

# NetApp E-Series Breaks Performance Barriers



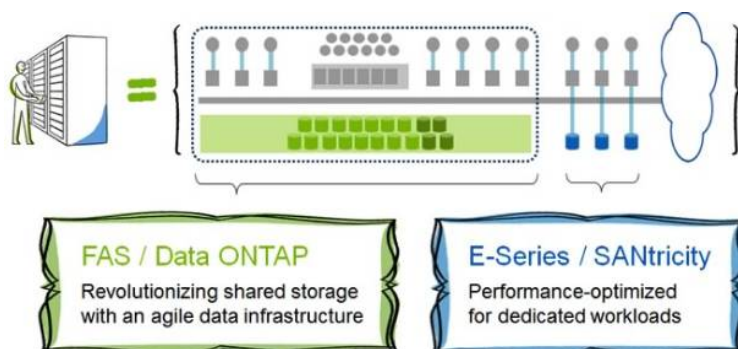
**Michael Jastad**

Sr. Technical Marketing Engineer

In March 2013, NetApp announced the latest addition to the E-Series line, the NetApp® E5500, designed to provide industry-leading performance, efficiency, and reliability for big data and high-performance computing (HPC). The E5500 lays the foundation for highly available high-capacity application workflows with half the operational costs and footprint of competing systems.

The E-Series became a part of the NetApp product portfolio when NetApp acquired the Engenio line from LSI in 2011. The E5500 is built on a proven legacy of over 650,000 deployed storage systems, including installations in some of the toughest computing environments in the world, such as the [Sequoia supercomputer](#) at Lawrence Livermore National Labs, currently ranked as the [second largest supercomputer](#) in the world.

All E-Series models are designed for use by data-intensive applications such as Hadoop, video surveillance, seismic processing, and other big data and high-performance computing applications that need dedicated storage. The E-Series is available through NetApp and its channel partners; branded versions of E-Series models are also available through OEM partners such as SGI and Teradata.



**Figure 1)** NetApp E-Series is designed specifically for dedicated workloads.

This article describes the performance of the E5500, provides an overview of the entire E-Series product line, and introduces key features such as Dynamic Disk Pools and SSD Cache.

## E5500 Breakthrough Performance

With the [March announcement](#), the E5500 became the new E-Series flagship system. Designed to meet the most demanding big data and HPC needs, the E5500 raises the bar for performance and performance density. A single E5500 delivers up to 12GB/sec read performance using an 8U, 120-drive configuration. That's an incredible amount of performance from a small footprint.

Performance for storage systems like the E5500 is often measured in terms of bandwidth rather than IOPS because many big data and HPC systems need to move large amounts of data with maximum throughput. The [Storage Performance Council's SPC-2 benchmark](#) is the most widely used benchmark in this space. According to the Storage Performance Council Web site:

*Explore*

### Learn More About E-Series

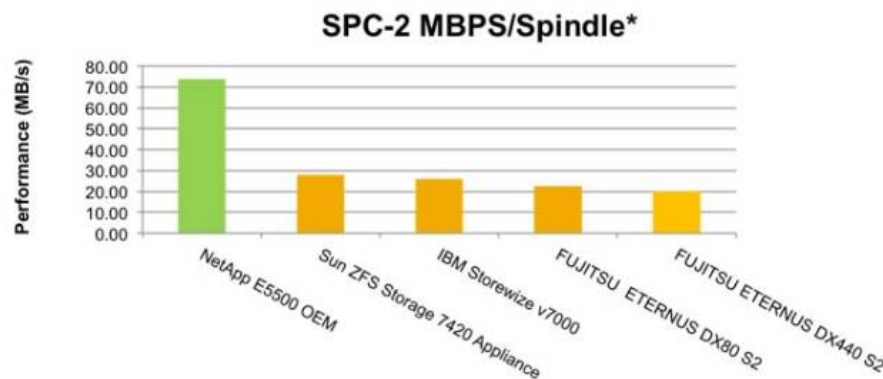
Want to get started with E-Series? Head to [NetApp.com](#) for the latest information on each of the E-Series platforms:

- [E2600](#)
- [E5400](#)
- [E5500](#)

SPC-2 consists of three distinct workloads designed to demonstrate the performance of a storage subsystem during the execution of business-critical applications that require the large-scale, sequential movement of data. Those applications are characterized predominately by large I/Os organized into one or more concurrent sequential patterns. A description of each of the three SPC-2 workloads is listed below as well as examples of applications characterized by each workload.

