

Date: September 5, 2003

To: T10 Technical Committee

From: Keith Holt (LSI Logic)

Subject: Straw Horse Proposal for T10/03-111r1, End-to-End Data Protection Proposal

1. Change History

Revision	Date	Section	Description of Change
0	09/05/03	All	Initial revision.

2. Introduction

A data integrity study group meeting was held on August 19-20, 2003. The joint proposal from Agilent Technologies, Intel Corp., LSI Logic Corp. and Seagate Technology, 03-111r0, End-to-End Data Protection Proposal, was among the topics discussed. The proposal had been previously discussed in T10 CAP meetings and T10 data integrity teleconferences. General feedback has been that the proposal is too complex, and has too many options. Various changes to the proposal have been discussed, but there have been no revisions to the joint proposal thus far. As documented in the data integrity study group minutes, several straw polls were taken regarding proposed changes. The straw polls were open to all meeting attendees. This document is intended to be a straw horse proposal for 03-111r1, with the proposed changes limited to those discussed in the data integrity study group based on the straw poll results documented in the meeting minutes and from notes taken during the meeting by the author of this document.

In many ways, the proposed 03-111r1 could be viewed as an extension of 03-176r4. The most notable difference is that 03-111r1 has a provision to allow the Reference Tag base to be specified in a 32-byte CDB, while 03-176r4 requires the Reference Tag to be the lower 4 bytes of the LBA. The other notable difference is that the two proposals share common data integrity field elements, but utilize a different format for placement of those elements within the protection block. This is due in part to the fact that the two-byte "application tag" found in 03-176r4 is covered by the CRC calculation, while the two-byte "meta tag" in 03-111r0 is not covered by the CRC calculation.

3. References

- [1] T10/03-111r0, End-to-End Data Protection Proposal
- [2] T10/03-281r0, Data Integrity Study Group Minutes, August 19-20, 2003
- [3] T10/03-176r3, Simplified End-to-End Data Protection Proposal
- [4] T10/03-176r4, Simplified End-to-End Data Protection Proposal

4. Study Group Meeting Consensus Changes to 03-111r0

4.1. Meta Tag Validation

Discussion: There is some ambiguity in 03-111r0 regarding whether the Meta Tag is owned by the initiator or the target. The anticipated usage model is that ownership is SCSI domain dependent. The META ECHO flag provides a mechanism to allow the target to check and return tag values specified by the initiator, while storing target-provided values on media. On writes, the target checks the data stream from the initiator against the Meta Tag specified in the CDB. The target has the option of storing a target-provided value on media instead of that provided by the initiator. On reads, the target ensures that the data stream contains the value specified by the initiator in the CDB regardless of what was stored on media.

In order to reduce complexity and eliminate ambiguity, it was decided to designate either the initiator or target as the owner, with the owner having sole responsibility for validating the Meta Tag. There are still some unresolved issues such as designation of reserved values that are independent of ownership. If the target owns the tag, there is also a question of what, if any, requirements exist for initiators on write operation tag contents.

Description of Change: Meta Tag validation will be the sole responsibility of the owner, with the owner established via a Mode Page.

Impact to 03-111r0: Since the Meta Tag owner is solely responsible for validating the Meta Tag, there is no need for the initiator to specify the Meta Tag in the CDB. It also means that the Meta Tag Primary, Alternate and Legacy Masks can be removed from the Data Integrity Mode Page. Ownership of the Meta Tag will be established via a Mode Page.

4.2. Data Guard

Discussion: 03-111r0 has a 16-bit data guard field. There are two data guard formats: CRC and checksum based. The usage model cited for the checksum form uses checksum guards to form the basis of a non-standard guard field. There was considerable discussion regarding the merits of checksum vs. CRC protection schemes, and implementation complexity and speed of software algorithms. In the end, consensus was to restrict the Data Guard to a single form, CRC.

