

External Path Protection Discussion

By
Curtis E. Stevens

18-Oct-2007





Agenda

- **SCSI Protection Information Overview**
- **SCSI Protection Information Usage**
- **SCSI Protection Information Usage Model**
- **ATA External Path Protection**
- **Development Questions**




SCSI Protection Information Overview



Sector Format Overview

■ Data Layout



- Data - User Data
- BLK_GRD –16-bit CRC on the user data (does not guard the 8 bytes of protection information)
- APP_TAG – 16 bits, application client specific, may be adjusted in type 2 protection information
- REF_TAG – 32 bits, depends on the protection information (PI) type



Type 1 Protection

- **Application – Standardized locality and CRC checking in systems where the application client communicates with a single drive or soft RAID**
- **BLK_GRD – 16 bit CRC on the user data**
- **REF_TAG – low order 32 bits of the LBA**
- **APP_TAG – application client specific information**
- **Protection only available for 6-, 10-, 12-, and 16-byte commands**
 - 32-byte commands are aborted
 - Protection information for 32-byte commands is type 2 only



Type 2 Protection

- **Application – Standardized locality and CRC checking in hardware RAID systems that receive an LBA from the host and then pass the data through multiple target/initiators. In this case, the REF-TAG retains its original value and is not necessarily related to the LBA.**
- **BLK_GRD – 16 bit CRC on the user data**
- **REF_TAG, APP_TAG, APP_TAG MASK provided in CDB**
 - REF_TAG – may NOT be low order 32 bits of LBA on destination target device
 - APP_TAG – application client specific information
 - APP_TAG MASK – may further qualify APP_TAG data
 - Protection only available to 32 byte commands
 - 6-, 10-, 12-, and 16-byte commands requesting type 2 protection information shall be aborted



Type 3 Protection

- **Application – Standardized CRC checking in systems where the application client provides additional protection in an application client specific manner**
 - Applies to systems where there is a value proposition for only checking BLK_GRD in the device.
 - Allows intermediary target/initiator devices to remap REF_TAG and APP_TAG as command moves through large system to adjust for different views of configuration
 - Provides a way for the host to do 48-bit locality checking
 - Provides a way for the host to do non-standard locality checking
- **BLK_GRD – 16 bit CRC on the user data**
- **REF_TAG and APP_TAG – Provided by application client and not checked by device**



SCSI Protection Information Usage



Discovery

- **Standard INQUIRY data – the PROTECT bit**
 - Informs the application client that the device is capable of supporting the Protect Information Model

- **Extended Inquiry VPD page**
 - SPT field – Indicates the protection types supported by the device
 - Only 4 options: None, Type 1, Type 1 and Type 2, Type 1 and Type 3.
 - GRD_CHK, APP_CHK, REF_CHK – Indicate which fields the device is capable of checking

- **READ CAPACITY (16)**
 - The PROT_EN bit indicates that the device has been formatted with protection
 - The P_TYPE field indicates the protection type the target is formatted with



Setting up the Device

■ FORMAT UNIT

- **FMTPINFO** – Enables/Disables protection
 - Since SBC-3 limits the combination of protection types the device may report, only 1 bit is needed to turn it on.
- **RTO_REQ** – The Reference Tag Owner distinguishes between Type 1 protection and other types
 - If the device owns the reference tag, the host shall supply a correct one and the device can check or generate based on the LBA.
 - If the application client owns the reference tag, the device shall never change the REF_TAG field and may check it for Type 2
- The device shall write the Protection Information as **FFFF_FFFF_FFFF_FFFFh** during the format process.
 - This initializes the protection information to the escape sequence for all protection types.

■ Control Mode Page – The ATO bit

- When the host performs a User Data only transfer, the ATO bit specifies that the Protection Information be the Escape Sequence or Valid protection information.