- **Large File Processing:** Applications in a wide range of fields that require simple sequential processing of one or more large files such as scientific computing and large-scale financial processing.
- **Large Database Queries:** Applications that involve scans or joins of large relational tables, such as those performed for data mining or business intelligence.
- **Video on Demand:** Applications that provide individualized video entertainment to a community of subscribers by drawing from a digital film library.

NetApp OEM partner SGI recently measured the performance of its SGI InfiniteStorage 5600, which is based on the NetApp E5500, using the SPC-2 benchmark. If you compare the performance demonstrated by the SGI SPC-2 submission to published SPC-2 numbers from competing vendors, the E5500 delivers the best price/performance and the lowest cost per unit of bandwidth (SPC-2 measures bandwidth in megabytes per second, or MBps). The E5500 also provides 2.5 times the performance per spindle, as illustrated in Figure 2.



\* Based on SPC-2 publications in 2011 or later with a total price of \$500K or less

**Figure 2)** Performance delivered per disk spindle on the SPC-2 benchmark shows the clear performance advantage of the E5500.

The E5500 achieves this tremendous level of performance using internal PCIe Gen 3.0 x8 buses. A unique ability to use both hardware and software RAID engines enables the E5500 to stream data from disk very efficiently and process tremendous amounts of I/O. SAS expansion ports are capable of delivering up to 48Gb/s, and each controller can access all drive ports.

The ability to drive more performance from fewer disks means that the E5500 is very efficient. In addition to delivering the highest throughput per spindle, it also delivers the highest throughput per footprint and watt of power. A single rack accommodates up to 10 E5500 storage systems and 600 drives (or up to 5 120-drive systems as configured for the SPC-2 submission described above).

Although the E5500 will most often be deployed for bandwidth-intensive workloads, its transactional performance is equally impressive. For 15K drives, the E5500 can deliver up to 150,000 IOPS on 4K random reads. Equally adept at IOPS and bandwidth, the E5500 is a great choice for a wide range of performance-oriented workloads.

## The E-Series Product Line

---

The full E-Series product line consists of three storage systems: E5500, E5400, and E2600. Preconfigured E-Series solutions are available for specific workloads including Lustre, Hadoop, surveillance, and media content management.

- The E2600 is the entry-level E-Series system designed to deliver performance value, reliability, and ease of use. It's the ideal system to support workloads like transaction processing, mail, and decision support. The E2600 is also

often deployed to provide metadata storage for clustered file systems like Lustre while E5400 and E5500 systems provide object storage.

- The E5400 has been the workhorse of the E-Series over the past few years, with 6GB/sec performance, great capacity, and a rich feature set. It has been deployed for a variety of big data and HPC applications including Hadoop, video surveillance, full-motion video, oil exploration, data mining, and a variety of government and scientific workloads. The E5400 is also well suited for transactional workloads.
- The E5500 delivers up to twice the performance of the E5400, offering the highest levels of throughput for data-intensive workloads demanding extreme bandwidth.

All E-Series systems are managed with enterprise-proven SANtricity® software, which lets you easily tune a system to achieve maximum performance and utilization. It provides features such as SSD Cache and Dynamic Data Pools (discussed later) to further enhance performance. Dual-active controllers with redundant I/O paths; automated failover; and fully redundant, hot swappable components to protect availability are standard.

