

# Reduce Data Corruption and Protect Your Investment

Build Architecture on Hitachi Virtual Storage Platform Family and  
the T10-PI With DIX Standard

By Hitachi Data Systems

May 2017

# Contents

Executive Summary .....	2
Silent Data Corruption.....	3
The T10-PI Standard.....	3
T10-PI With DIX Testing .....	4
Hitachi Virtual Storage Platform Family.....	5
Oracle Linux .....	6
Conclusion.....	6

## Executive Summary

Businesses must ensure that the data that is used to make strategic decisions is accurate. Data corruption reduces the accuracy of the information that is needed and can cause a loss of revenue due to inaccurate information. Data corruption is one of the sources of inaccurate information.

This paper focuses on data corruption and how the following computer system architecture using the T10-PI standard meets this challenge:

- Hitachi Virtual Storage Platform (VSP) G series and VSP F series, powered by Hitachi Storage Virtualization Operating System (SVOS).
- Oracle Linux with the Unbreakable Enterprise Kernel (UEK).
- QLogic host bus adapter (HBA).

This architecture reduces data corruption and protects your investment in today's complex data processing environment.

## Silent Data Corruption

With data flowing from the application layer to the storage array layer and back again, there is a possibility that the data may become corrupt. Undetected data errors, where there is no indication that the data has been corrupted, are known as silent data corruption.

Hitachi Virtual Storage Platform (VSP) G series and VSP F series systems, powered by Hitachi Storage Virtualization Operating System (SVOS), are designed and engineered with the ability to detect, recover and/or correct data corruptions. Nevertheless, data corruption can still occur at any level external to Hitachi's storage systems, such as the host, the host bus adapter (HBA) or the SAN.

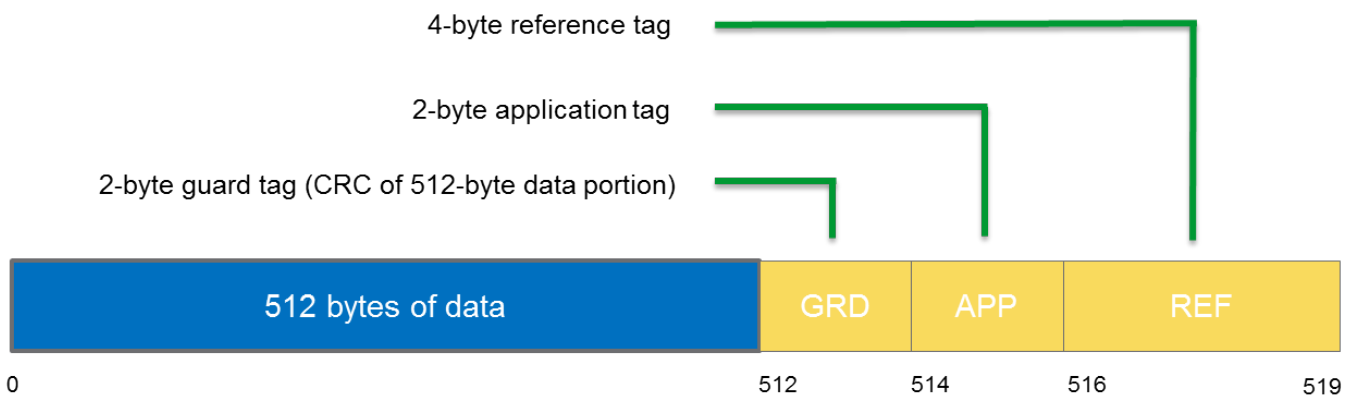
## The T10-PI Standard

The InterNational Committee on Information Technology Standards (INCITS, pronounced "insights") operates under the rules established by the [American National Standards Institute \(ANSI\)](#). T10, which is the technical committee of INCITS, has developed a process for finding and correcting errors in data that is flowing from the application layer to the storage layer. The T10 committee is also responsible for the creation of the SCSI interface architecture.