Description of Change: The only data guard form will be a 16-bit CRC, polynomial 0x18bb7, zero seed.

Impact to 03-111r0: Remove the checksum guard method.

4.3. READ/WRITE/VERIFY CDB Formats

Discussion: 03-111r0 proposes the addition of new 16-byte and 32-byte CDB formats. The 32-byte CDB format is required to address devices with > 2 TB capacity if the Reference Tag base is not locked to the LBA. The proposed 16-byte formats have a multi-format 8-byte LBA/Reference tag field that can be used with certain restrictions. The format allows the usage of 16-byte CDBs if the Reference Tag base for a particular I/O request or logical unit is indeed the lower four bytes of the LBA. In this case, there is no need to specify the Reference Tag base in the CDB, so the field is viewed as an 8-byte LBA and an implicit Reference Tag base. A second format allows usage of 16-byte CDBs if the I/O request or logical unit could be addressed with a 4-byte LBA. The other 4 bytes are used for Reference Tag base. The proposed 16-byte CDBs also have a 2-byte Meta Tag field, and a 2-byte Transfer Length field. (Note that existing Read/Write (16) CDBs have a 4-byte Transfer Length field.)

With the decision to make Meta Tag validation the sole responsibility of the owner, there is no need for a Meta Tag field in the CDB. (Refer to Section 4.1 of this document for details.) With the Meta Tag field removed from the CDB, the only reason to keep the new 16-byte form would be to allow a 4-byte LBA with a 4-byte Reference Tag base. Consensus was that the 32-byte form with an 8-byte LBA and a 4-byte Reference Tag base would be sufficient to meet this need. For I/O requests that have the Reference tag locked to the LBA, existing CDB formats can be used as long as per I/O controls can be provided to enable/disable checking on a per I/O basis.

Description of Change: Eliminate the new 16-byte CDBs. Add protection enable/disable control fields to existing READ/WRITE/VERIFY (10/12/16) CDBs as described in 03-176r3.

Impact to 03-111r0: Eliminate the new 16-byte CDB forms. Add protection enable/disable control fields to existing READ/WRITE/VERIFY (10/12/16) CDBs as described in 03-176r3. Remove the Meta Tag and Meta Tag Mask fields from the 32-byte CDB. Change the protection control fields in the 32-byte CDB to match those described for READ/WRITE/VERIFY (10/12/16) CDBs in 03-176r3.

4.4. Data Integrity Mode Page

4.4.1. STOR_DIF and DI_AVAIL Fields

Discussion: The specification of whether or not to format with DIF fields can be done via the FORMAT UNIT command. An indication of whether a device is DIF capable and DIF formatted if preferably returned in Inquiry data as opposed to a mode page. This approach eliminates the need for the STOR_DIF field.

Description of Change: Remove the STOR_DIF field from the Data Integrity Mode Page and provide control and status information via FORMAT UNIT and INQUIRY commands.

Impact to 03-111r0: Self-explanatory on the mode page. Add two-bit data protection support field to INQUIRY command as per 03-176r3. Add protection control field to FORMAT UNIT command as per 03-176r3.

4.4.2. Vendor Specific Field

Discussion: No precedent for designating a field as "vendor specific."

Description of Change: Remove the Vendor Specific field from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.3. LBA_40b Field

Discussion: Essentially a reserved bit for a non-existent 40-bit LBA proposal.

Description of Change: Remove the LBA_40b field from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.4. EXCL_Bytes Field

Discussion: Eliminate exclusion bytes in the interest of reducing complexity.

Description of Change: Remove EXCL_Bytes field from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.5. META ECHO Field

Discussion: Replace with Meta Tag ownership control as per Section 4.1 of this document.

Description of Change: Replace META ECHO with Meta Tag Ownership field.

Impact to 03-111r0: Define requirements for Meta Tag Ownership field.