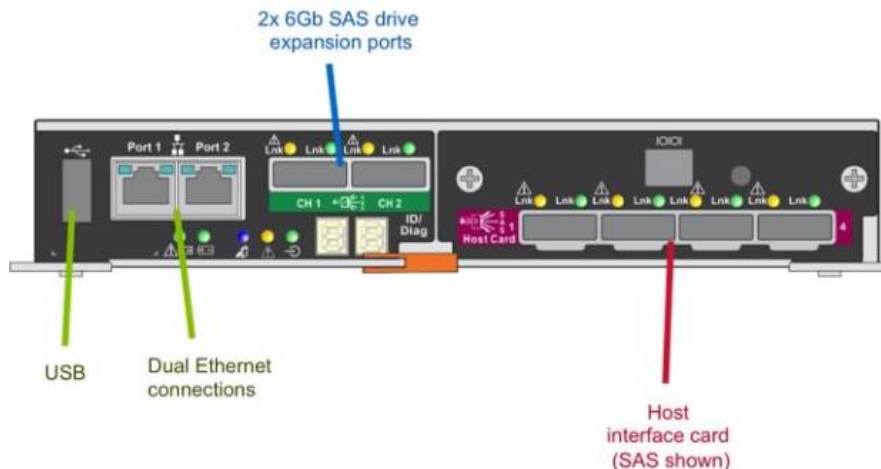
The three E-Series storage systems are distinguished by their performance, capacity, and connectivity options.

**Table 1)** Comparison of E2600, E5400, and E5500.

|                                | E2600  | E5400  | E5500                      |
|--------------------------------|--|--|----------------------------|
| <b>Max. Performance (read)</b> | 4GB/sec                                      | 6GB/sec  | 12GB/sec                   |
| <b>Max. Drives</b>             | 192  | 384  | 384                        |
| <b>Max. Capacity</b>           | 576TB  | 1152TB   | 1152TB                     |
| <b>Host Connectivity</b>       | 6Gb SAS<br>8Gb FC<br>10Gb iSCSI<br>1GB iSCSI | 6Gb SAS<br>8Gb FC<br>10Gb iSCSI<br>40Gb InfiniBand | 6Gb SAS<br>40Gb InfiniBand |

Note that performance of these systems is typically reported in GB/sec rather than IOPS. This is because the important applications in the E-Series space need to move large amounts of data with maximum throughput.

Each E-Series controller features dual Ethernet connections for management. The E5500 has dual SAS ports on each controller for expansion while the E2600 and E5400 have a single SAS port per controller. Host cards are optional on the E2600 and E5400; the E5500 requires a host card.



**Figure 3)** E5500 controller, rear view.

### E-Series Models and Expansion Options

Each storage system includes three separate models, as shown in Table 2.

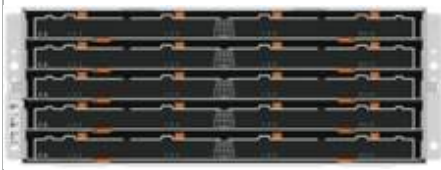


**Table 2)** Configurations of E-Series models.


|  |  |
|--|--|
|  |  |
|--|--|

| Model                      | Configuration  |
|----------------------------|--|
| <b>E5560, E5460, E2660</b> | 4U/60-drive configuration supports high-capacity 3.5" 7.2K HDDs, high-performance 2.5" SFF 10K HDDs, and 2.5" SFF SSDs         |
| <b>E5524, E5424, E2624</b> | 2U/24-drive configuration uses 2.5" SFF drives (SSD and 10K HDD) to deliver great performance/watt and bandwidth per rack unit |
| <b>E5512, E5412, E2612</b> | 2U/12-drive configuration supports 3.5" LFF 15K and 7.2K HDDs and provides the lowest entry price                              |

Base systems can be expanded through the addition of any of three disk shelves, as shown in Table 3.

**Table 3)** E-Series expansion options.

| DE6600  | DE5600   | DE1600  |
|---|--|---|
|    |    |                |
| A 4U/60-drive shelf that provides industry-leading density with a unique 12-drive drawer design. Each of its 5 drawers can be opened for drive access and replacement without disrupting data access, providing a level of reliability, ability, and serviceability unmatched by competing offerings. | A 2U/24-drive shelf that accommodates low-power 2.5" drives and delivers IOPS performance and throughput density. Can be configured with SSDs for extreme performance. | A 2U/12-drive shelf that can include high-performance 15K drives as well as high-capacity drives. |



**Figure 4)** E-Series systems can be homogeneous or heterogeneous. (Each configuration shown includes 120 disk drives.)

### E-Series Reliability, Availability, and Serviceability

The E-Series combines field-proven technology and leading reliability, availability, and serviceability features, protecting your valuable data and delivering uninterrupted availability.

