

Draft Minutes  
Automation/Drive Interface (ADI) Working Group  
Ad Hoc Meeting  
T10/03-295r0  
8-9 September 2003  
1:00 PM – 8:00 PM (8 September)  
8:00 AM – 11:00 AM (9 September)

1. Introductions: Group

Paul Suhler called the meeting to order at 1:12 PM PDT on 8 September 2003. He thanked Microsoft for hosting the meeting. A table of the attendees appears at the end of these minutes.

2. Approval of the agenda: 03-294r0 Paul Suhler

Paul Suhler discussed the order of the discussion items.

Paul Suhler made a motion for acceptance of the modified agenda. Rod Wideman seconded the motion. The group passed the motion; no one objected or abstained.

3. Approval of previous meeting minutes: Paul Suhler

7-8 July 2003 meeting	03-236r0
21 July 2003 teleconference	03-262r0
4 August 2003 teleconference	03-270r0
18 August 2003 teleconference	03-276r1

Paul Suhler requested comments for the minutes of the 7-8 July 2003 meeting, the 21 July 2003 teleconference, the 4 August 2003 teleconference, and the 18 August 2003 teleconference: 03-236r0, 03-262r0, 03-270r0, and 03-276r1 respectively.

Paul Suhler moved for approval. Paul Entzel seconded. The group passed the motion; no one objected or abstained.

4. Review of action items: Michael Banther

- a. Bob Griswold to follow up with SNIA Interoperability Conformance Test Program (ICTP) Subcommittee regarding test/emulation tool. *Closed*
- b. Paul Entzel will write an appendix to ADT to describe an example login. *Carryover*
- c. Michael Banther will revise 03-239r1 per discussion item (a) of 03-270r0. *Closed*
- d. Rod Wideman will incorporate 03-239r1 as revised into ADC. *Closed*
- e. Rod Wideman will incorporate 03-263r1 as revised into ADC. *Closed*
- f. Rod Wideman to modify Tables 28, 33, and 39 of ADC per discussion item (a) in 03-276r0. *Closed*

- g. Rod Wideman to modify clause 4.2.4 of ADC per discussion item (c) in 03-276r0. *Closed*
- h. Rod Wideman to modify clause 5.2 of ADC per discussion item (c) in 03-276r0. *Closed*
- i. Rod Wideman to modify Table 1 rows h & i, Table 2, and the DAcc description of ADC per discussion item (c) in 03-276r0. *Closed*
- j. Rod Wideman to modify the paragraph of ADC describing DENOVR to mention Select Write Density in the lead-in per discussion item (c) in 03-276r0. *Closed*
- k. Rod Wideman to modify clause 4.2.2.1.1 of ADC to mention the ENABLED bit in the SMC Logical Unit descriptor per discussion item (c) in 03-276r0. *Closed*
- l. IBM to propose changes to clause 4.2.2.1.3 of ADC. *Closed*
- m. Rod Wideman to delete last sentence of 4.2.9 in ADC. *Closed*

5. Discussion items:

- a. SCSI Application Layer clause for ADT                      03-286r1                      Paul Entzel

Paul Entzel introduced the SCSI Application Layer clause. He pointed out that much of the text comes from SAS.

Paul Entzel pointed out that the use of the SCSI Transfer Ready IU for both Data In and Data Out renders the SAM requirement for the application client having ready all data or space associated with a command unnecessary. He stated that anyone implementing the ADT application layer interface will need additional transport protocol services (see SAM3r08, 5.4). We debated the wisdom of:

- a) Adding these services,
- b) Reverting to the existing SAM SCSI model, i.e., removing use of the SCSI Transfer Ready IU by the initiator,
- c) Keeping the requirement that initiators issue SCSI Transfer Ready IU's but not documenting the additional services needed, or
- d) Documenting the additional services needed elsewhere.

After some discussion we agreed to define (and hence document) these additional transport protocol services with the assumption that we will define them in sub-clause 8.2.

As several of us did not understand the intricacies of the transport protocol services, we worked through an example and explored variations to the existing model. Paul Entzel led that way. The fundamental problem occurs because ADT allows the receiver of the data to manage the data transfer while the SAM SCSI model requires the target device to manage the data transfer.

We considered adding four additional transport protocol services called by the initiator that mirror the Data Transfer transport protocol services currently called by the target.

