

October 30, 1999

T10/99-325 revision 0



To: T10 Membership
From: Bob Snively, Sun Microsystems
Subject: FCP-2 changes from revision 03 to revision 04.

At the FC-TAPE/FCP-2 meeting of October 5, 1999, the document T10/99-247r1 was extensively considered. FCP-2 revision 3 was published at the same time. The discussion of FCP-2 revision 3 and T10/99-247r1 created a number of requirements for changes in FCP-2, which are incorporated in revision 4 of the FCP-2. The changes are documented and explained here. The changes incorporate all the action items requested in the following E-mail:

Stewart Wyatt, T11.3/T10 Joint Activity Group Action Items - 5 Oct., dated 19 Oct., 1999, 10:57.

Additional items have been presented in the following E-mails:

Stewart Wyatt, Review of FCP-2 Revision 03, dated 29 Oct., 1999, 08:30

Jim Coomes, FC Mode Page DSA/RHA Bit, dated 12 Oct., 1999, 14:53

Horst Truestedt, Ken Hallam, Dennis Talluto, FCP IUs, thread beginning Oct. 5, 1999

Dale LaFollette, REC, S_ID in payload, dated 12 Oct., 1999, 11:05

1.0 Resolution of items from T10/99-247r1

The following are items in T10/99-247r1 that were discussed and resolved in the meeting of October 5, 1999. Those changes to FCP-2 documented by T10/99-247r1 that were not changed by the discussion or by subsequent E-mail proposals are not included here. Where applicable, the resolutions are identified with the corresponding action item from Stewart Wyatt's E-mail.

1.1 Rules for ELS generation before Login

Most FCP devices compliant with FC-PLDA limit the ELS codes that may be used before a Login has been successfully completed. At present, FCP is silent on this and FC-TAPE has expressed rules similar to FC-PLDA. Robert Kembel's comment #29 on FC-TAPE, which requires clarification on this issue, has not been resolved.

The committee has requested Bob Snively and Bob Kembel to prepare a list of ELSs that do not require implicit or explicit login. This list, when approved, will be included in an annex of revision 02 of FCP-2 until it is transferred to FC-FS.

Bob Kembel indicates:

I was wondering which Extended Link Services that you feel should require PLOGI before they are accepted. The only one that I would accept without argument is PRLI. Are there others?

This is partially included in section 11.10, but more work remains.

This work item is not yet complete and should be included in FCP-2, revision 04.

1.2 References for error examples

To assist in understanding the error recovery procedures, cross references between the error recovery examples in Annex C and the error recovery descriptions in clause 11 are provided. So far, only the general cross-reference has been completed. **I will work on these as time permits, but it will not be complete until FCP-2 revision 04.**

1.3 DSA bit clarifications

A proposal has been put in place changing the name and clarifying the description of the DSA bit. The proposal, 99-226r2 from Jim Coomes, was approved in the last working group meeting. The changes are now installed in FCP-2 revision 3.

See Jim Coomes Letter.

My note to myself says: Bit has no effect on targets not performing loop initialization.

1.4 Name server extensions

Two recent documents have addressed FCP-2 specific capabilities in the Fibre Channel name server. One possible approach to these is to include them as a normative annex in FCP-2 as the first of several protocol specific extensions. A second possible approach is to include these as an FCP-2 device server definition within the body of FCP-2. I still personally prefer that the name server accept these in some format or other, possibly pre-documented by the appropriate protocol documents.

After considerable discussion, the following approach has been selected.

Create an FC-4 specific object name space which is defined by the individual protocol document.

Note that initiator attributes must also be defined in this way.

Charles Binford suggests that zoning access to this information from the viewpoint of the target may be critical. He points out that some vendors already provide separate images of Inquiry strings and logical unit lists on the basis of initiator. This can only be determined when asked for from the initiator that actually is interested. The meeting group determined that this was outside the scope of FCP-2.

A proposal document is required for this function.

1.5 LSI 004 Restriction of FCP_CONF usage (Technical)

page 9, section 4.4, 4th paragraph Last sentence restricts FCP_CONF from being used for non-queued, non-check FCP_RSPs. Why? This sounds like a profile type restriction, not a standard.

This was discussed in the August meeting. Dal Allan provided a model for use in defining and refining the usage of FCP_CONF. The model had to be extended to allow the described behavior.

There was further discussion at the October 4 meeting. The text must specify that the FCP_CONF function is never used unless the FCP_STATUS field is valid. In particular, it is never used for task management functions or when an FCP_RSP_DATA is provided without status. This was accepted by the committee and is installed in Revision 4.

