

Date: October 21, 2003

To: T10 Technical Committee

From: Keith Holt (LSI Logic)

Subject: End-to-End Data Protection Interoperability

References

- [1] T10/03-176r8, Simplified End-to-End Data Protection Proposal
- [2] T10/03-309r2, SBC Default Data Protection Block
- [3] T10/03-307r2, SBC 32 Byte Commands for SBC 32 Byte Commands for End-to-End Data Protection

Introduction

The proposed 32-byte commands for end-to-end data protection allow the application client to arbitrarily specify the INITIAL DATA BLOCK REFERENCE TAG on a per I/O basis (see Reference [3]) using algorithms that are unknown to the device server. In theory, each application client is free to choose any algorithm. Realistically, this can be done only in a homogeneous system where the algorithm is controlled by a single entity such as the file system or the Host Bus Adapter. In order to ensure interoperability, heterogeneous systems require that all application clients in a given SCSI domain use the same algorithm.

Interoperability with Legacy Hosts

The system example in Figure 1 shows a system with a mix of E2E capable and legacy hosts. In this example, all logical units in SCSI domain A are assumed to be formatted with E2E protection. If both host systems wish to access a common set of logical units in SCSI domain A, the device server has two basic choices on how to handle legacy writes:

1. Reject legacy writes.
2. Write a default pattern.

If the device server chooses to permit access by legacy hosts by writing a default pattern, the device server again has two basic choices:

1. Write a default pattern that indicates that the block was written by a legacy write.
2. Write a default pattern using the same algorithms as the E2E capable hosts.

If the device server is an intermediate device as shown in the system example in Figure 1, it is desirable to write a default pattern using the same algorithms as the E2E capable hosts. This provides E2E protection on the internal data paths of the intermediate device. In addition, this protection can be extended to the second SCSI domain, B, which the storage media resides. If the E2E capable application client uses LBA-locked REFERENCE TAGS, the device server can ensure interoperability with legacy hosts by inserting valid protection information with LBA-locked REFERENCE TAGS into the data stream received from legacy hosts. This model requires that the APP_TAG_OWN bit be set to zero for the logical unit, assigning ownership of the APPLICATION TAG to the device server.

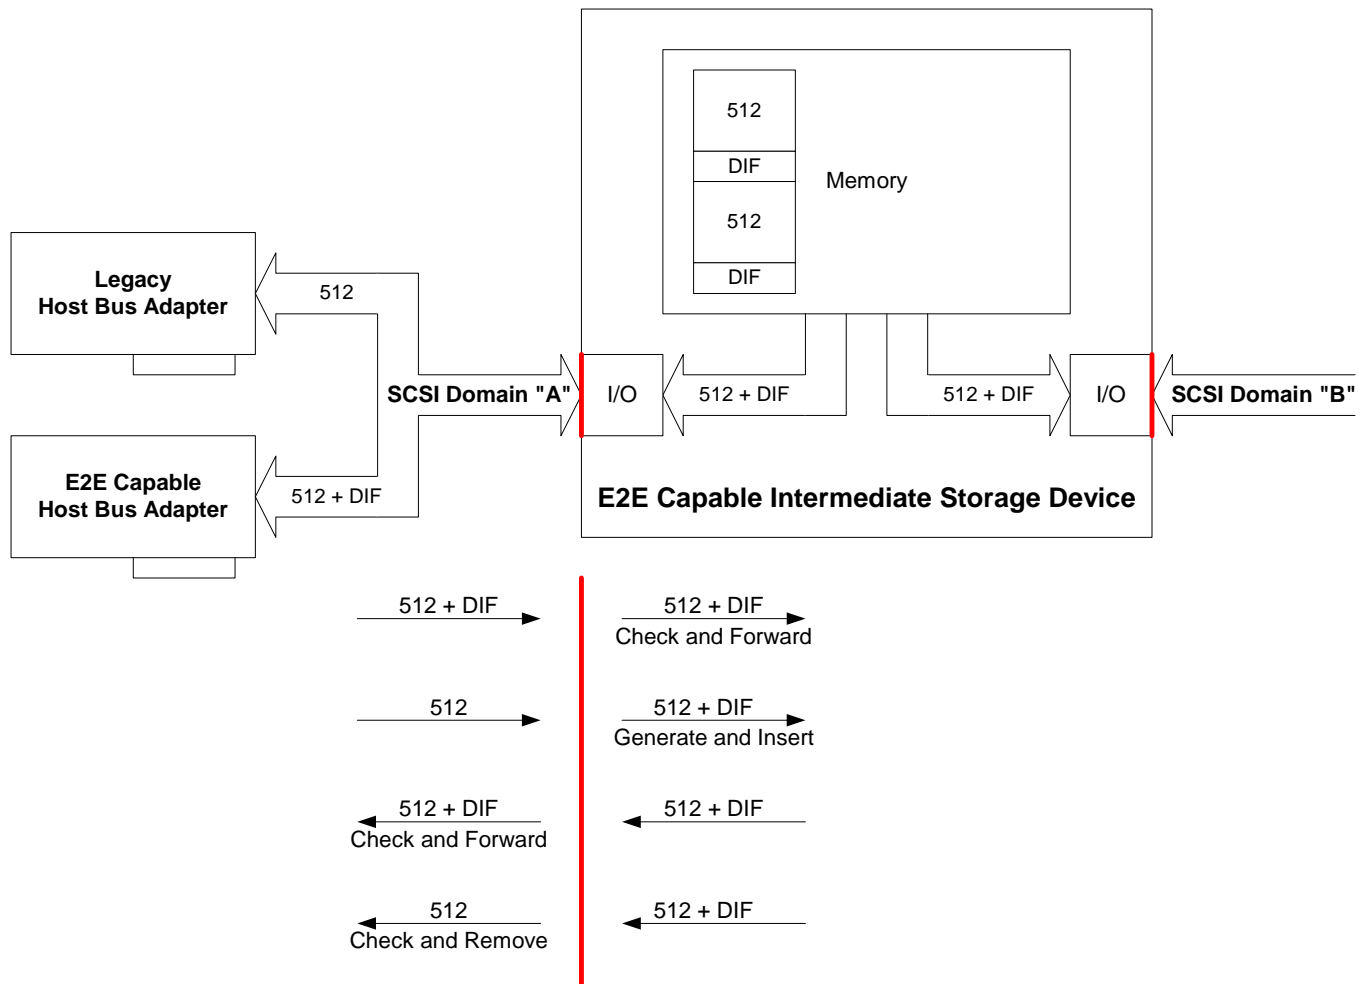


Figure 1 – Interoperability Model for System with mix of E2E Capable and Legacy Hosts

SPC-3 Additions

The E2E capable/legacy host interoperability model described earlier requires E2E application clients to use LBA-locked REFERENCE TAGS. The following changes are proposed for SPC-3.

4.1.2 Protection Information VPD page

Byte 4 bit 3: VAR_REF

If the variable reference tag (VAR_REF) bit is set to one, the device server supports variable REFERENCE TAGS via the 32-byte commands for end-to-end data protection.

Byte 4 bit 4: LCK_REF

If the locked reference tag (LCK_REF) is set to one, the device server appends LBA-locked REFERENCE TAGS to legacy write commands.