageTek.

## 8. Unscheduled Business

### 8.1 Where do we go from here?

Rod Wideman summarized his impression of where we are headed as:

- encapsulated 'SCSI' like command set, but engineered for low overhead;
- commands that are not SCSI CDB-like, but allow the things we do today;
- which commands are supported on primary (host) interface as opposed to supported on ADI?

Lee Jesionowski indicated that we might be able to think of this interface as a distinct logical unit within the drive and/or library.

- what is the drive state machine (load/unload) that can be used, independent of specific drive technology?
- need to define transport layer protocol
- what is our 'surrogate' or 'passthrough' model (as employed by attached media changers)?

Paul Suhler asked if we could expand on 'zero byte' status query. The group concurred that we do not want to preclude the use of this facility.

What will our status look like? It was pointed out that the new status format (SPC-3 Long sense data format) is a self-describing format. Then it was pointed out that we should probably first define what status we need to return before we define the format. The group agreed to this suggestion.

Erich Oetting pointed out the difference between static information and status that changes. The sense of the group was that it was somewhat improper to return 'status' information upon an 'inquiry' command.

action item - Paul Entzel has an action to disclose Quantum's physical layer specification.

action item - Joe Breher has an action to investigate whether he is able to disclose Exabyte's physical layer specification, and to disclose it to the group if able to.

action item - Erich Oetting will has an action to investigate whether he is able to disclose StorageTek's physical layer specification, and to disclose it to the group if able to.

action item - Michael Banther will generate a straw man summary of the various vendors' physical layers.

action item - Lee Jesionowski will discuss exceptions to SAM with George Penokie.

### 8.1.1 Consensus building on transport concepts

The group engaged in a rather undirected and wide-ranging discussion with the intent to build consensus on what we believe we know about our transport. Some of the points raised are covered below. The secretary has attempted to group these points into logically related topics. Consequently, the points are not necessarily covered in the same sequence as they came out in the discussion.

#### 8.1.1.1 Addresses and Channels

The physical addresses embodied in each frame are distinct from channel (DID, SID in Quantum's proposal). There can be multiple channels associated with each physical address pair.

It was pointed out that perhaps we should use the term 'Process Associator' for the 'channel' Concept, as that concept from Fibre Channel seemed to match closely. This was originally agreed to. However, most members of the group lapsed to the term 'channel' in natural conversation. The suggestion to use the term 'Process Associator' was withdrawn in favor of 'channel'. It was then pointed out that SRP uses the term 'channel' for this concept.

We attempted to define examples of usage of independent channels between an 'initiator'-'target' pair. Examples forwarded included:

- 1 channel for encapsulated SCSI sent from drive to library
- 1 channel for encapsulated SCSI sent from library to drive
- 1 channel for 'fastpath' status (similar to 'zero byte' status)
- 1 or more channels to implement vendor unique diagnostics, encapsulated legacy protocols, etc.

We will implement the concept of 'well-known channels', and define well-known channels that handle the above cases.

Shall we employ Queue Tags at frame header level? It was pointed out that SAM sets precedence for yes.

Any given 'device' may communicate with multiple source and destination 'devices'. Each 'device' will have a unique physical address.

ADI will allow multiple channels per physical address. These may be likened to Logical Units.

ADI will allow multiple concurrent 'operations' (SAM Task – like constructs) per channel.

#### 8.1.1.2 Frame identifier

The discussion turned to identifying a frame within the scope of a stream of frames. The secretary found himself confused from time to time, due to his preconceived notions of the meanings of certain overloaded terms. An attempt is made herein to use consistent terminology for the same concept. The reader will have to judge for him- or herself how successful that attempt is.

What is a sequence number? Is it a unique frame identifier? Almost – a frame, and its acknowledgement (ACK) or negated acknowledgement (NAK), each contain the same frame identifier. This frame identifier is otherwise unique.

The frame identifier appears at the transport layer only. The command layer should remain oblivious to its presence, contents, and usage.

Discussion over philosophy ensued -- is ADI to be a layered interface, or is it simple? Sergio Encarnacao indicated these are mutually exclusive.