Should FCP_CONF be allowed for FCP_RSP with FCP_RSP_INFO only?

As part of this study, it was proposed that status and response fields be mutually exclusive in FCP_RESP. The editor will study whether any conflicts occur if this is accepted.

To be done in FCP-2 Revision 04

1.6 LSI 012 Reject of retransmission requests (Technical)

page 30, 7.1, 1st paragraph at top of page Why is a target prohibited from rejecting retransmission requests for XFER_RDY or RCP_RSP. This seems an unreasonable requirement. It is not obvious to me that tape drive (for example) can in all cases successfully recover an interrupted write command. Further, FCP_RSP may not be available if a device is a bridge and serving both disk and tape. The device may support SRR, yet the LUN with the error may be a disk.

If this comment is rejected, then the flavor of the XFER_RDY (i.e. Read or Write XFER_RDY) needs to be clarified (specify Write XFER_RDY).

The general principle that should be followed is that retransmission requests should always be accepted unless there is an error or other condition that prevents their execution. A device should not commit itself to retry, then reject all possible attempts to perform a retry. In view of this, the suggested modification is probably desirable.

[done in Rev 3](#)

Additional issues concerning clarification of the valid responses to SRR were raised in the discussion. These are clarified in the appropriate sections.

- 1) The target may choose to end a command with FCP_RSP rather than retry.

[included in 11.2.6, rev 02](#)

- 2) SRR ACC timing with respect to the response is not defined. This should be the same value specified for normal ELS responses and need not be stated. After further discussion in the October 4 meeting, the following clarifications were made.

- a) The SRR ACC must be transmitted before any recovery operation frames are transmitted. It is possible that garbage frames may continue to flow until SRR ACC is transmitted.
- b) Retry is only made on the requested data. The target cannot make changes in the specification of the data to be transmitted, even if the restriction makes the recovery impossible.

Needs to be done in Rev 04.

3) FCP_RSP may not be retryable on some LUNs of a device.

done in rev 03

1.7 LSI 030 ** RO during recovery (Technical)

page 58, 11.2.5 FCP_DATA Recovery Write The last sentence (above the note) says the SRR contained an RO. However that is not how SRR is currently defined. The RO field is only valid if requesting data, in this case we are requesting a data descriptor. However, because Writes have a handshake between every data IU, I don't think it is a problem to not have the RO specified. The implication is retransfer the last write data IU. Need to change the words in this paragraph.

SRR is presently defined incorrectly. RO is presently defined only for solicited data retry. It is now clarified that it applies for an XFER_RDY missing after a solicited data transfer as well. It will also be clarified that RO is that of the lowest unsuccessfully transmitted data.

done (see 7.1 and 11.3.5)

After further review at the October 5 meeting, it was decided that the FCP_XFER_RDY must have the same relative offset specified by the SRR if the recovery is to be allowed at all. The SRR beginning relative offset is required to be on a 4-byte boundary.

Modified text installed in revision 04.

1.8 LSI 031 * FCP_RSP retransmission (Technical)

page 59, 11.2.6 FCP_DATA Recovery Read Wording in second paragraph assumes the target has already transferred FCP_RSP once. This may not be true. Current words: The Target shall retransmit the requested data in a new Sequence, and then retransmit the FCP_RSP. Suggested new words: The Target shall retransmit the requested data in a new Sequence, and then complete the I/O as normal, including transmitting or retransmitting the FCP_RSP.

The suggested change is accepted.

done

After further review at the October 5 meeting, it was decided that the FCP_XFER_RDY must have the same relative offset specified by the SRR if the recovery is to be allowed at all.

Modified text installed in revision 04.

1.9 LSI 032 ** Recovery RO (Technical)

page 59, 11.2.6 FCP_DATA Recovery Read Wording in second paragraph implies target may modify the RO. The current SRR definition requires the target to start at the specified RO.

After further review at the October 5 meeting, it was decided that the FCP_XFER_RDY must have the same relative offset specified by the SRR if the recovery is to be allowed at all.

Modified text installed in revision 04.

1.10 LSI 034 * OX_ID and RX_ID in REC (Technical)**

page 74, B.3.1 Read Exchange Concise (REC) The paragraph at the top of the page tells initiators to check the OX_ID and the targets to check the RX_ID and both to ignore the other half of the X_ID. I believe both sides need to check the entire X_ID. (See LSI 033 for an example)

done

After further review and E-mail discussions after the October meeting, it was concluded that the S_ID must be available because the REC may be performed by the target in some recovery cases, and the OX_ID/RX_ID context may not be clear from the exchange containing the REC ELS. No change is made in the S_ID definition.