With no good solution in sight, Paul Suhler built consensus in the group to bring Ralph Weber into the discussion. Ralph advised that we define the additional transport protocol

services in a separate table in the same sub-clause. The benefit is that we fulfil the SAM requirements whilst making the extensions and their justification clear. He indicated that we should have a model that explains the purpose of the extensions.

Paul Entzel agreed to add a separate table for the extended transport protocol services. He expects to add definitions for the new services as sub-clauses in 8.2.

Paul Suhler asked if we wanted to define transport protocol services for the Fast Access protocol. Several of us agreed that such definitions will add significant value. However no one expressed willingness to provide a proposal. Similarly we agreed that having a model section for use of the Fast Access protocol is a good idea, but no one stepped forward to provide a proposal.

Rod Wideman made a motion for the inclusion of 03-286r1 as modified, but without the additional transport protocol services, into ADT. Kevin Butt seconded the motion.

Paul Entzel indicated that he will address the additional transport protocol services as a separate proposal. Michael Banther asked whether that proposal will include a model section. We agreed that a model section should go into the text but did not reach agreement on whether or not Paul's proposal will include it.

Paul Entzel asked if any other concerns exist with this proposal (i.e., 03-286r1). Rod Wideman raised an editorial concern which Paul Entzel corrected in real time.

The group passed the motion; no one objected or abstained.

b. SCSI Request IU Buffer Allocation Length                      Email 2003/09/02 Paul Entzel

Paul Entzel described the uselessness of the BUFFER ALLOCATION LENGTH field in the SCSI Request IU given the presence of the FIRST BURST LENGTH field in that IU and the requirement to not send data until having received a SCSI Request IU with a non-zero FIRST BURST LENGTH or a SCSI Transfer Ready IU with a non-zero BURST LENGTH. The target device server knows the buffer allocation length after parsing the CDB. The FIRST BURST LENGTH and SCSI Transfer Ready IU throttle the movement of data.

Paul Entzel moved for removal of the BUFFER ALLOCATION LENGTH field from the SCSI Request IU, movement of the FIRST BURST LENGTH field of the SCSI Request IU into bytes 20 through 23, and removal of the paragraph that describes the BUFFER ALLOCATION LENGTH field. All changes occur in sub-clause 7.1.2 of ADTr06. Paul Suhler seconded the motion.

Paul Suhler asked if we wanted the FIRST BURST LENGTH field to remain four bytes. Paul Entzel replied that it matches the size of the BURST LENGTH field in the SCSI Transfer Ready IU. The group discussed the implications of the proposal.

The group passed the motion; no one objected or abstained.

c. Expected Frame Number Clarification                      Email 2003/08/21 Paul Suhler

Paul Suhler described the problem documented in the email. The group walked through the proposed changes in detail.

Michael Banther suggested changing item (d) in 4.7.1.3 such that it references the Expected Frame Number Counter instead of 'non-sequential frame numbers'. Paul Entzel proposed changing the text of item (d) in 4.7.1.3 to say, 'Frames with a frame number that does not match the Expected Frame Number counter (see 4.6.3).'

Kevin Butt proposed that we change the NAK status code of 'Non-sequential frame number' to 'Unexpected frame number' throughout ADT.

Paul Suhler moved for inclusion of the proposal in his email plus the two proposals described above. Kevin Butt seconded. The group passed the motion; no one objected or abstained.

d. Miscellaneous ADT Issues

Email 2003/09/02 Paul Suhler

NAK frame status codes

Paul Suhler reviewed his question/concern. Lee Jesionowski suggested changing the status code 87h from 'Maximum frame size exceeded' to 'Negotiated maximum payload size exceeded'. Rod Wideman and Susan Gray suggested using 'Maximum payload size exceeded'. We agreed on Susan's phrase. Paul Suhler stated that we can make this change as editorial.

NOP IU

We reviewed the purpose of NOP, i.e., it provides a way to wake up a Paused port. No change to the existing text is needed.

Limiting Auto Sense Data Size

We discussed the issues around creeping sense data sizes and the implications for minimum frame size support. We agreed that, in order to use autosense, the minimum frame size must be big enough to support the sense data because no mechanism currently exists to break up the sense data into multiple SCSI Response IU's.

Paul Entzel asked if everyone could live with limiting the sense length value in the SCSI Response IU to 256 bytes. Susan Gray suggested 252 to make a 256 byte IU. SPC3r10 calls out 252 bytes in sub-clauses 6.25 and 7.24.1.

