

White paper End-to-end Data Protection Using Oracle Linux with the ETERNUS DX S3 series and Emulex HBA

Greatly improving the reliability of the entire system



Content

Preface	2
1. Data Protection with T10 PI	3
■ 1.1 What are T10 PI and DIX?	3
■ 1.2 T10 PI Data Protection Field Format	3
■ 1.3 I/O Processing With T10 PI	3
■ 1.4 I/O Processing With DIX	4
2. Components that Enable End-to-end Data Integrity	4
3. Combination of Components and Functions that Achieves Data Integrity	5
4. Summary	5

Preface

As demand for data protection has grown in recent years, ensuring the integrity of the customer's data is a very important task. To carry out this task, data protection is not only required for the storage system but the entire system.

This White Paper explains how data protection can be implemented for the entire system using T10 Protection Information (T10 PI) that comes with the Oracle Linux with Unbreakable Enterprise Kernel (UEK), FUJITSU Storage ETERNUS DX S3 series, and Emulex HBA.

ORACLE
LINUX

ORACLE Diamond
Partner



EMULEX



FUJITSU Storage
ETERNUS DX S3 series

1. Data Protection with T10 PI

1.1 What are T10 PI and DIX?

Data protection is nothing new. Each system component has data protection functions such as memory parity detection, error detection with communication protocols, CRC error detection of the disk storage system, and redundancy. (Figure 1.1, I. Conventional data integrity functions)

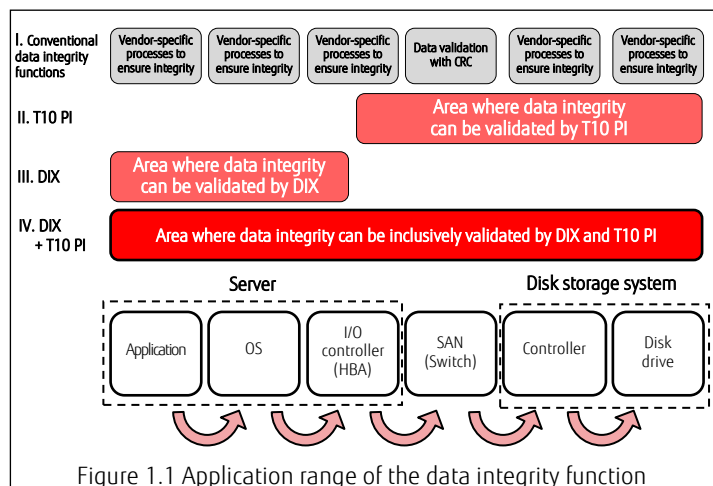
T10 PI defines the mechanism for data integrity protection that is standardized through communication paths and devices. End-to-end data integrity from applications to disk drives is implemented with the combination of T10 Protection Information (T10 PI) and Data Integrity Extensions (DIX).

T10 PI adds protection information to data in the I/O controller (HBA) layer, transfers the protection information to the devices that support T10 PI in the layers lower than the I/O controller (HBA), and validates data integrity, including data written to the disk drive. (Figure 1.1, II. T10 PI)

However, because protection information is added to data in the I/O controller (HBA) layer, data cannot be protected in the higher layers (applications, OS).

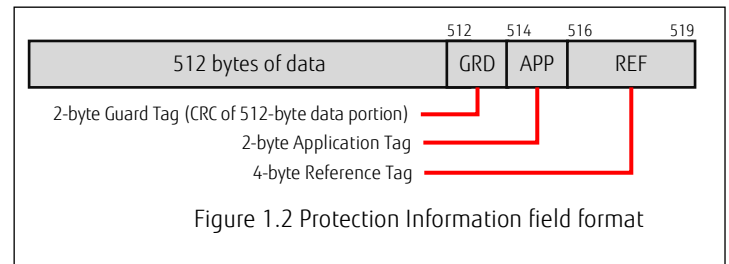
Data protection, including applications and OSs, is implemented with the Data Integrity Extensions (DIX) which enable the T10 protection information to be exchanged between the Linux kernel and the HBA. (Figure 1.1, III. DIX)

The combination of T10 PI and DIX provides complete end-to-end data protection. (Figure 1.1, IV. DIX + T10 PI)



1.2 T10 PI Data Protection Field Format

An 8-byte protection information that is added by T10 PI during I/O operations consists of Reference Tag, Application Tag, and Guard Tag as shown in Figure 1.2.