1.11 HP 14: Correct figure C.9

Annex D, page 95, “Figure C.9 - Lost Read Data, Last Frame of Sequence” The Class 3 Error detection drawing has the REC and ACC arrows in the wrong direction. After further review, it is apparently the proximity of the label to the arrow that must be corrected.

To be done in revision 04.

1.12 Specification of formats for ELS

Section 11.6 (now 11.7) is a hold-over from profile days. Should we move this to an informative annex? It specifies the details of the contents of the ELSs used in the recovery procedures, probably in a redundant manner with respect to FC-FS and other documents.

After discussion at the October 5 meeting, it was decided to move this text to an informative annex.

to be done in revision 04

1.13 Multi-initiator

Section 11.8 (now 11.9) contains some multi-initiator behavior definitions. This should be moved to an informative annex. It specifies behaviors that may conflict with SAM-2 and SPC-2.

After discussion at the October 5 meeting, it was decided to move this text to an informative annex.

to be done in revision 04

2.0 Review actions required

2.1 Clarification that link error recovery works if in-order

11 - It needs to be specified, in big bold letters, that the link error recovery procedure specified here **ONLY WORKS ON AN IN-ORDER TOPOLOGY**.

Accepted

In the July meeting, there was considerable discussion about this question. Dave Peterson feels that most of the work required to make out-of-order operation behave correctly is already included. Dal Allan and Carl Zeitler believe that out-of-order operation should be allowed if at all possible. Please review section 11 carefully for discrepancies that may cause failures during recovery of out-of-order transfers. If there are none, we can remove the “in-order-only” restriction.

2.2 Behavior of PRLI

There is an implicit assumption in the choice of bits in the PRLI request payload and in the PRLI accept payload that the PRLI request is always performed by an initiator. Since devices can label themselves as both and since there is no explicit rule that says the PRLI request is always done to a device that is only a target, I assume that the bits useful for initiators should be placed in both the PRLI request and the PRLI accept payload.

The following bits were copied over from table 9 to be placed in section 6.2.7, table 10.

Confirmed Completion Allowed

Data Overlay Allowed

I have not yet adjusted the text to clearly identify the bits as being sourced by initiators and not set by targets. The reason is that the PRLI image creation capabilities seem to be somewhat at odds with the informative and negotiative intent of the capabilities bits in FCP-2. This will be addressed as a separate issue in 2.3, which proposes that process associators be made obsolete in FCP-2.

After discussion, it was concluded that either 0 or 1 Process Associators may be created for an FCP compliant SCSI device. That simplification is sufficient to guarantee that the present text is acceptable with respect to Process Associators.

To be done in revision 04.

2.3 Obsolete process associators

There is an informal proposal for making process associators obsolete, at least for FCP-2. I will be making that proposal formal for the next FC and FCP-2 meetings.

Note that the Process Associator definitions do not create a consistent architecture with SCSI and with PRLI. The problem is:

- 1) Process associators do not take part in the SCSI LUN or initiator addressing.
- 2) Process associators do not take part in separating CRN or exchange recovery.
- 3) The theoretical basis for process associators implies that independent processes are operating in the host. However, reservation protocols use as their primary parameter various initiator port identifiers, implying that the independent processes are not independent for at least that major part of the SCSI behavior.
- 4) PRLI has some problems separating initiator/target capabilities by process associator, since the process associator is not part of the addressing structure.

The best way to avoid having to figure out rational answers to all these questions (which will inevitably violate other standards) is to simply make them obsolete and not use them.

This is made more compelling by the fact that they are essentially unusable with their present definitions.

The resolution depends on review work being done by at least one company on the possible uses of Process Associators.

not done in revision 03

2.4 Incorrect use of Recovery Abort

Section 11.4.1 (now 11.5.1) defines a number of cases where the recovery abort protocol is supposed to be executed. Many of these definitions are incorrect. In particular, ABTS should not be invoked following CHECK CONDITION status for resets or microcode changes. These are normal SCSI behaviors. Many of them conflict with section 8.1.4, which requires ABTS-LS only for exchanges whose state is ambiguous. I propose that we do the following:

- 1) a-1 should apply only to ambiguous exchanges.
- 2) a-2, both sections should be deleted.
- 3) b-1 and b-2 should apply only to ambiguous exchanges.
- 4) b-3 should be deleted.

to be done in revision 04.