In their 2003 proposal, the Protection Information (PI) subcommittee of T10 added a section on data integrity, formally known as T10-DIF (Data Integrity Field), to the SCSI standards. In the T10-DIF specification, the SCSI block device protocol is extended by 8 bytes. The additional information is attached to each data block that is transmitted between the HBA and the storage array and serves as a checksum. The layout of the DIF is shown in Figure 1.

The Data Integrity Extension (DIX) improves things further by moving the checksum calculation up to the application for a full end-to-end data integrity check. HBA vendor QLogic implements Oracle's DIX specification by enabling the exchange of T10-PI between the Hitachi VSP storage array and the operating system (Oracle Linux with the Unbreakable Enterprise Kernel) and the application (Oracle Database). The data is verified as it travels through the I/O stack, where the extra 8 bytes in each sector leaving the HBA are used by the storage array firmware for integrity checking. In the 8 byte extended sector, the guard tag protects the data portion of the sector. The data in the reference tag is used to protect against out-of-order and misdirected write. The application tag is used by the operating system and the application or relational database management system (RDBMS) for error correction, when data is written to storage and then checked after it is read back.

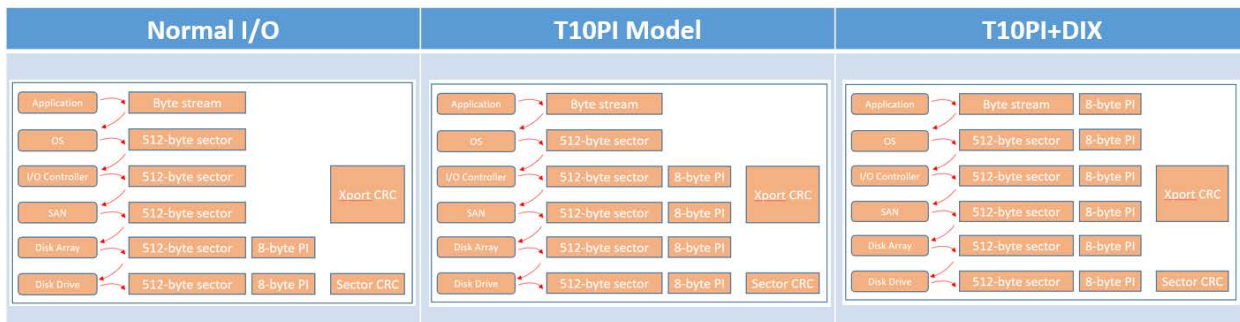
Figure 1.



Therefore, to achieve end-to-end protection, the operating system, HBA and storage array use this common format. Checksums created by Oracle’s ASMLib can be verified by the host adapter and storage array firmware. Any mismatches between the data and the protection information can be corrected before the corrupted data is used in your business applications.

Figure 2 shows (from left to right) normal I/O, what the I/O looks like with the T10-PI model implemented, and what the I/O looks like with the DIX extension implemented.

**Figure 2. Normal I/O in Comparison to T10-PI and T10-PI+DIX Models**



## T10-PI With DIX Testing

Hitachi storage systems are engineered to have no single points of failure. Our engineering and quality assurance teams conduct extensive quality assurance tests prior to each release.

Hitachi’s VSP G series and VSP F series DIX and T10-PI capable storage arrays have undergone rigorous end-to-end testing. Our T10 and DIX test environments consist of:

- Oracle Real Application Clusters (RAC) and single instance Oracle Database 11g Release 2 (11.2.0.4).
- Oracle Linux Release 6 Update 5 (x86\_64) with UEK Release 3 (UEK R3).
- T10 and DIX capable HBA (QLogic 16Gb/s - QLE2672 firmware 7.04.01).
- T10 and DIX capable Hitachi storage arrays.
- Online transaction processing (OLTP) application simulators such as Oracle’s SwingBench.