Based on the discussion, we agreed to change the text in the next to the last paragraph of 6.5.4 of ADT to state, 'a port shall be capable of supporting a frame payload size of 256 bytes to accommodate a SCSI Response IU with a maximum autosense length of 252 bytes.' We agreed to change the reference in the last paragraph of sub-clause 7.1.3 from SPC-2 to SPC-3.

Kevin Butt made a motion for inclusion of the proposal in the preceding paragraph into ADT. Paul Suhler seconded the motion. The group passed the motion without objection or abstention.

4.3 ADT Port States, Table 1

Moving on to his second email, Paul Suhler reviewed his confusion when comparing the text for Port State P3: Recovering and item (b) in sub-clause 4.7.2.3. Paul Entzel

explained that the two statements support each other although a clarification to the existing text is in order. We agreed to an editorial change to both sections of text.

#### 4.7.1.3 Error detection by the frame receiver

We discussed Paul Suhler's concerns about the Awaiting Initiate Recovery IU NAK Status and the existence of a new port state when waiting for an Initiate Recovery IU. We considered abolishing the Awaiting Initiate Recovery IU NAK Status. We also considered adding a state machine in the receiver that would run upon detecting an error.

We agreed that we do need an additional state. Without a new state, a port that receives a bad frame with frame number  $n$  and then receives a good frame with frame number  $n$  cannot indicate that it expected to receive an Initiate Recovery IU. Paul Suhler agreed to bring in a proposal to add the state.

#### 4.7.2.2 Error Recovery for Port Login IU's

After a short discussion, Paul Suhler agreed that the existing text covers his concerns. No change to the ADT text is needed.

#### 7.1.2 Table 19

Paul Suhler asked if we really need to support Query Task. Paul Entzel stated that we do need to support it given that CAP intends to eliminate untagged tasks.

Paul Suhler proposed adding the Task Management Function Query Task with a value of 80h.

Paul Suhler made a motion to incorporate the proposal in the preceding paragraph into ADT. Rod Wideman seconded the motion. The group passed the motion without objection or abstention.

- e. Link Negotiation Error Recovery Email 2003/09/03 Michael Banther

Michael Banther described the concern. The first sentence of sub-clause 4.7.2.2 doesn't make it clear what state the port transitions to if it detects an error during a Port Login sequence. Lee Jesionowski and Michael Banther proposed changing the first two sentences to read, 'The transmission is detected on a Port Login IU, the recovery process is accomplished by remaining in P1 state and initiating a Port Login IU with a frame number of zero and a new exchange ID value.' Paul Entzel agreed to make this change as editorial.

- f. ADT N3 Necessary Rod Wideman

Rod Wideman asked the question, how does a port enter state N3? He pointed out that from state N2, a returning Port Login IU with the Accept bit set (and no parameter changes) causes a port state transition to P2: Active. Paul Entzel and Susan Gray pointed out that the port enters N3 from N2 when it has received a Port Login IU with Accept bit set and has sent a subsequent Port Login IU with the Accept bit set but has not yet received the ACK IU for the Port Login IU it just sent. Rod accepted Paul's explanation. The group agreed to leave the text as is.

## g. ADC Time of Day

Erich Oetting

Erich Oetting described a request for a command to set the time of day and set the volume identifier (e.g., barcode label). Lee Jesionowski stated that the Volume Serial Number concept should involve storage in MAM. Susan Gray asked where it gets reported. Lee replied that it's important in logs in field failures.

We generally brainstormed these ideas.

We agreed to defer work on these ideas until ADI-2. Paul Suhler will add them to the list of pending topics, 03-133.

## h. ADT Technical Questions

Email 2003/08/05 Paul Suhler

6.5.3.3 Table 14

Paul Suhler asked what constitutes a framing error. Paul Entzel replied that it's anything the UART would detect as a framing error. We agreed to clarify the existing text by changing to 'byte framing error' as an editorial change.

7.1.2 SCSI Request IU

Paul Suhler described the possible race condition. One can occur if the initiator wants to receive more data and sends a SCSI Transfer Ready IU at the same time as the target decides that it has sent all of the data it is going to send and sends a SCSI Response IU. The condition arises only if the target sends less data than the initiator expects. Paul Suhler walked us through an example where the initiator wants to receive 9k bytes and the target provides only 8k bytes.