2.5 Read error recovery examples

Dale LaFollette completed an action item to provide some examples for the recovery of errors in multiple block read operation. Dale has completed this action item. The editor will make some minor editorial corrections and included this in revision 04.

to be done in revision 04.

2.6 Clearing effects of PRLI/PRLO

This change resulted from discussions caused by last months observation by Rob Basham that a PRLI would cause exchanges in progress from previously logged in initiators to be aborted. The PRLI would also reset Mode pages to their power on value.

Bob Snively noted that a PRLO changes the status of a node from a SCSI target or initiator to an undefined Fibre Channel node unless an implicit PRLI is in effect.

During the following discussion the differences between shared and unshared node pages were reviewed. The following changes were agreed upon.

- 1) The effect of a PRLI on shared mode pages: Pickup current values if any other initiator is logged in. If not, the current value is the saved or default value.
- 2) The effect of a PRLI on unshared mode pages: Use saved values or default values.
- 3) The effect of a PRLO on shared mode pages: Clear current mode pages only if this is the last initiator to logout.
- 4) The effect of a PRLO on unshared mode pages: Clears current mode pages for the initiator originating logout.

Bob Snively noted that this was only necessary if word 0 bit 13, establish image pair, equals 1. If no image pair is defined by the PRLI, the command is only establishing capabilities and no clearing is required.

to be done in revision 04.

3.0 Comments on FCP-2 Revision 03 by Hewlett Packard

Stewart Wyatt of Hewlett Packard has provided the following comments on FCP-2 revision 03 in an E-mail dated October 29, 1999.

3.1 SCSI-3 s/b SCSI (editorial)

In Clause 2.2, References under development, the titles of the SCSI-3 Architectural Model and the SCSI-3 Primary Commands have been changed to the SCSI Architectural Model and SCSI Primary Commands. If the SCSI committee intended this change to be universal, a search should be made for "SCSI-3" in the document since this occurs in a number of places.

The change is accepted.

to be done in revision 04.

3.2 Editorial

In Table 1 on page 7, the last entry is missing the "IU". It should read "Confirmation IU (FCP_CONF)".

The change is accepted.

to be done in revision 04.

3.3 FCP_RSP_INFO not sent with FCP_SNS_INFO (Technical)

Clause 4.2 Device management, 5th paragraph, second sentence, “That payload contains the SCSI status and, if an unusual condition has been detected, THE SCSI REQUEST SENSE information and the FCP response information describing the condition.” In the October meeting, I understood that a decision was made that only one of the FCP response and the SCSI Status would be valid.

The change is accepted. See 1.5.

to be done in revision 04.

3.4 Editorial

Clause 4.2 Device management, last paragraph, first sentence: “Both FC-PH and SCSI allow the initiator function in any FCP_Port and the target function in any FCP_Port.” Does the reference to FC-PH belong here?

The change is accepted.

to be done in revision 04.

3.5 Editorial, FCP_CONF correction

Table 2 - Discovery of FCP capabilities. In the Capability Column referencing FCP_CONF, I would think the “Initiator performs the FCP_CONF” (at the targets request) and the “Target accepts FCP_CONF”.

The wording will be modified to:

Initiator generates FCP_CONF

Target requests FCP_CONF

The change is accepted in principle.

to be done in revision 04.

3.6 Clarify RX_ID is required (Technical)

Technical. Clause 5.6.10 RX_ID. The text associated with the unassigned value of FFFFh should be modified to note that this value is prohibited for Class 2.

The proposed change appears to be required if REC is to be used, but there does not seem to be any requirement for it if REC is not used. (See 1.10.)

This change needs to be reviewed by the committee before it is accepted for FCP-2 revision 04.

3.7 Define and explain IPA (Editorial)

Clause 6.2 Process login (PRLI). The headings of clauses 6.2.1, 6.2.2, 6.2.3 and 6.2.4 contain an abbreviation “IPA”. The abbreviation is in clause 3.2, but there is no explanation of or introduction to initial process associators in the text. A brief explanation in clause 6.2 would be helpful.

The change is accepted.

to be done in revision 04.

3.8 ACA s/b CA (technical)

Clause 8.1.1.4 Task Management flags, Byte 10. In the Clear ACA text fourth paragraph, “If the ACA bit in the CDB field is set to 0, the automatic sense operation performed by the presentation of the FCP_RSP IU shall clear the ACA condition.” Actually since the ACA bit is 0, a conditional allegiance (CA) exists, not an ACA. (SAM-2 rev 12 clause 5.6.1 page 53.)