Escape Sequence

- **Type 1 and Type 2 protection**
 - when the APP_TAG field is FFFFh, the device shall not check the BLK_GRD or REF_TAG fields

- **Type 3 protection**
 - when the APP_TAG and REF_TAG are FFFFh and FFFF_FFFFh respectively, the device shall not check the BLK_GRD field



APP_TAG

- **This field is normally host vendor specific information and is simply stored by the device with the other protection information fields**
- **When the application client provides a read/write command that does not transfer protection information what does the device do?**
 - **The ATO (application tag owner) bit in the Control mode page determines the behavior**
 - **When the application client owns this field, the device shall insert FFFFh in APP_TAG field. This has the effect of disabling checking for BLK_GRD and REF_TAG**
 - **When the device owns this field, the device may insert a vendor specific value. This has the effect of allowing the device to place a value other than FFFFh and provide valid BLK_GRD and REF_TAG fields where appropriate.**



Protected Media Access

- **ORWRITE, WRITE, READ, VERIFY, WRITE AND VERIFY**
 - **Fieldnames - ORPROTECT, WRPROTECT, RDPROTECT, VRPROTECT**

- **A value of zero invokes a legacy operation**
 - **Only user data is transferred at the interface, no protection information is transferred**
 - **Protection Information is generated or stripped as necessary**

- **A value other than zero indicates that**
 - **Protection information is transferred, if the target is formatted with protection information**
 - **The type of checking that the target will perform on the protection information**

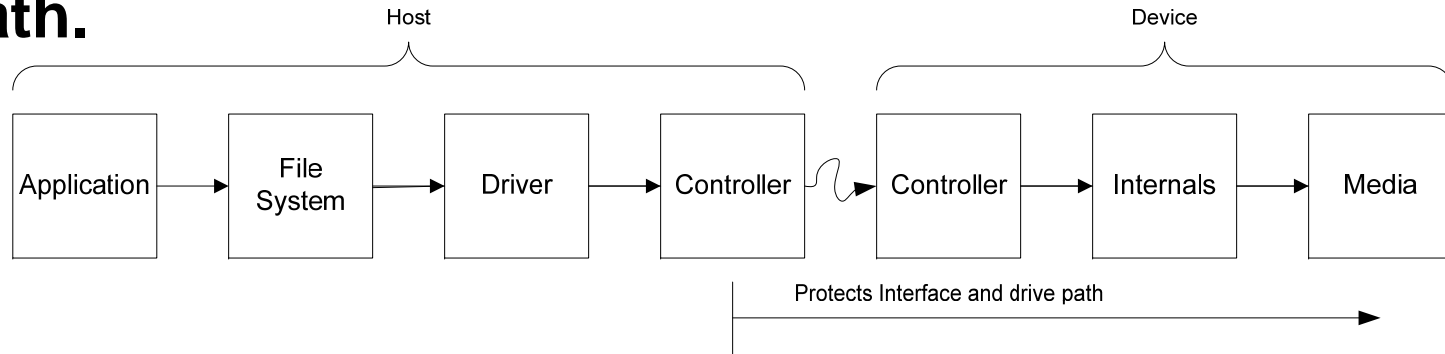


Protection Information Usage Models



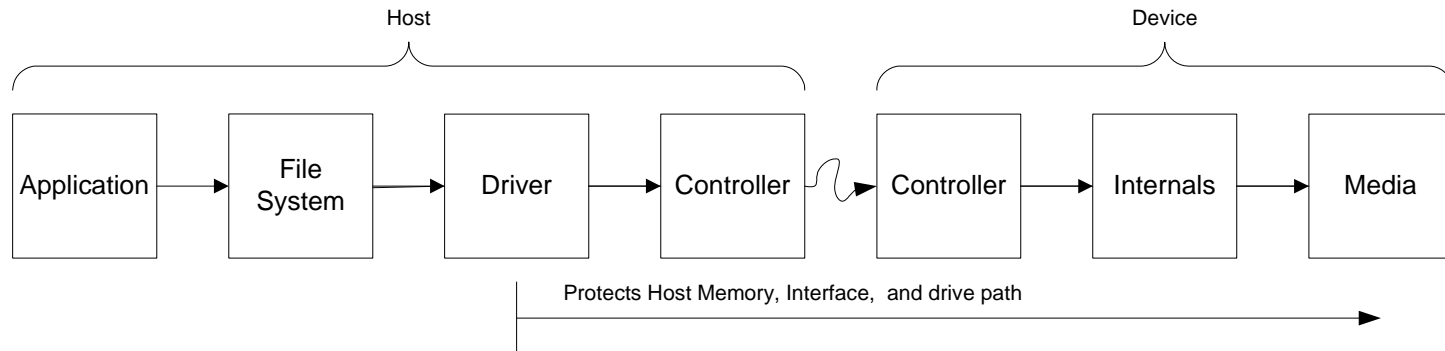
Sample Usage Models

- Protects data from the controller through the drive path.



