


# White Paper

## All-Flash Arrays for Improving Database Performance to Open Up Potential for New Applications

### ETERNUS AF series

The ETERNUS AF series is an SSD-only storage system (all-flash arrays) that was developed to take full advantage of the SSD's capabilities. The ETERNUS AF series offers a dramatic improvement in database performance that ensures seamless responsiveness. This performance also simplifies operation and maintenance.



<b>Contents</b>	
<b>Preface</b>	<b>3</b>
<b>1 Database Requirements</b>	<b>4</b>
1.1 Hardware Performance Improvement	4
1.2 Emergence of All-Flash Arrays	5
1.3 ETERNUS AF series All-Flash Arrays Purpose-Developed for Use with SSDs	5
<b>2 Dramatic Improvement in Database Performance with All-Flash Arrays</b>	<b>6</b>
2.1 Database Performance and Server and Storage System Resources	6
2.2 Using an All-Flash Array to Boost Database Performance	6
2.2.1 Response Improvements	7
2.2.2 Database Performance Improvement	8
2.2.3 Consolidation of Database Storage by Maintaining Low Latency	9
2.3 Performance Improvements without Changing the Design and Operation of the Existing Storage System	10
<b>3 System Benefits of ETERNUS AF series Performance</b>	<b>11</b>
<b>4 Conclusions</b>	<b>11</b>

<b>Figures and Tables</b>	
Figure- 1 Database requirements	4
Figure- 2 Improvements in hardware performance	4
Figure- 3 FUJITSU Storage ETERNUS AF series	5
Table- 4 Random access performance (IOPs) of HDD and SSD relative to a 10,000rpm HDD	6
Figure- 5 How response time influences processing speed	7
Figure- 6 Data warehouse database benchmark	8
Figure- 7 Storage area consolidation	9

## Preface

As the information society continues to grow, the data it deals with is also growing in volume, detail, and length of storage. Examples include image recognition, sensing, IoT, and other realtime data as well as more detailed management and enhanced traceability in existing production and distribution management applications. Just as databases are essential to utilizing this huge volume of data, storage systems are indispensable for managing and archiving all of this data. Most of the data in databases is stored in a storage system. This makes the performance of the database highly dependent on that of the storage system.

Database performance needs to improve to keep pace with this growing volume of data. The dramatic performance boost provided by all-flash arrays solves this database performance challenge.

The product lineup and product information stated in this document are current as of December 2017.

### ■ Target Readers

This document targets readers who are familiar with the following topics.

- General knowledge of databases and disk storage systems

### ■ Applicable Models

This document relates to the following storage systems (all-flash arrays).

- FUJITSU Storage ETERNUS AF250 S2/AF650 S2
- FUJITSU Storage ETERNUS AF250/AF650

### ■ Naming Conventions

The following abbreviations are used in this document.

- ETERNUS AF series ..... FUJITSU Storage ETERNUS AF250 S2/AF650 S2 and ETERNUS AF250/AF650
- AFA ..... All-Flash Array
- HDD ..... Hard Disk Drive
- SSD ..... Solid State Drive
- DB ..... Database

## 1 Database Requirements

Database environments are subject to ever-evolving requirements year by year. The ways in which databases are used continue to change. For example, along with integrated management and advanced data leveraging by using ERP, SCM, and CRM, data use in distribution, production, and IoT is becoming both more fine-grained and more diverse by advances in sensing.

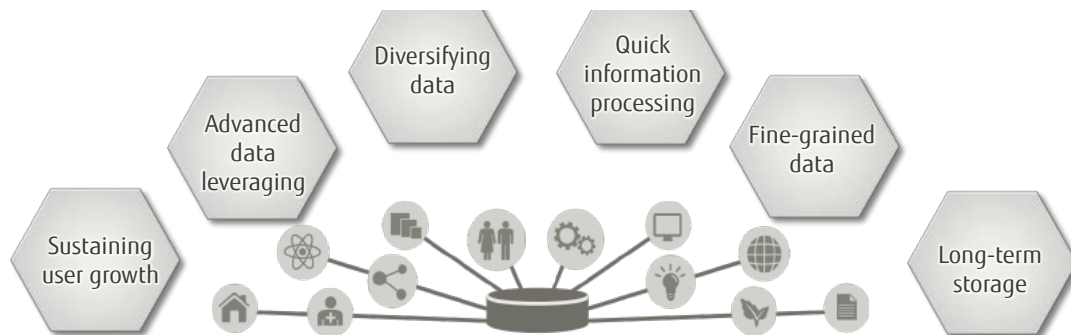
The demands placed on databases change year after year.