Hitachi’s VSP G series and VSP F series systems that are DIX and T10-PI capable have minimal impact on system performance. They provide the storage administrator the flexibility to enable or disable the T10-PI or DIX support in the storage array.

## Hitachi Virtual Storage Platform Family

All-flash VSP F series systems and hybrid flash VSP G series systems give you more choice in how you design your data center, while delivering industry-leading data management efficiencies.

From entry models to full-scale enterprise, the VSP G series models are hybrid flash arrays that deliver data faster to improve user experiences. They are the right choice for organizations that are making a seamless transition to flash, tier by tier. VSP F series models are all-flash arrays designed to become the foundation of your data center modernization needs.

The entire VSP family provides support for both block and file workloads. SAN and NAS protocols are supported, enabling organizations to serve a wide variety of internal and customer-facing applications from a single high-performance array. And for mainframe users, VSP G1500 and VSP F1500 models can be integrated with mainframe applications for an unprecedented level of performance and operational resiliency.

The VSP family also delivers global storage virtualization, enabling over 100 different types of arrays to be used as extended capacity. It allows administrators to use a single data management offering to control their storage infrastructure, simplifying day-to-day operations.

Built on legendary Hitachi reliability, VSP family offers complete system redundancy and is backed by the industry's only 100% data availability guarantee.

**Storage Virtualization Operating System v7** is the powerful engine at the heart of every VSP system. It includes a new set of capabilities engineered specifically for flash.

- **Adaptive data reduction:** software that reduces capacity requirements by two times, five times or more, which lowers acquisition and maintenance costs.
- **Quality of service:** controls that help deliver consistent application performance.
- **Direct connect to cloud:** functionality that enables unparalleled on-site capacity reduction with transparent movement of files via data migrator to cloud (DM2C).

For more information about Hitachi VSP family, please visit

<https://www.hds.com/en-us/products-solutions/storage/virtual-storage-platform-g-series.html>.

## Oracle Linux

The Oracle Linux operating system is engineered for open cloud infrastructure. It delivers a secure and easy to manage infrastructure environment for enterprise SaaS and PaaS workloads, as well as traditional enterprise applications. It is the foundation for Oracle's public cloud. Oracle Linux Support offers access to award-winning Oracle support resources and Linux support specialists, zero-downtime updates using Oracle Ksplice, additional management tools such as Oracle Enterprise Manager, and lifetime support, all at a low cost. Unlike many other commercial Linux distributions, Oracle Linux is easy to download, completely free to use, free to distribute, and free to update.



Oracle Linux comes with a choice of two kernels, the Unbreakable Enterprise Kernel (UEK), which is installed and enabled by default, and the Red Hat compatible kernel. The UEK tracks the latest Linux kernel releases, supplying more innovation than other commercial Linux kernels, while providing binary compatibility with applications certified to run on Red Hat Enterprise Linux. The UEK is open source (GPL) and is designed for enterprise workloads, such as Oracle Database.

Oracle Linux with UEK efficiently runs on systems with many cores, threads and NUMA nodes, is optimized for solid state drives, supports data integrity, and provides hardware fault management.

For more information about Oracle Linux with the Unbreakable Enterprise Kernel, please visit:

<http://www.oracle.com/us/technologies/linux/product/overview/index.html>

<http://www.oracle.com/technetwork/server-storage/linux/technologies/uek-overview-2043074.html>

## Conclusion

With many years of experience building and delivering enterprise-class storage systems, Hitachi's T10/DIX capable storage arrays meet the requirements that allow business to build highly reliable and complex compute environments. Hitachi VSP G series and VSP F series DIX or T10-PI capable storage platforms protect your valuable data. They minimize downtime caused by silent data corruption that may occur in data transmission, ensuring that your data is accessible and accurate.



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 95050-2639 USA  
[www.HDS.com](http://www.HDS.com) | [community.HDS.com](http://community.HDS.com)

Regional Contact Information  
**Americas:** +1 866 374 5822 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)