Lee Jesionowski advocated simple counting identifier per each frame. If the receiver misses a sequential value, it has detected that it has missed a frame. The term of this in the IBM proposal is 'section number'.

Michael Banther indicated that HP has found that multiple frame sequences have not been worth the trouble of implementing them.

#### 8.1.1.3 Frames

All communication at the transport layer is to take the form of frames. We will not employ any 'primitives'.

The concept of 'Byte Stuffing' was discussed. As the secretary understands it, this is assigning a unique character to SOF and to EOF to simplify hardware acceleration of the lower protocol layers. Consensus seems to be that this is important. Frame synchronization is simplified by this mechanism. There was a consensus for employing EOF. Different SOFs for determination of frame 'type' was considered, and discarded.

There was a consensus that frame length should be indicated at the transport layer, independent of SOF and EOF delimiters.

It was initially decided that all frames would generate an ACK, NAK, or timeout. Sergio Encarnacao then pointed out that, as ACK and NAK are frames, we would of necessity ACK each ACK. It was thereby decided that all frames except for ACK and NAK shall generate ACK, NAK, or timeout.

Straw man frame:

- Byte Stuffing yes
- SOF, Length, EOF yes
- Hamming distance for SOF, EOF, esc  $\geq 2$
- 2 byte checksum?
- Default min, negotiable max payload size
- Frame level flow control

Lee Jesionowski asked if we support 'zero byte' minimum status. The rough consensus was that all communication under ADI would be in the form of a frame. The ADI protocol will not support any primitives, such as 0x00. Any nonframe communication is outside the scope of the standard.

There was consensus that the frame header shall include a type field.

#### 8.1.1.4 Flow Control

The following question was addressed: Do ACKs have SOF, EOF? The consensus was yes. ADI will make everything a frame for consistency.

There seemed to be widespread support for a buffer-to-buffer credit flow control model.

The following question was addressed: Are NAKs subclassed? Consensus was that detail within NAKs might not be useful for anything other than development debugging. There was a consensus to specify detailed NAKs, despite the fact that the subtypes may never be used. Michael Banther asked that all NAKs share a common nybble that can be masked off. This would allow a single test for all NAKs, and still distinguish all NAKs from all ACKs. There was consensus on this point.

It was decided that a sender shall receive an ACK, NAK, or timeout event before sending a subsequent frame for any given 'ITLQ nexus'-like construct.

The specific timeouts are to be determined at a future session.

Lee Jesionowski brought up the question of whether we have captured all the necessary identifiers in the header that allow recovery - IBM has implemented 'section number' and 'last section flag'.

action item - Bob Griswold to take concepts developed above as a straw man, and hammer it into an initial spec for ADT.

action item - Rod Wideman to develop straw man for minimum drive state polling frame.

action item - Paul Suhler to develop initial list of SCSI commands that should be mandatory and/or optional.

## 9. Next meeting requirements

The requirements for the next meeting were determined. The facilitator was directed to request 8 hours for the next meeting.

## 10. Review of action items

- 10.1 Paul Suhler to find out whether T10 rules require individual patent applications to be disclosed.
- 10.2 Erich Oetting to bring a proposal for automation/drive interface from StorageTek.
- 10.3 Bob Griswold to present the three project proposals at the T10 Plenary meeting.
- 10.4 Paul Entzel to disclose Quantum's physical layer specification.
- 10.5 Joe Breher to investigate whether he is able to disclose Exabyte's physical layer specification, and to disclose it to the group if able to.
- 10.6 Erich Oetting to investigate whether he is able to disclose StorageTek's physical layer specification, and to disclose it to the group if able to.
- 10.7 Michael Banther to generate a straw man summary of the various vendors' physical layers.
- 10.8 Lee Jesionowski to discuss exceptions to SAM with George Penokie.
- 10.9 Bob Griswold to take concepts developed above as a straw man, and use them to develop an initial specification proposal for ADT.
- 10.10 Rod Wideman to develop straw man specification for minimum drive state polling frame.
- 10.11 Paul Suhler to develop initial list of SCSI commands that should be mandatory and/or optional.

## 11. Adjournment

The meeting was adjourned at 4:00 PM.