- Integrated management and advanced data leveraging by using ERP, SCM, and CRM
- Greater requirements for fine-grained and realtime data in production and distribution management driven by advances in sensing
- Growing diversity of data as a result of technical innovations in fields like sensing and the IoT, and its use in analytics

The ways in which databases are used are changing significantly, with data increasing in volume and becoming more detailed and diverse.

Databases are used for more than just managing data, being called on to deliver high-speed access in response to many and varied requirements in such diverse applications as the management of image data and other media content, big data analytics, the use of data from mobile devices, and high-speed electronic payments.

While the performance of the hardware on which databases run continues to improve in response to the demand for faster database access, storage systems remain a bottleneck.



**Figure- 1 Database requirements**

### 1.1 Hardware Performance Improvement

Ongoing improvements are being made in hardware performance to keep up with growing data volumes. Unfortunately, there is a major mismatch between performance trends in servers and storage systems.

#### • Servers

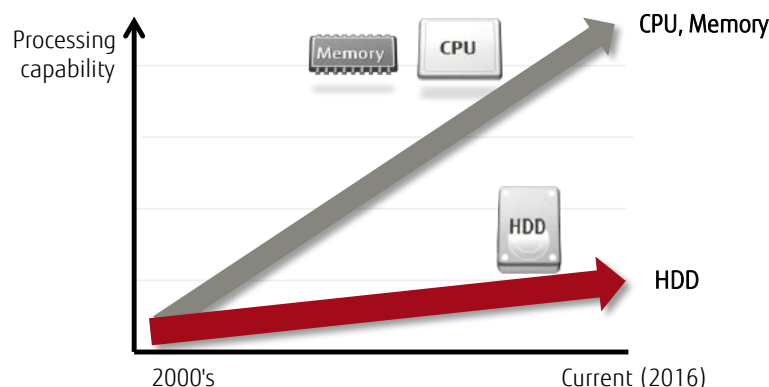
CPU and memory specifications clearly show the improvements in server performance. Per-core CPU performance has improved by two to three times over the last decade while advances in core density have increased the number of cores that can be installed per CPU, to the extent that per-CPU performance has increased more than tenfold. Memory, meanwhile, has achieved a more than fivefold increase in data transfer speed while the amount of memory that can be installed on a server has increased by a factor of as much as one hundred.

It is fair to say that server hardware has delivered major performance improvements in response to the demand for faster databases and larger data volumes.

#### • Storage systems

In the case of storage system performance, meanwhile, enhancements have been made by increasing the cache size a hundredfold or more over the last decade, with higher cache hit rates and a larger number of HDDs per stripe. However, while operating more HDDs in parallel increases throughput, the random access response has remained at an average of 10ms, meaning that these enhancements have not kept pace with the improvements in server performance.

The root of the problem lies in the HDDs. While the increase in cache sizes and recording densities has led to faster data transfer speeds, there has been little improvement in random access performance because the disk rotational speed remains at 15,000rpm, the same as it was 10 years ago.



**Figure- 2 Improvements in hardware performance**

## 1.2 Emergence of All-Flash Arrays

SSDs entered the market about 10 years ago, offering a new generation of drives.

SSDs use a form of non-volatile memory called NAND flash memory. Access is considerably faster than HDDs because data is stored in solid-state memory and there are no disk platters to rotate into position. While this gives SSDs excellent random access performance, adoption of SSDs has been constrained by downsides, including capacity, lifespan, and cost.


In recent years, SSDs have achieved capacities equal to those of HDDs, while the ongoing fall in prices has brought them into an affordable range.

On the other hand, a number of characteristics unique to SSDs remain, which include a finite number of write cycles, blocks that must be erased before writing, writing on a per block basis, and garbage collection that must be performed by the drive (elimination and reorganization of unused blocks).