The change is accepted.

to be done in revision 04.

3.9 Editorial

8.1.1.4 Task Management flags, Byte 10 The entry for Clear Task Set, first paragraph, last sentence, mistakenly repeats a sentence from the previous entry It should state that CLEAR TASK SET bit is mandatory rather than the TARGET RESET bit.

The change is accepted.

to be done in revision 04.

3.10 Editorial, Recovery_Qualifier terminology

10. Editorial. 8.1.2.2 Recovery Abort. There are several uses of the term “recovery qualifier” in this clause. It is usually printed as Recovery_Qualifier in FC-PH.

The change is accepted. The appropriate wording from FC-FS will be used.

to be done in revision 04.

3.11 Correct FCP_XFER_RDY reference on read (Technical)

8.3 FCP_DATA IU. The fourth paragraph, first sentence, states that “If required by the PRLI FCP service parameters, each inbound and/or outbound FCP_DATA IU shall be preceded by an FCP_XFER_RDY IU...” This sentence and the remaining paragraph may need to be rewritten since FCP now prohibits FCP_XFER_RDY for reads.

The change is accepted.

to be done in revision 04.

3.12 Consideration of write overrun (Technical)

8.3. FCP_DATA IU. The sixth paragraph which starts with the phrase, “During a write operation”, the third sentence is confusing. “If the write operation requires a total amount of data less than the amount of data provided by the initiator, the target shall discard the excess bytes and indicate that an overrun has occurred by setting the FCP_RESID_OVER bit in the FC_RSP IU.” I think that this sentence is addressing the issue of an initiator requesting more data in the CDB than is allowed by the FCP_DL. Since the initiator is the one providing the data, it should know the precise length of data to transfer. If this is a case where the FCP_DL is not equal to

the length of a variable block of data to be transferred or to the command length times the fixed block length, the target should reject the command with an RSP_CODE of “FCP_CMND fields invalid”. Otherwise assuming that the target allocated a buffer of length FCP_DL to receive the data, this statement requires the target to accept all of the data and overrun and corrupt its buffer. I would think it would be better for the target not to request more than FCP_DL bytes of data and then send the response. The initiator can then clean the residuals out of its own buffers.

To fix this problem I would prefer the sentence be removed.

Analysis:

The case should actually be treated as an under-run, since the data required to be transferred is less than FCP_DL and less than the amount of data offered. The case is actually a case where it is uncertain whether the FCP_DL was generated incorrectly, the CDB contained the incorrect information, or the interpretation by the logical unit was either flawed or encountered an error. Therefore, the only change that needs to be made is to change the words “FCP_RESID_OVER” to “FCP_RESID_UNDER”.

to be done in revision 04.

3.13 Unlimited FIRST BURST SIZE (technical)

9.1.1.10 FIRST BURST SIZE. “A value of zero indicates that there is no first burst size limit.” This only seems practical for a target with an infinite buffer. Is this what is intended?

This case requires either an infinite buffer or a commitment not to exceed the known buffer capabilities. Such a commitment is made known by mechanisms outside the scope of this standard.

No change is required in the document.

3.14 MCM parameters (technical)

Clause 9.1.3.10 Control MCM. The abbreviations MCM, CMR and BMCM are not defined anywhere in this document. These entries need more explanation and cross references to other standards. The cross references should also be added to clause 2.2 References under development. The reader needs some clue about what the relevance of these fields.

The change is accepted.

to be done in revision 04.

3.15 Editorial

Table 29 Timer Summary. Note 4 FCP_CONFIRM should be FCP_CONF.

The change is accepted.

to be done in revision 04.

3.16 Clarify Sequence_Qualifier (technical)

Clause 10.2 Resource Allocation Timeout (R_A_TOV). The term “Sequence qualifier,” is referred to as Sequence_Qualifier in FC-PH. The parenthesis in

the text implies that this is the SEQ_ID and the SEQ_CNT. In FC-PH 18.1.2 Sequence Identification, The Sequence_Qualifier is defined to be the S_ID, D_ID, OX_ID, RX_ID and SEQ_ID.

The change is accepted.

to be done in revision 04.

3.17 Editorial

Clause 11.2.2. Formatting problem of too many blank lines after the list.

The change is accepted.

to be done in revision 04.

3.18 Editorial

Clause 11.4 FCP Error Recovery (Target, class 2 and other acknowledged services). The paragraphs need to be reordered for the text to make sense. The fourth paragraph needs to immediately follow the second, since the second paragraph raises a problem that is addressed in the fourth paragraph and the third paragraph itemizes a different issue.