Paul Suhler suggested having the receiver discard SCSI Transfer Ready IU's that appear outside of an open exchange. This suggestion relies upon the un-stated assumption that the exchange opens upon exchange of the SCSI Request IU and closes on exchange of the SCSI Response IU when using the Encapsulated SCSI protocol.

Rod Wideman asked about the sequence where the target receives a SCSI Transfer Ready IU but has no data to send. It sends the SCSI Response IU. We agreed that this sequence represents correct behaviour. Rod then explored the case of the target receiving the SCSI Transfer Ready IU just after sending the SCSI Response IU. Rod suggested having the target positively acknowledge the SCSI Transfer Ready IU even if the target subsequently ignores it. Paul Entzel pointed out that the initiator transport layer should report the completion of the upper-layer function call that the application used to cause the sending of the SCSI Transfer Ready IU *before* it reports the receipt to the upper-layer application of the SCSI Response IU.

We explored the possibility of having the target transport layer wait for the last SCSI Transfer Ready IU before sending the SCSI Response IU. If the target transport layer doesn't have anymore data to send, it responds with a SCSI Data IU with zero data and then sends the SCSI Response IU. This solution keeps the target transport layer from having to receive a SCSI Transfer Ready IU for an open exchange after sending the SCSI Response IU for that same exchange. However it requires the target transport layer to know the amount of data requested by the initiator for the entire exchange. Hence it

argues for putting the BUFFER ALLOCATION LENGTH field back into the SCSI Request IU (see discussion item [b]).

Susan Gray suggested adding a SCSI Final Data IU that replaces the last SCSI Data IU and the SCSI Response IU.

Kevin Butt suggested using the Sequence Initiative concept from Fibre Channel. This idea requires an extra bit in each frame. The bit indicates which side has the right to send the next frame in the sequence (i.e., exchange).

Paul Suhler recapped the possibilities:

- a) Add a SCSI Final Data IU.
- b) Require the target transport layer to wait for the next SCSI Transfer Ready IU before sending the SCSI Response IU.
- c) Allow the target to positively acknowledge and ignore a SCSI Transfer Ready IU received outside of an open exchange.

We compared the three options.

Kevin Butt raised the possibility of intermixed exchanges and the consequences thereof.

Paul Suhler agreed to bring in a proposal to sort out the acknowledgement and discarding of a SCSI Transfer Ready IU. The group will reconsider this issue in the light of Paul's proposal.

i. ADT Comments

Email 2003/09/03 Michael Banther

4.4 ADT Link Negotiation States

Michael Banther described the question. Paul Entzel stated that the text describing the correct behaviour upon receiving a Port Login IU in a new exchange, and currently in the description of negotiation state N1, should apply to all negotiation states. Kevin Butt agreed. Paul Entzel agreed to move the text into a sub-clause as an editorial change.

6.5.3.2 ACK Information Unit

Michael Banther described the concern.

Paul Entzel moved for adoption of the proposal stated in his response email copied below. Kevin Butt seconded the motion. The group passed the motion without objection or abstention.

[Michael Banther] The last sentence in Table 1, state P5 (sub-clause 4.3, p. 14) states, 'All other frames received while in this state shall be ignored.' However 6.5.3.2 states, 'Except for acknowledgement IU's, a port shall send an ACK IU for every frame that it receives without error.' If the receiver in state P5 is going to ignore frames received without error, then 6.5.3.2 needs modification to remove this inconsistency. (editorial)

[Paul Entzel] I can't remember why we decided to ignore frames in the Logged Out state. Wouldn't it be better to respond to other frames with a NAK IU and

status of Rejected, port is logged out (85h)? Since this is a contradiction in the standard I think it qualifies as technical. Let's talk about this next week.

j. Completion schedule for ADT

All

Michael Banther stated that he believes that plenty of good work remain to be done but that he is unwilling to sign up to write proposals under imminent pressure to get to letter ballot. Kevin Butt stated that George Penokie had complained about the quality of the ADC draft standard that is now in the letter ballot comment phase. A general consensus emerged that we should not expect to submit ADT to a letter ballot until at least the November meeting.

Michael Banther agreed to provide a proposal to define when exchanges open and when they close.