Accordingly, the ETERNUS AF series All-Flash Arrays were purpose-developed for use with SSDs to compensate for these characteristics. The ETERNUS AF series extracts maximum performance from its SSDs while also maximizing SSD lifespan.

## 1.3 ETERNUS AF series All-Flash Arrays Purpose-Developed for Use with SSDs

Existing disk storage systems use striping and controller-based caching to make up for the high latency of HDDs. The high performance of SSDs, however, outstrips that of the controller when used in a RAID configuration. Developed specifically for SSDs, the ETERNUS AF series has been made lighter and faster by excluding HDD-specific operations, having been designed as a reliable and convenient storage system that takes maximum advantage of the performance of SSDs.



**Figure- 3 FUJITSU Storage ETERNUS AF S2 series**

The ETERNUS AF series is equipped with the following functions to fully utilize the SSDs.

- SSD access optimization  
Alignment is used to optimize SSD block writes (4KB) to minimize the number of blocks needed. (lower latency, less garbage collection, longer life)
- Minimized performance degradation during garbage collection  
The controllers of the ETERNUS AF series monitor the garbage collection performed by each SSD to minimize performance degradation by avoiding access to those drives where garbage collection is in progress and accessing other drives instead. (Because SSDs must erase each block before writing to it, they reorganize and eliminate unused blocks through a process called garbage collection, with a consequent degradation in drive performance while this process is taking place.)
- Flexible write-through  
Controller operation is streamlined and the overall performance of the storage system is improved by optimizing the cache to suit the SSDs. This is done by maintaining data equivalence between the two controllers in a storage system to the bare minimum needed to ensure reliability.

While all-flash arrays have been used in applications that take advantage of their high speed, such as scientific computation, data analysis, and VDI, the falling price of SSDs is opening up a wider range of potential applications.

In addition to functions that extract 100% of the performance of high-speed SSDs, the ETERNUS AF series also utilizes high-performance CPUs and multiprocessing techniques to make dramatic improvements in the performance of the storage system.

The ETERNUS AF series delivers major benefits in the form of shorter execution times for business systems and faster databases.

## 2 Dramatic Improvement in Database Performance with All-Flash Arrays

The performance of all-flash arrays also plays a major role in the performance of applications. Along with the servers, storage system performance is an important factor for database systems that handle large quantities of data and require frequent drive access. This chapter describes how an all-flash array boosts database performance.

### 2.1 Database Performance and Server and Storage System Resources

Databases are used to extract the required information from large amounts of data at high speed.

To achieve high-speed data access, database operations should be performed on the server as much as possible. The access frequency to the storage system can be reduced by keeping and buffering an index in server memory that uniquely identifies data stored on the storage system. Nevertheless, the saving and accessing of data on the storage system is still needed because of the limited resources of servers and their inability to store data permanently.

However, while a storage system is a vital resource for databases, it is also a bottleneck.

The data transfer speed of the server memory (RAM), which consists of electronic circuitry, is unparalleled and faster (by more than 4000 times\*) than that of an HDD, which depends on mechanical operation.

In contrast, while RAM has dramatic performance advantages, it is roughly 80 times more expensive than HDD per unit of data stored. While server hardware resources have been improved in performance and expanded in capacity, making greater use of the server hardware resources is costly. In addition, the reality is that storage systems are still needed for storing data.

The storage systems that store the bulk of the data are a major factor in database performance. As the amount of data handled by the systems keeps increasing yearly, the increased amount of access to the storage system causes the system performance to degrade. Therefore, additional performance is required for the storage systems.

\*Data transfer speeds were compared using a standard value RAM and a randomly accessed HDD

### 2.2 Using an All-Flash Array to Boost Database Performance

SSDs have a performance approximately 100 times higher than HDDs. They are also much cheaper than RAM, which are roughly twice as expensive as HDDs per unit of data stored.

All-flash arrays take maximum advantage of the performance of SSDs. The performance of an all-flash array also makes itself apparent in databases (which are high level applications).

Drive type	Read	Write
10,000rpm HDD	-	-
15,000rpm HDD	1.7 times (approx.)	1.5 times (approx.)
SSD	150 times (approx.)	100 times (approx.)

