An Oracle White Paper
September 2013

Data Integrity Validation Feature of StorageTek T10000 Tape Drives

ORACLE

Executive Overview

It is critical to ensure that stored data has been recorded accurately. Oracle's Data Integrity Validation, a feature of Oracle’s StorageTek T10000C and StorageTek T10000D tape drives, takes this one step further by validating CRC checksums generated at the host.

Introduction

In today’s business environment almost all enterprise data is stored digitally. This valuable information must be correct with accuracy guaranteed for the life of the data. Most storage devices do an excellent job protecting data once it is at rest.

In tape drives, data is protected with read-after-write verification as it is written, and error correction code (ECC) is added to ensure data recovery once it is on the medium. In addition, a typical tape drive adds cyclic redundancy code (CRC) protection, as soon as a record is received. This ensures the record does not get corrupted while moving between internal memories.

Unfortunately, these technologies do not protect data that is being moved outside the storage device. As a result, there is a chance for data corruption as it is migrated across the storage landscape.

Oracle knows how important your data is, and has developed Data Integrity Validation for StorageTek T10000 tape drives, to protect enterprise data from end to end.

The Problem with Moving Valuable Data

The movement of data typically occurs across many different pieces of hardware (switches, back planes, buses, adapters, and various memory buffers), with each component claiming data cannot be corrupted while under its control. Various storage protection technologies like disk RAID 6 and tape drive internal data checks are correctly advertised as features that protect your data. Technologies like error correction codes (ECCs), cyclic redundancy codes (CRC) and other hash-based data check methods are used as safeguards, ensuring user data is never corrupted. Figure 1 is a selective illustration of data movement from disk to tape.

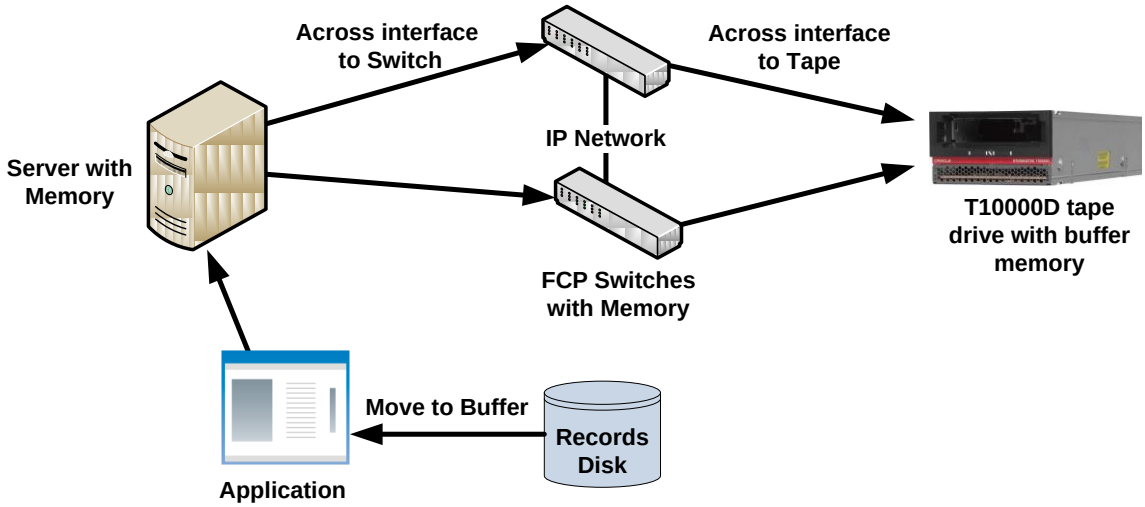

Figure 1: Records moving through a typical environment.

1. The records are moved from disk memory to server buffer memory via the application. (Disk drive checks integrity.)
2. The records are framed and moved from buffer memory to an interface transport (Fibre Channel host bus adapter or HBA) to Fibre Channel Protocol (FCP) switches. (Server checks integrity.)
3. The records are moved across the interface transport to buffer memory in a switch. (FCP interface ICs check integrity)
4. The records are moved from buffer memory in a switch to another interface transport (FCP switch checks integrity.)
5. The records are moved across the interface transport to buffer memory in a tape drive. (FCP ICs check integrity.)
6. The records are moved from buffer memory to tape. (Tape drive checks integrity.)