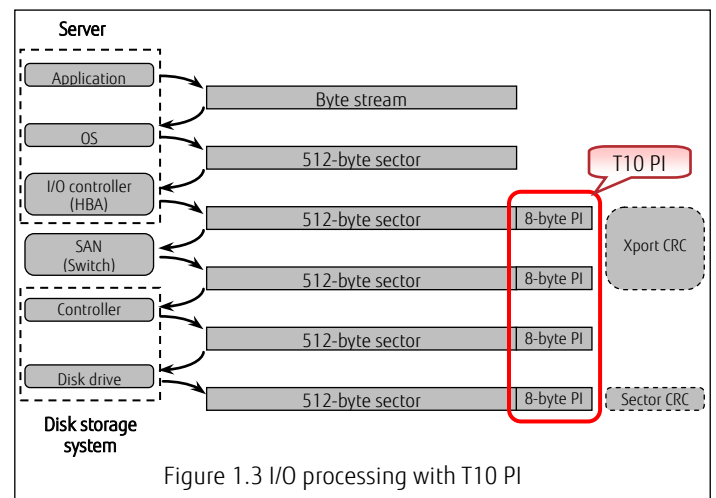


- Reference Tag: Detects data writing to incorrect blocks.
- Application Tag: Specific to applications. Application Tag defines the purpose of data.
- Guard Tag: Protects the data portion of the sector with CRC.

1.3 I/O Processing With T10 PI

T10 PI adds protection information (8-byte PI) to data in the I/O controller (HBA) and validates data integrity within the layers lower than the I/O controller.

Figure 1.3 indicates I/O processing with T10 PI.



- The application transfers data to be written to the OS.
- The OS (file system) transfers the data to be written to the I/O controller (HBA) as a logical block with 512 bytes per sector.
- The I/O controller generates and adds the protection information, and sends data with 520 bytes per sector to the SAN.
- The controller (disk storage system) validates the protection information and destination.
- The disk drive validates the protection information before writing the data.

■ 1.4 I/O Processing With DIX

DIX validates data integrity in applications and OSs.

Figure 1.4 indicates I/O processing with DIX.

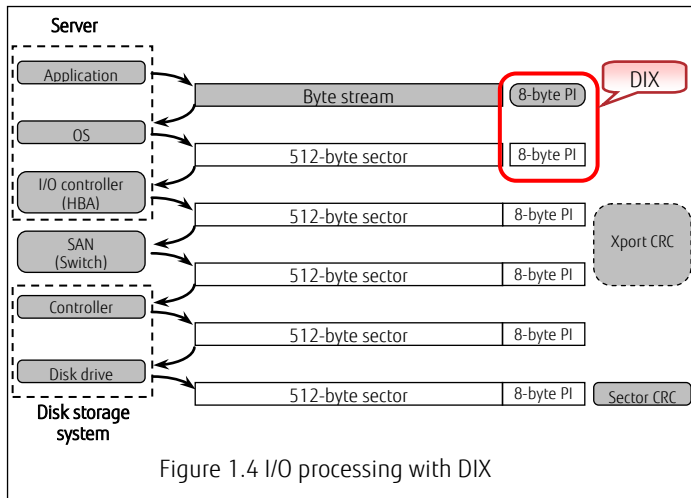


Figure 1.4 I/O processing with DIX

- Data integrity processing that is specific to applications generates protection information for data and transfers both the data to be written and the protection information to the OS as required.
- The OS transfers the data and the protection information to the I/O controller (HBA).
If the protection information is not generated, the OS automatically generates and adds the protection information to the data.
- The I/O controller merges the data and the protection information, and sends them as data with 520 bytes per sector to the SAN.
- The following processes are the same as those described in "■ 1.3 I/O Processing With T10 PI".

The combination of T10 PI and DIX enables end-to-end data integrity as described above.

2. Components that Enable End-to-end Data Integrity

This section explains the three components that enable end-to-end data integrity (OS, I/O controller (HBA), disk storage system).

- Oracle Linux with Unbreakable Enterprise Kernel (UEK)
- Emulex LightPulse® LPe16000 series of 16Gb Fibre Channel HBA
- FUJITSU Storage ETERNUS DX S3 series

■ Oracle Linux

Oracle Linux integrates latest innovations in the Linux mainline and provides excellent performance, high scalability, and reliability for enterprise systems and applications.

Oracle Linux provides the latest tools, and features that enable customers to innovate, collaborate, and create solutions across traditional, cloud-based, and virtual environments.

Oracle Linux with UEK efficiently runs on systems with many cores, threads and NUMA nodes.

Developed and tested to support demanding workloads, such as the Oracle Database, it is one of the most reliable, fast, and scalable OS for the enterprise implementations.