**Table- 4 Random access performance (IOPs) of HDD and SSD relative to a 10,000rpm HDD**


### 2.2.1 Response Improvements

All-flash arrays deliver extremely low latency.

Disk storage systems are typically designed to achieve a response time of 20ms or less. A technique used to keep the response time below 20ms as access volumes increase is to add dozens of additional HDDs and operate them in parallel.

In contrast, the ETERNUS AF series can deliver access performance in the hundreds of thousands of IOPs, a level unheard of in disk storage systems, together with response times of 0.5ms or less, even under heavy load. That is, even under heavy load, the ETERNUS AF series has a response time that is 1/40th or less than that of a disk storage system.

As the server is kept waiting during this response time, a faster response means faster processing, with major benefits for database performance.



**Figure- 5 How response time influences processing speed**

All-flash arrays can deliver greater benefits for databases with a low CPU load in situations where the disk storage system is a bottleneck. The bottleneck imposed by disks can be seen by looking at the %io wait value in the Linux sar command, and at Avg. Disk sec/Read, Avg. Disk sec/Write, and Avg. Disk sec/Transfer in the Windows performance monitor.


Lower latency also speeds up the execution of batch and other jobs. For systems where the server's CPU usage rate does not reach 100% and random access to the disks is a bottleneck, the disk access time is reduced to 1/40 and the execution time of the batch jobs can also be reduced by a maximum of 75%.

Along with reducing the waiting time in the CPU and making better use of the server resources, this also opens up the potential to reduce server resources in situations where faster processing is achieved due to the performance of the server resources.

### 2.2.2 Database Performance Improvement

In a Microsoft SQL Server benchmark based on a row store database used as a data warehouse, an all-flash array with only eight SSDs demonstrated an improvement in database performance (throughput) of approximately 2.2 times compared to a disk storage system made up of 72 HDDs (CE + DE × 2).

The all-flash array achieved higher performance despite using only one-ninth as many drives as the disk storage system, with a per-drive performance improvement of approximately 20 times.



**Figure- 6 Data warehouse database benchmark**

A performance increase of approximately 2.2 times was achieved despite the database system already making full use of the server resources, with the test being conducted on a high-performance server with more than 20 CPU cores and on which the database system operated with a server CPU utilization rate of around 80%.

This demonstrates how replacing an existing disk storage system with an all-flash array is, on its own, enough to significantly improve the performance of many database systems.

The ETERNUS AF series can access data with lower access performance (IOPs) degradation even with block sizes of 4KB or larger. Even database's 8KB or larger block size access can be handled.