Susan Gray agreed to provide a proposal to define state diagrams for the Port states (4.3) and the Link Negotiation states (4.4)

6. Unscheduled business:

a. Action Item (a)

Kevin Butt pointed out that action item (a) had remained open for a very long time with low probability of closure by Bob Griswald. After a short discussion, Paul Suhler moved to drop action item (a). Erich Oetting seconded. The group passed the motion; no one objected or abstained. Paul Suhler pointed out that anyone who wanted to volunteer to pick up this work in future would be welcome.

b. Port Login process versus Negotiation process

Susan Gray pointed out that ADT uses the phrases 'Port Login process' and 'negotiation process' interchangeably. We searched the document and discovered only one instance of 'negotiation process'. Hence Paul Entzel will change that instance to 'Port Login process'.

c. 4.4 ADT Link Negotiation States

Kevin Butt asked about the 15 second time in the description of N1: Negotiating state. He asked if we shouldn't state the time period in terms of characters sent since the existing text doesn't specify the baud rate. Lee Jesionowski stated that he likes the absolute limit as did Rod Wideman. Paul Entzel pointed out that the text doesn't prohibit a port considering the lack of a response an error even though 15 seconds hasn't passed. We agreed to keep the text as is.

d. ADT figures 2 and 3

Rod Wideman pointed out that these figures no longer match ADC. He asked whether we want ADT and ADC to match. Kevin Butt stated that they should. Rod will forward the ADC figures to Paul Entzel. Paul Entzel will replace the existing figures in ADT with the copies from Rod.



## 7. Next meeting requirements:

Paul Suhler

The group will hold teleconferences on 22 September 2003, 6 October 2003, and 20 October 2003. These teleconferences will begin at 8:00 AM PST and conclude at 10:00 AM PST.

The group will hold a meeting 3-4 November 2003 during T10 plenary week in Austin, TX. The meeting will begin on the 3<sup>rd</sup> at 9:00 AM CST and conclude at 6:00 PM CST. The meeting time on the 4<sup>th</sup> start at 9:00 AM CST and will adjourn when the FCP-3 working group meeting begins.

## 8. Review new action items:

Michael Banther

- a. Paul Entzel will revise 03-286r1 per discussion item (a).
- b. Paul Entzel will incorporate 03-286r1 as revised into ADT.
- c. Paul Entzel will propose definitions for the four additional transport protocol services documented in discussion item (a). This proposal will include a new table for sub-clause 8.2 showing the relation of the new transport protocol services to the SCSI client-server model. It will also include a model describing the use and purpose of the new services.
- d. Paul Entzel will incorporate the proposal documented in discussion item (b) into ADT.
- e. Paul Entzel will incorporate the proposal documented in discussion item (c) into ADT.
- f. Paul Entzel will incorporate the proposal documented in discussion item (d) into ADT.
- g. Paul Suhler will write a proposal to add the Pending Recovery port state.
- h. Paul Suhler will add the time-of-day and volume identifier features to 03-133 ADT-2 features.
- i. Rod Wideman will provide updated figures 2 and 3 from ADC for ADT.
- j. Paul Suhler will provide a proposal to positively acknowledge and discard a Transfer Ready IU received under a variety of conditions: before sending a SCSI Response IU when the target has no more data to send, after sending the SCSI Response IU but before receiving the subsequent acknowledgement, and after both sending the SCSI Response IU and receiving the subsequent acknowledgement.
- k. Michael Banther will provide a proposal to define when exchanges open and when they close.
- l. Susan Gray will provide state diagrams for the Port states (4.3) and the Link Negotiation states (4.4) as a proposal.

## 9. Adjournment:

Group

Kevin Butt made a motion for adjournment. Michael Banther seconded the motion. The group passed the motion unanimously. Paul Suhler adjourned the group at 10:58 AM PDT on 9 September 2003.

## Attendees:

Name	Organization	E-mail
Rod Wideman	ADIC	rod.wideman@adic.com
Paul Suhler	Certance	paul.a.suhler@certance.com
Michael Banther	HP	michael.banther@hp.com
Kevin Butt	IBM	kdbutt@us.ibm.com
Lee Jesionowski	IBM	ljesion@us.ibm.com
Susan Gray	Quantum	susan.gray@quantum.com
Paul Entzel	Quantum	paul.entzel@quantum.com
Erich Oetting	StorageTek	erich_oetting@stortek.com