The change is accepted.

to be done in revision 04.

3.19

19. Technical. Clause 11.8 SCSI Target device level error behavior. Third paragraph - what if ULP resources are unavailable for a non-queued target (a disk drive doing an XOR operation or a tape doing a copy command). Is TASK SET FULL still the appropriate status?

This section is moved to an informative annex and corrected to comply with SAM, SAM-2, and SPC-2. The text contains some obsolete profile statements. (See 1.13). This issue is also being discussed as part of SPC-2, where the QUEUE FULL status is defined incompletely.

After review, to be done in revision 04.

3.20 Editorial

C.2 FCP write example, frame level. The first sentence refers to figure A.1. This reference and the 4 following figures should be “C” not “A”.

The change is accepted.

to be done in revision 04.

3.21 Editorial

C.3 FCP read example Figure A.3 (which should be C.3) has the FCP_RSP arrow in the middle of the text.

The change is accepted.

to be done in revision 04.

3.22 Editorial

Annex D. The arrow for the FCP_CMND extends too far for the Class 3 Error Detection examples in Figure D.3, D.5, D.7, D.8, D.9, D.10

The change is accepted.

to be done in revision 04.

3.23 Clarify Discovery Procedure (technical)

Annex E SCSI Device Discovery Procedure. Steps E.2 and E.3 seem initiator centric. It is not clear to me that a target should perform all of these steps. For example, E.2 number 5 Register for State Change Notification, in some reflector discussion someone suggested that this was inappropriate for a target. In E.3 number 1 Obtain a map of the loop. (Or poll all if a loop map is not available.) This requirement is clearly an initiator requirement as targets don't poll initiators. It seems this needs to define the different responsibilities for targets and initiators.

The change is accepted.

to be done in revision 04.

4.0 Comments on FCP-2 Revision 03 by Store Tek

Dave Peterson of Storage Technology reviewed FCP-2 revision 03 and made notes on the document. From his notes, the editor has extracted the following comments.

4.1 Editorial, 11.1.1

“attaching” s/b “communicating with”

The change is accepted.

to be done in revision 04.

4.2 Change from optional to recommended (technical)

Clause 11.1.1 presently treats the error detection query response as optional. While in use, the FCP-2 device should make every attempt to satisfy the error detection query.

This proposed change is contrary to previous treatment of error recovery capabilities as optional.

please review. The change should be rejected.

4.3 Clarify usability of error recovery (technical)

In Clause 11.1.2, the implied restrictions on which types of devices use error recovery procedures should be relaxed. The first sentence of the first paragraph is changed to:

SCSI devices may use the mechanisms described in this chapter to detect the presence of link errors, then perform retransmission procedures that will allow the commands to be completed without requiring complex higher level recovery algorithms. Such recovery may be required for the proper op-

eration of SCSI logical units that depend critically on command ordering and maintaining records of internal device state.

The change is accepted.

to be done in revision 04.

4.4 Simplification of in-order delivery requirement (technical)

In Clause 11.1.2, third paragraph, the text should be changed to read:

Frames shall be delivered in order when the error detection and recovery mechanisms specified by this clause are used.

The change is accepted.

to be done in revision 04.

4.5 Clarification of recovery requirements

Clause 11.2.1 considers recovery mechanisms for all classes of service.

Clause 11.2.2 extends the recovery capabilities for acknowledged services.

At present, the last sentence of 11.2.1 confuses the issue and should be deleted. In addition, a new first sentence should be added to the beginning of clause 11.2.2.

Acknowledged classes of service provide the following additional error detection mechanisms.

The change is accepted.

to be done in revision 04.

4.6 Definition of inbound and outbound

In section 11.2.2, the words inbound and outbound are used. What do they mean?

By SCSI convention, outbound transfers are from the initiator to the target and inbound transfers are from the target to the initiator. To clarify this, a search will be made to see if the words are used often. If they are, a glossary entry will be made for these words. If they are not, then the words will be eliminated and the complete descriptive statement will be used.

The change is accepted.

to be done in revision 04.

4.7 Clarify use of REC

The first three sentences of 11.3 should be rewritten as follows:

REC is transmitted by the initiator to periodically poll each outstanding exchange to determine if a SCSI task is progressing properly and if any sequences have been received incorrectly. The following optional error detection and recovery procedures are described for acknowledged services and for Class 3 service.

The change is accepted.

to be done in revision 04.