- Protects host memory while the data is controlled by the driver

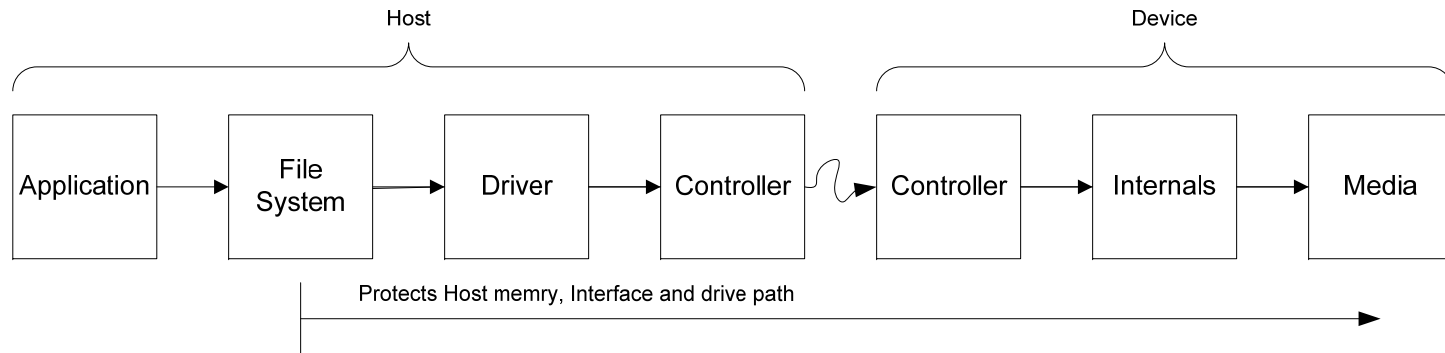
- Remains transparent to the rest of the system



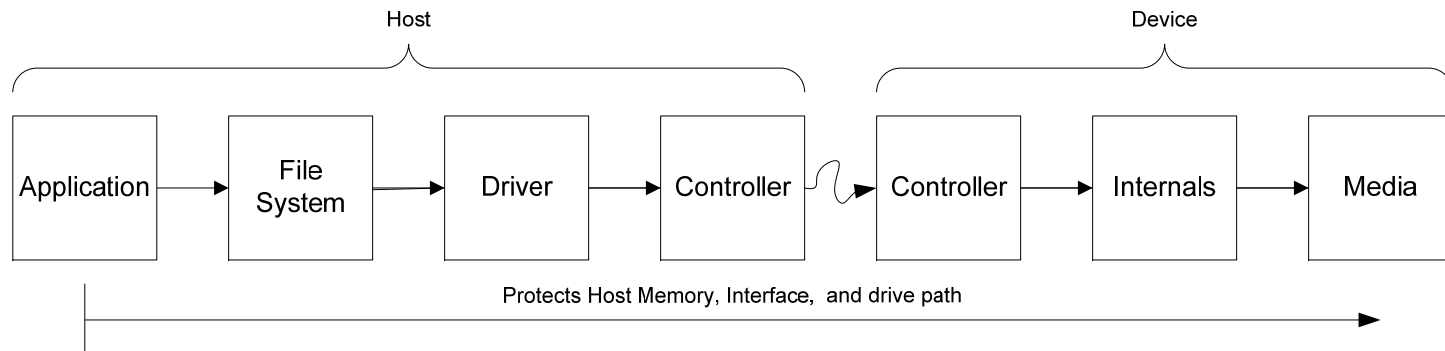


Sample Usage Models

- Protects host memory while the data is controlled by the filesystem and driver
 - Remains transparent to applications



