

White Paper

FUJITSU Storage ETERNUS DX S4/S3 series 512e HDDs: Technology, Performance and Configurations

The new technology of physical 4k sectors for HDDs increases storage density, performance, and the capability to correct errors. The first of these HDDs is now available for the ETERNUS DX S4/S3 series. Suitable versions of OS and applications should be used for optimal performance.



Contents

Abstract.....	2
Technology and examples.....	2
Performance	3
Examples of 512e aware Operating Systems	4
Non-512e-compliant Operating Systems.....	5
About applications	5
About middle wares	5
Others	5

Abstract

The ETERNUS DX S4/S3 series supports Advanced Format (4k sectors) HDDs. This document describes how to avoid the potential performance impacts when the Advanced Format HDDs are used in the ETERNUS DX S4/S3 series.

Technology and examples

For a long time HDDs have been physically structured in sectors that have 512 bytes of user data and a certain amount of overhead (including an “Error Correction Code”, or ECC). However, this sector format has in recent years reached its limits. It is no longer compatible with today’s requirements for storage density, performance, and the capability to correct errors. To overcome these limits the format for physical sectors of 4096 bytes was defined in 2009 as an industry standard (“Advanced Format”, 4k sectors).



Advanced Format

Advanced Format is defined by the International Disk Drive Equipment and Materials Association (IDEMA). This technology changes the user data size of a sector from 512 bytes to 4096 bytes as shown the figure-1 below. It provides for more efficient use of the drive media.

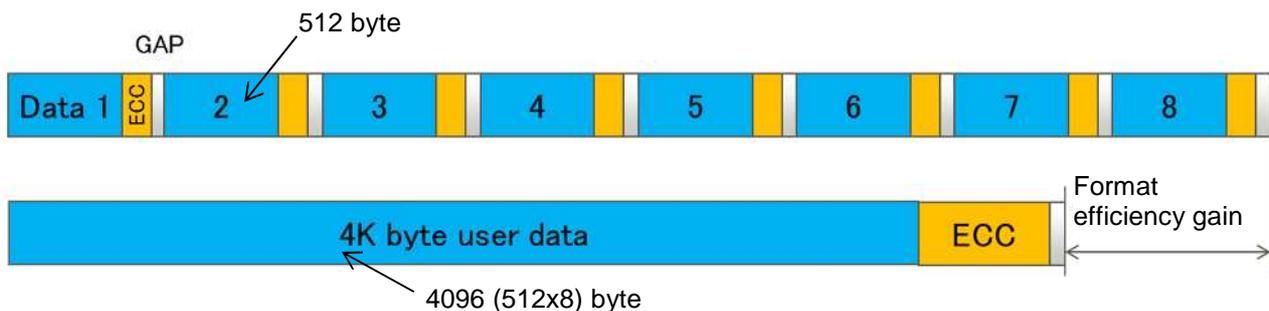


Figure-1 Difference between physical sector size 512 bytes and 4096 bytes

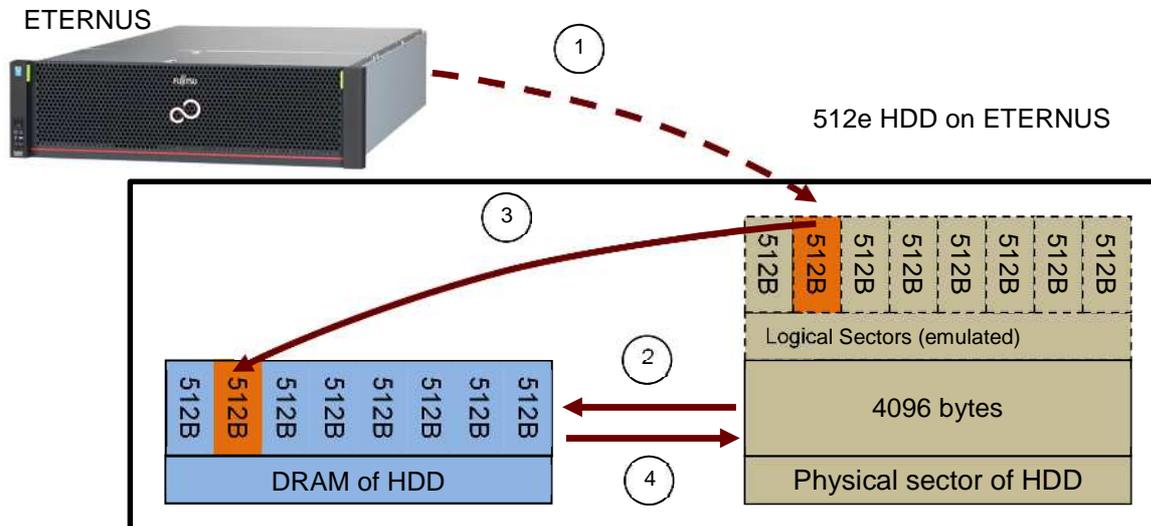
The use of this format requires support for all affected HW and SW components:

HDDs, disk controllers, operating systems, drivers, applications and other tools.

Since the conversion of all these components by the manufacturers will take some time, the technology of 512-byte sector emulation has been introduced for the interim period. This means that although HDDs already have internal physical 4k sectors, from an external viewpoint they will emulate logical 512-byte sectors. This ensures that these new HDDs can also be run in existing environments. Such HDDs sometimes also are nicknamed “512e” (“e” for “emulation”). Equally, conventional HDDs with physical sector size 512 byte are nicknamed “512n” (“n” for “native”).