4.4.6. REF METHOD Field

Discussion: This field implicitly allows a target to require LBA locked Reference Tags by virtue of providing a non-changeable option. Given the changes described in Section 4.2, a target can now implicitly require this mode of operation if it doesn't support the 32-byte CDB form. It was left unresolved as to whether or not the target should indicate this other than by rejecting 32-byte CDBs.

Description of Change: Remove the REF METHOD field from the Data Integrity Mode Page.

Impact to 03-111r0: Remove the REF METHOD field from the Data Integrity Mode Page. Open issue regarding mechanism to indicate support for 32-byte CDBs.

4.4.7. GUARD METHOD Field

Discussion: This field is no longer required since there is only a single guard method as per Section 4.2 of this document.

Description of Change: Remove GUARD METHOD field from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.8. DIF Stacking Fields

Discussion: 03-111r0 provides controls for DIF stacking. Consensus is that this can be done at the discretion of the target, outside the scope of the standard.

Description of Change: Remove DIF stacking controls from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.9. Verification Control Fields

Discussion: The Primary, Alternate and Legacy Verification Controls are no longer required since the CHK_OPTN controls in the CDB will be replaced with explicit per I/O controls as described in Section 4.2 of this document.

Description of Change: Remove Primary, Alternate and Legacy Verification controls from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.10. Legacy Write Control Field

Discussion: Study group consensus is that initiator control of legacy operations was not required, but the target should minimally indicate how it would treat legacy Read/Write operations.

Description of Change: Remove the Legacy Write Controls from the Data Integrity Mode Page. Replace with legacy operation status/support bits.

Impact to 03-111r0: Remove the Legacy Write Control field from the Data Integrity Mode Page. Final resolution pending on definition of legacy status field.

4.4.11. META TAG MASK Fields

Discussion: The META TAG MASK Primary, Alternate and Legacy fields are provided to limit checking of the META TAG value. Since META TAG checking is now the responsibility of the tag owner as per Section 4.1 of this document, there is no need to provide mask fields.

Description of Change: Remove the META TAG MASK fields from the Data Integrity Mode Page.

Impact to 03-111r0: Self-explanatory.

4.4.12. META TAG DEFAULT Field

Discussion: The FORMAT UNIT command allows an initialization pattern to be specified. This mechanism can be used to specify how the Meta Tag fields should be initialized during a format operation. With this approach, this field is not need to specify an initialization pattern. Behavior regarding legacy write operations is an open issue.

Description of Change: None at this time.

Impact to 03-111r0: None at this time.

4.5. INQUIRY Command

Discussion: An indication of whether a device is DIF capable and DIF formatted if preferably returned in Inquiry data as opposed to a mode page. See Section 4.4.1 for more information.

Description of Change: Add two-bit data protection support field to INQUIRY command as per 03-176r3.

Impact to 03-111r0: Self-explanatory.

4.6. FORMAT UNIT Command

Discussion: Control of device format is preferably done via the FORMAT UNIT command as opposed to a mode page. See Section 4.4.1 for more information.

Description of Change: Add protection control field to FORMAT UNIT command as per 03-176r3.

Impact to 03-111r0: Self-explanatory.

4.7. Data Integrity Field Format

Discussion: The layout of the fields within the Data Integrity Field has been discussed in prior T10 CAP meetings and in T10 data integrity teleconferences. The topic was not discussed at the study group meeting, however. As such, there is no consensus to change the layout from that shown in 03-111r0.

Description of Change: No changes.

Impact to 03-111r0: None. The DIF field format for an n-byte user data block remains as shown below. The CRC calculation covers only the user data block.

Table 1 – Protected Data Block Format

Byte/Bit	7	6	5	4	3	2	1	0
0	User Data Block							
n								
n+1	(MSB) Reference Tag (LSB)							
n+2								
n+3								
n+4								
n+5	(MSB) Meta Tag (LSB)							
n+6								
n+7	(MSB) CRC (LSB)							
n+8								