- Provides full system round-trip data protection





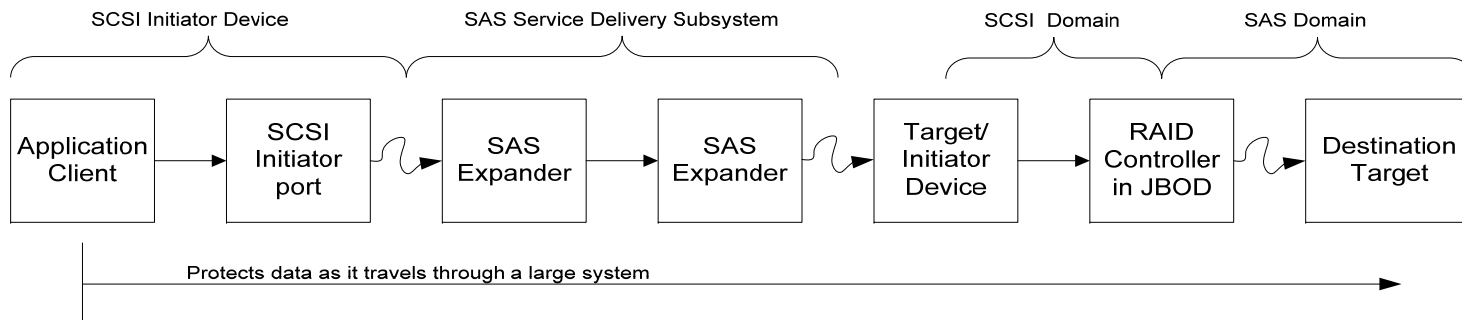
Differentiating Type 1/3 and Type 2

■ Type 1 and Type 3 protection

- REF_TAG may be changed in the PI data by target/initiator devices (e.g., RAID controllers) as the data travels through a system
 - For Type 1, the destination target may compare the PI data to the LBA in the CDB
 - For Type 3, the REF_TAG is not checked by the destination target

■ Type 2 protection

- REF_TAG in the CDB remains the same from application client to destination target as the data travels through a system



ATA External Path Protection



Foundational Principles

- **Assumes no FORMAT UNIT command**
- **Assumes devices are pre-formatted with either valid protection information or the escape sequence**
- **Read and write commands will not be modified**
- **All changes in protection field transfer, checking and generation will be “modal”**



Summary of SCSI Protection Information

- **Transfer of protection information may be changed on a command by command basis**
- **BLK_GRD and REF_TAG checking may be changed on a command by command basis**
- **APP_TAG is a field only useful to the application client**
 - **If the protection information is not transferred then ATO may be used to place the device in a mode where an escape sequence is inserted, or valid protection information is inserted.**



ATA w/Type 1 and Type 3 protection

- **Provide SET FEATURES for**
 - **Enable/Disable Protection Information Transfer (following user data)**
 - **Enable/disable Escape Sequence – provides functionality of ATO bit**
 - **Escape Sequence type – differentiates between type 3 and other types**
 - **BLK_GRD and REF_TAG checking enable/disable**
 - **Disabling REF_TAG checking is the same as Type 3 operation**
- **SCT Write Same could be used to force valid protection information onto the media.**



Escape Sequence

- Follows the same requirements as SCSI
- If **BLK_GRD** or **REF_TAG** checking is enabled and the escape sequence is encountered in the protection information, then the **BLK_GRD** and **REF_TAG** shall not be checked by the device.
- If Escape Sequence Type is set to **APP_TAG=FFFFh** then whenever the **APP_TAG=FFFFh**, the protection information shall not be checked by the device.
- If Escape Sequence Type is set to **APP_TAG=FFFFh** and **REF_TAG=FFFF_FFFFh** then whenever the **APP_TAG=FFFFh** and **REF_TAG=FFFF_FFFFh**, the protection information shall not be checked by the device.



Open Items



ATA w/Type 2 Protection

- In SCSI, when the media is formatted with Type 2, the 32 byte media access CDB's work. The other ones (6, 10, 12, and 16) only work in legacy mode
 - This is because the REF_TAG field is provided as a part of the CDB.
- Still studying usage model for mapping into T13 commands
 - *Is it reasonable to program in an offset that applies to the entire device until changed?*
- How do we deal with APP_TAG and APP_TAG MASK?
 - *Since SET FEATURES has 32 bits available, these could be set with SET FEATURES as well.*