The ETERNUS DX S4/S3 series supports 512e only as Advanced Format HDDs.

In 512e, all accesses to media are processed by a physical sector (4096 bytes) unit. Write operations to smaller size than the physical sector size are processed by Read-Modify-Write (RMW) cycle. Figure-2 describes how the RMW works.



- 1) 512-byte sector write requested by ETERNUS to the 512e
- 2) HDD reads 4kB sector into its DRAM
- 3) HDD updates requested 512-byte sector in DRAM
- 4) HDD overwrites 4kB sector on media

Figure-2 How the Read-Modify-Write works

Performance

Since "Read-Modify-Write" cycles require an additional full disk rotation on the HDD, they reduce performance. Thus, in the interest of best possible performance on the part of the application and operating system, accesses that only write a part of a physical 4k sector should be avoided. There are two such cases of access: (1) when a write access takes place with a block size that is not a full multiple of 4k; (2) when a write access does take place with a block size that is a full multiple of 4k, but this access is not aligned to the boundaries of the physical sectors.

These accesses are called "Unaligned accesses". On the other hand, accesses with a block size, which is a full multiple of 4k and aligned to the boundaries of physical sectors, are called "Aligned accesses".

The performance of 512e is equal to or higher than 512n on Aligned accesses. However, operating systems and applications that do not support 512e can perform Unaligned accesses. Fujitsu recommends using 512e aware operating systems and applications when 512e HDDs are used on ETERNUS systems to avoid the performance impacts.

Examples of 512e aware Operating Systems

To achieve optimal performance, operating systems and applications need to be 512-aware (to support 512e). Table-1 shows the examples of the 512e aware OSs.

512e-aware Operating Systems	
Name	Additional Requirements
Windows Server 2016	
Windows Server 2012 R2	
Windows Server 2012	
Windows Server 2008 R2	Requires application of "KB 982018" or "SP1". Hyper-V is not supported.
Windows Server 2008	Requires application of "KB2553708". Hyper-V is not supported.
Red Hat Enterprise Linux 7	
Red Hat Enterprise Linux 6	2.6.32 or later
Red Hat Enterprise Linux 5	RHEL5.6 or later
SUSE Linux Enterprise Server 12	
SUSE Linux Enterprise Server 11	SLES11 SP2 or later
Oracle Solaris	Oracle Solaris 10 or later
Oracle Linux 7	
Oracle Linux 6	UEK:2.6.32 or later UEK2:2.6.32 or later
Oracle Linux 5	Oracle Linux 5.6 or later
AIX7.2	
AIX7.1	TL03 SP3 + IV56367 or later
AIX6.1	TL09 SP3 + IV56366 or later
Oracle VM 3.X	V3.2, UEK2:2.6.39 or later
VMware vSphere™ 6.5	

Table-1 Examples of 512e-aware operating systems

Some components may not be available in all countries / sales regions.

Comments :

- Some OSs require partition configuration changes to use 512e HDDs. Please check on to the OS vendor.
- The following applies for all applications: If the application works with "unbuffered writes", this must take the physical sector size into account for optimal write performance when accessing 512e HDDs.
- About the other OS not listed in Table-1, please ask to the OS vendor.

Non-512e-compliant Operating Systems

OSs listed in Table-2 do not support 512e HDDs. Fujitsu recommends using 512n HDDs on these OSs.

Note: The both of VMware vSphere™ 5 and 6 technically support ETERNUS DX and AF with 512e drives, though it is not recommended by VMware due to possible performance degradation.

Non-512e-compliant Operating Systems	
Name	Reminder
VMware vSphere™ 6 VMware vSphere™ 5.5 VMware vSphere™ 5.1 VMware vSphere™ 5	Fujitsu recommends using 512n on these OSs. Usage of ETERNUS DX and ETERNUS AF with 512e drives in VMware 5.5 and 6 environments is supported but due to possible performance degradation not recommended by VMware.
AIX5.3	
HP-UX 11iv3 HP-UX 11iv2	
XenServer 7 XenServer 6	

Table-2 Non-512e-compliant operating systems

About applications

As long as an operating system is 512e aware, applications running on the OS can work as 512e aware. However, some data base applications, for example “Oracle Server”, access to HDDs in the unique way. In those cases, the applications also need to be 512e aware. Please confirm the application vendor if it supports Advanced Format.

About middle wares

512e-aware middle wares	
Name	Reminder
Oracle Database 12c R2 (12.2.0.1.0) Oracle Database 12c R1 (12.1.0) Oracle Database 11g R2 (11.2.0)	The operating system need to be 512e aware. 512n and 512e can not be mixed in an ASM disk group.

Table-3 512e-aware middle wares

About the other middle wares not listed in Table-3, please confirm the middle ware vendor if it supports Advanced Format.

Others

There is a possibility of unexpected write performance down if there are RAID groups configured by 512n and configured by 512e in a same system, and a non-512e-compliant OS accesses to the system, the following functions can be affected; Thin Provisioning, Flexible Tier, Logical Device Expansion, LUN Concatenation, and Wide Striping. Please see manuals for each functions for details.