In addition to this, Unbreakable Enterprise Kernel contains many new features that are relevant to Linux running in the data center, including support for T10 PI and the Data Integrity Extensions.

Oracle Linux with the Unbreakable Enterprise Kernel, including the data integrity features, is provided under the GNU General Public License (GPL) and is available to anyone in both binary and source form.

Oracle Linux with the default Unbreakable Enterprise Kernel can be downloaded for free from edelivery.oracle.com/linux

Existing or new Oracle Linux support customers will receive full support for this kernel as part of their existing support subscriptions.

Bug fixes and security errata are delivered via ULN and announced via the el-errata mailing list.

■ Emulex LightPulse LPe16000 series Fibre Channel HBAs

Emulex LightPulse LPe16000 series of 16Gb Fibre Channel HBA implements the data integrity functions of T10 PI and DIX using BlockGuard technology.

■ Data Integrity Offload Support

- Emulex LightPulse Fibre Channel HBA family implements T10 PI as standard using the Data Integrity offload function.
- The integrity of user data that is transferred from applications to the SAN is maintained by T10 PI.
- Data Integrity offload operates in the UEK and in the Emulex LightPulse Fibre Channel HBA driver that resides in the kernel.

■ Emulex LightPulse LPe16000 series Fibre Channel HBA

- 16Gb/s Fibre Channel
- PCI Express 3.0
- Single-Port Host Bus Adapter (LPe16000)
- Dual-Port Host Bus Adapter (LPe16002)
- Supports hardware offload of T10 PI tasks
- The end-to-end data integrity function is offloaded to the hardware by using T10 PI and provides a 40% higher performance than a firmware implementation.

■ FUJITSU Storage ETERNUS DX S3 series

The hybrid and unified ETERNUS DX S3 series feature data protection and flexibility functions to meet highest business-centric requirements and provide excellent reliability.

The protection information generated by the OS is secured in the UEK and the HBA using T10 PI and DIX functions. It is then transferred to the disk storage system, ensuring data integrity until all data is written to disk.

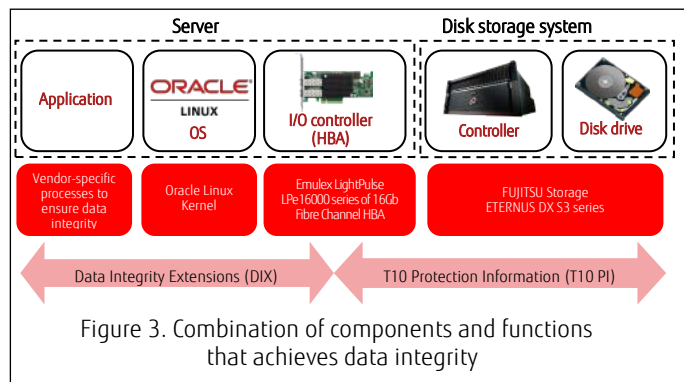
Besides T10 PI, ETERNUS DX series improves data reliability with the unique "Data Block Guard" function. This function adds a check code to each block of data being written to disk and validates data integrity at multiple check points. In the case of drive failures incorrect data will be detected, thus data integrity is validated and ensured throughout the disk storage system.

The current ETERNUS DX series supporting T10 PI lineup consists of:

- Entry class systems
ETERNUS DX100 S3, ETERNUS DX200 S3
- Midrange class systems
ETERNUS DX500 S3, ETERNUS DX600 S3
- Enterprise class systems
ETERNUS DX8000 S3 series
- All-flash Array
ETERNUS DX200F

3. Combination of Components and Functions that Achieves Data Integrity

The relationship between components and data integrity functions that achieves data integrity is shown in Figure 3.



- DIX validates data integrity from the application to the I/O controller (HBA).
- T10 PI validates data integrity from the I/O controller (HBA) to the disk drive.

4. Summary

The T10 PI and DIX functions enable data protection for the entire system including data write processes from applications to disk drives.

Data protection for the entire system can improve system reliability by protecting the system from damaged or inconsistent data to minimize downtime and reduce end-user maintenance costs.

For other specifications and functions of the components that are described in this document, refer to the following URLs:

- Oracle Linux
<http://www.oracle.com/us/technologies/linux/product/overview/index.html>
- Emulex LightPulse LPe16000 series Fibre Channel HBA
<http://www.emulex.com/products/fibre-channel-hbas/>
- FUJITSU Storage ETERNUS DX S3 series
<http://www.fujitsu.com/global/products/computing/storage/>