In this example the disk drive guarantees integrity in step 1. The server guarantees integrity in step 2. The FCP ICs guarantee integrity in step 3. The switch guarantees integrity in step 4. The FCP ICs guarantee integrity in step 5. The tape drive guarantees integrity in step 6.

Even in this basic example, six different technologies guarantee integrity within their domains but no single technology is in place to guarantee data integrity from end-to-end (steps 1- 6).

Oracle's Data Integrity Validation Feature

To solve this data movement problem, Oracle has introduced the capability to check end-to-end data integrity. Oracle's StorageTek T10000 tape drives are capable of checking the 32-bit CRC specified in ANSI X3.139. This is the same CRC used in the Fibre Channel Protocol and the Fiber Distributed Data Interface (FDDI) for optical transmissions. As a result, the generation polynomial is readily available. While this is a standard interface CRC, it is important to note that this check is performed outside the interface protocol. In addition, the drive can use and generate the iSCSI CRC32c (supported by Intel Xeon and Oracle's SPARC T4 chips) and use and generate the Reed Solomon CRC (not supported by Intel Xeon and SPARC T4 chips).

Use of this capability is straightforward. The application or file system generates a CRC on data at rest and appends the CRC to records sent to the StorageTek T10000 tape drive. If the drive is in Data Integrity Validation mode, the drive will check this CRC for every record it receives.

This CRC is stored with the record on tape, and appended to the record when it is returned to the application on a subsequent read. This solution ensures the tape drive checks end-to-end data integrity when a record is written, and allows the application to check end-to-end data integrity when the record is read at a later time.

Data Integrity Validation: CRC Implementation

To support Data Integrity Validation, the application, or driver, places the StorageTek T10000 tape drive in Data Integrity Validation mode by sending a specified SCSI mode select command. When the drive is in this mode, the last 32 bits of any record are treated as a CRC and used to check the integrity of each record. If the CRC check fails, a write error is reported to allow the application to resend the record. A bad record never will be written to tape. If the CRC is correct, that CRC is stored with the record on tape and checked every time the record is read. All of this is done with zero performance loss on the tape drive.

If a deferred write error has been reported to the application, the application can determine which record was in error by using one the following methods. Either the deferred error Residuals information or the Last Block Location field from the Read Position command can determine the record in error. The application only needs to do a Backspace Block command, to delete the trapped write data, then a Forespace Block command to reposition at End of Data. The recovery is completed when the application resends the previously failed record and the remainder of the data records.

If the drive is in Data Integrity Validation mode during a subsequent read, the CRC will be appended to the record. This allows the application or driver to perform its own data integrity checks, and ensure the data has not been corrupted months, or even years, later. The Intel CRC32c format allows very fast CRC processing and checking by the application. The user application, or driver, can use Data Integrity Validation as follows.

- Write with Data Integrity Validation and read with Data Integrity Validation
- Write with Data Integrity Validation and read in normal mode
- Write in normal mode and read in Data Integrity Validation mode (Note: in this case, the DIV-read CRC, which is generated by the drive on the fly, was not stored on tape.)

Data Integrity Validation: Verify Implementation

Another advantage of writing a tape in Data Integrity Validation mode is the ability of the tape drive to use the Verify command to check an individual record, one file, multiple files, or the entire tape, without having to send all the data to the application to verify the validity of that data. This can be done because the Data Integrity Validation CRC is recorded on the tape with each record, and the tape drive has the ability to verify each record with that CRC. Because the Data Integrity Validation CRC is only 32 bits, checking only the CRC saves valuable processing resources and time.

Summary

With vast amounts of critical data being stored digitally, it is essential that the content of the data remain unchanged during movement and when stored for long periods of time. For legal and preservation purposes, the fixity of this data must be verified. Oracle's Data Integrity Validation feature of the StorageTek T10000C and StorageTek T10000D tape drives allows applications, drivers, and file systems to use CRC checksums to ensure the integrity of important data is preserved.

Oracle's Data Integrity Validation Feature of
StorageTek T10000 Tape Drives
September 2013

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together