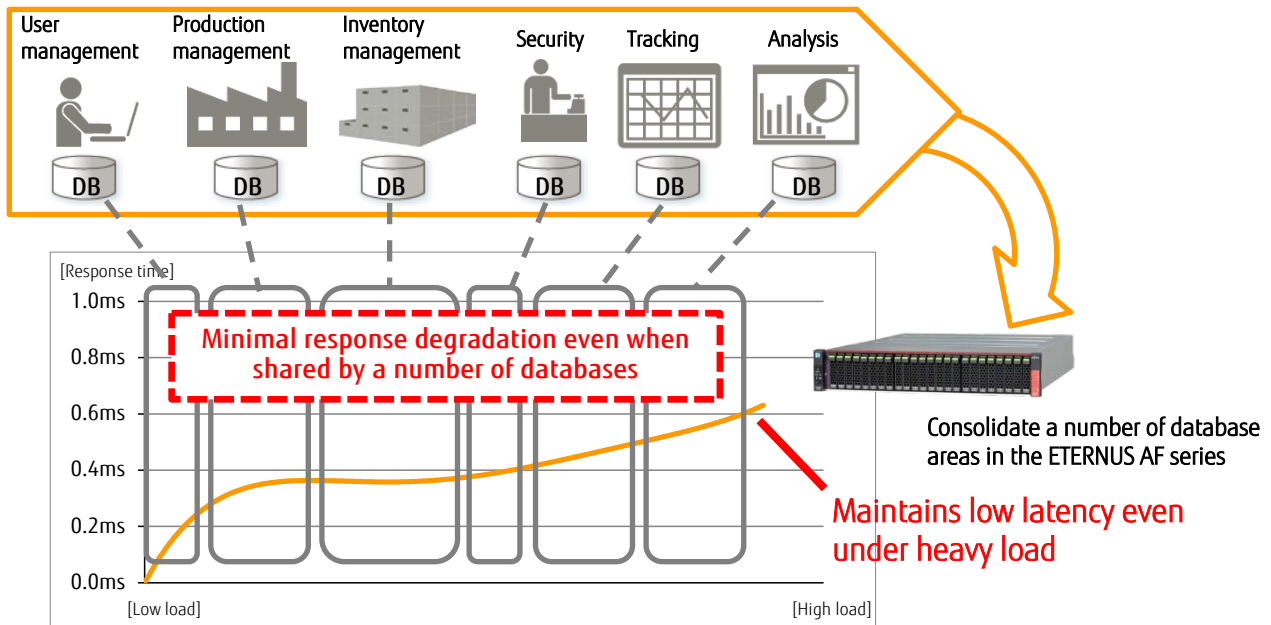


### 2.2.3 Consolidation of Database Storage by Maintaining Low Latency

The ETERNUS AF series can maintain low latency even under heavy load.

The response of disk storage systems slows significantly when heavily loaded, with response times increasing to over 100ms. In contrast, the ETERNUS AF series can maintain a low latency of approximately 0.6ms even under heavy load.

Due to this feature of maintaining low latency under heavy load, the ETERNUS AF series can consolidate the storage areas of the database systems that require high storage system performance into a single storage system and can also allow the performance of the ETERNUS AF series to be shared across multiple databases.



**Figure- 7 Storage area consolidation**

It is also possible to consolidate databases that remain on physical servers to ensure the storage system performance.

The consolidation of a large number of systems is also possible in the case of systems with stringent performance requirements by using the QoS automation function to ensure that the required level is achieved, and using the Thin Provisioning Pool function to optimize capacity.

As storage area consolidation can keep system databases independent of each other, unlike database federation, consolidation can be achieved using the existing system environment and operation, without the need for a new operational design or other work.

## 2.3 Performance Improvements without Changing the Design and Operation of the Existing Storage System

Among the methods used to improve database performance are new techniques such as column store databases or key-value stores, and techniques such as in-memory databases that utilize server resources. Unfortunately, all of these require major system modifications. The impacts of these modifications affect not only the applications but also extend to the system configuration, operation, and associated infrastructure.

An all-flash array boosts the database performance by improving the performance of the storage systems, which are bottlenecks on existing platforms. In addition to dramatically increasing the performance of the storage systems, the ETERNUS AF series also replicates the functions of a disk storage system to enable existing systems to continue using the same configuration and operation. Accordingly, database performance can be improved simply by replacing an existing disk storage system with an ETERNUS AF series.

Improvements in database performance not only make systems easier to operate, they also facilitate cost-cutting in areas such as development, operation, and maintenance.

- Applications ..... Modifications are unnecessary for platform modifications, less work needed for SQL tuning and other forms of optimization to achieve performance requirements.
- Operation ..... Faster execution enables different backup practices, such as the range of the backup area and switching from differential to full backups (allowing faster recovery), more frequent batch processing, and securing retry times.
- Maintenance ..... Tasks such as data export, reconfiguration, and system modification can be completed in a shorter amount of time and the retry and confirmation times can be secured. By securing the confirmation times, operational errors can be prevented.

Improving the performance of a database system leads to ongoing reduction of risks and costs by ensuring that adequate time is available for performing work and that work can be completed more quickly.

The ETERNUS AF series works with backups, OSs, and virtualization software, and are equipped with redundancy and other functions to keep the same operation and design as the existing ETERNUS DX S4/S3 series.

- Advanced Copy functions ..... A wide range of copy functions are available to satisfy different backup requirements
- Remote Advanced Copy functions ..... Copying between storage systems enables data backups to a remote ETERNUS AF series or ETERNUS DX S4/S3 series storage system  
Suitable for remote backup and disaster recovery practices
- Storage Cluster function ..... Business continuity can be enhanced by using redundant storage systems
- Thin Provisioning Pool ..... Virtualizes storage resources to optimize physical capacity
- Deduplication/Compression function ..... Increases the data capacity of the storage system by reducing the amount of data stored on physical drives
- Software integration functions ..... Numerous integration functions are available, including Offloaded Data Transfer (ODX) for Windows, VAAI and VASA for VMware vSphere, VMware vSphere virtual volumes function, and VolumeDriver for OpenStack volume management

### 3 System Benefits of ETERNUS AF series Performance

By improving database performance, the ETERNUS AF series facilitates cost savings on the associated systems, hardware resources, and infrastructure.

#### Benefits of performance improvements

- Reduced cost of database system performance tuning  
The improvement in performance provided by the ETERNUS AF series significantly reduces the cost of database tuning. Whereas in the past, satisfying the performance requirements has required that tuning be performed as part of application development, with ongoing tuning such as data dispersal and index rebuilding also being required once the system starts operation, the performance provided by the ETERNUS AF series reduces these tuning costs.
- Reduction in server resources made possible by greater storage system performance  
The benefits of a reduction in server resources include fewer CPU licenses and the application of lower license grades.
- A faster database response improves the response of other systems that use the database  
This improves convenience and can be used to reduce hardware resources on these other systems. Database response times become faster relative to the response performance requirements, making more processing time available to other systems.
- Shorter execution times made possible by improvements in database performance open up the potential for new system applications  
Improvements in database performance enable shorter execution times. Shorter execution times in turn enable batch jobs to be run more frequently than just once a day, resulting in shorter lead times and lower yield. Shorter analysis times also open up possibilities for new strategies based on a more diverse analysis that is to be run more often with larger data sets.
- Improvements in the performance of the storage system enable easier maintenance  
Improvements in performance of the storage system shorten the time taken for maintenance tasks that involve the processing of large amounts of data, such as data exporting or the data patches that need to be installed during system upgrades. Being able to perform maintenance work more quickly also cuts costs because work can be performed more reliably and work schedules can be shortened.

#### Reduction in the storage system's hardware

In the case of disk storage systems, large numbers of HDDs have needed to be installed to satisfy performance requirements. As shown in "Figure- 6 Data warehouse database benchmark", the ETERNUS AF series significantly reduces the number of drives required to satisfy the requirements.

This lower number of drives cuts normal operation costs by reducing the installation footprint and power consumption.

The ETERNUS AF series also permits additional drives to be added one at a time, making it possible to minimize up-front costs by starting small and to expand capacity as the system and performance requirements grow.

#### Consolidation of database systems

The ETERNUS AF series has a very short latency of 0.5ms and maintains reliable response performance even under heavy loads. This means the ETERNUS AF series can be used for the consolidation into a single storage system of database systems that use dedicated storage and those that have not been virtualized due to the need for a high level of IO performance.

In addition to making full use of the high IO performance of the ETERNUS AF series, consolidation of database systems enables performance sharing and improved convenience and operation of the individual systems.

### 4 Conclusions

A lot of time and money has been devoted to maintaining the performance of databases in the face of ever-increasing quantities of data. The performance of the ETERNUS AF series resolves the challenges of database performance while also improving convenience and ease of operation.

With database software licensing evolving to deal with new developments such as virtualization and the cloud, the performance of the ETERNUS AF series is also contributing to the ongoing trend toward the virtualization and consolidation of database systems.

All-flash arrays already provide platforms on a similar level to disk storage systems while also significantly improving the performance provided by those platforms. The performance of the ETERNUS AF series is opening up possibilities for new system applications and supporting business expansion.

---

**Contact**  
FUJITSU Limited  
Website: <http://www.fujitsu.com/eternus/>

■ Trademarks

Microsoft®, Windows, SQL Server, and other Microsoft product names are trademarks or registered trademarks of Microsoft Corporation in the United States, and other countries. Linux® is the registered trademark of Linus Torvalds in the United States, and other countries. The product names and company names in this document are registered trademarks or trademarks of their respective companies.

Trademark symbols such as (R) and (TM) may be omitted from system names and product names in this document.

■ Disclaimer

FUJITSU Limited is not responsible indemnity that might be caused by the content in this documentation or any damage related to the content in this documentation.