#### Hardware Features

Each array is architected to deliver enterprise levels of availability with:

- Dual-active controllers, fully redundant I/O paths, and automated failover

- Battery-backed cache memory that is destaged to flash upon power loss
- Extensive monitoring of diagnostic data that provides comprehensive fault isolation, simplifying analysis of unanticipated events for timely problem resolution
- Proactive repair that helps get the system back to optimal performance in minimum time

The rear view of the E5560 shows the redundant controllers, power, and cooling.



**Figure 5)** E5560 rear view showing dual controllers, power supplies, and cooling fans. In the DE6600 disk shelf, the controllers are replaced with Environmental Service Modules providing SAS connectivity.

### AutoSupport

NetApp has been offering the [AutoSupport™](#) tool on the FAS product line almost from the very beginning. On FAS systems the AutoSupport tool has been shown to improve storage availability and to reduce the number of priority-1 support cases by up to 80%.

This capability has now been extended to the E-Series, including the new E5500. AutoSupport enhances customer service and speeds problem resolution by tracking configuration, performance, status, and exception data. With AutoSupport enabled, messages are sent either based on events or time (weekly, daily, other).

### Online Management

All management tasks can be performed while E-Series systems remain online with complete read/write data access. This allows you to make configuration changes and conduct maintenance without disrupting application I/O or scheduling planned downtime.

### Advanced Tuning

The E-Series includes advanced tuning functions so you can optimize performance with minimal effort. We've been handling storage for high-performance computing and other demanding applications for a long time, so we know which capabilities need to be tunable. You can tune any attribute to meet specific application needs. For example, if you initially configure a volume to use RAID 5 but later decide that RAID 10 would better suit your application, you can dynamically convert the volume with no disruptions.

Graphical performance tools give you several viewpoints into array I/O activity. The ability to view data in real time helps you make better-informed decisions.

### Data Protection

The E-Series provides a flexible data protection and disaster recovery architecture that includes enterprise features such as:

- **High-speed, high-efficiency snapshots** that let you protect data in seconds and reduce storage consumption by storing only changed blocks
- **Synchronous mirroring** for no-data-loss protection of business-critical data
- **Asynchronous mirroring** for long-distance protection and compliance with business requirements

### Dynamic Disk Pools (DDP)

The E-Series includes two features that are ideal for transaction-oriented workloads: Dynamic Disk Pools and SSD Cache. (SSD Cache is discussed in the section that follows this one.)

Dynamic Disk Pools increase the level of data protection, provide more consistent transactional performance, and improve the versatility of E-Series systems. DDP dynamically distributes data, spare capacity, and parity information across a pool of drives. An intelligent algorithm (seven patents pending) determines which drives are used for data placement, and data is dynamically recreated and redistributed as needed to maintain protection and uniform distribution.

### Consistent Performance During Rebuilds


DDP minimizes the performance drop that can occur during a disk rebuild, allowing rebuilds to complete up to eight times faster than with traditional RAID. That translates to your storage spending more time in an optimal performance mode to maximize application productivity.

Shorter rebuild times also reduce the possibility of a second disk failure occurring during a disk rebuild and protects against unrecoverable media errors. Stripes with several drive failures receive priority for reconstruction.

Overall, DDP provides a significant improvement in data protection; the larger the pool, the greater the protection.

### How DDP Works

With traditional RAID, when a disk fails, data is recreated from parity on a single, hot spare drive, creating a bottleneck. All volumes using the RAID group suffer. With DDP, each volume's data, parity information, and spare capacity are distributed across all drives in the pool. When a drive fails, data is reconstructed throughout the disk pool, so no single disk becomes a bottleneck.



**Figure 6)** When a disk fails in a Dynamic Disk Pool, reconstruction activity is spread across the pool. Rebuilds complete up to eight times faster.

### Increased Versatility

DDP provides flexible disk pool sizing to optimize shelf utilization. There are several ways to implement pools. A single pool for all volumes maximizes simplicity, protection, and utilization. Smaller pools with one volume/pool maximize performance for applications that need maximum bandwidth and for clustered file systems. You can create numerous pools to meet different requirements, and you can intermix traditional RAID and DDP.

### SSD Cache

---

SANtricity® SSD Cache is designed to accelerate random I/O for transactional workloads, similar to the [Flash Pool™](#) intelligent caching used in FAS systems. SSD Cache automatically caches data blocks on SSD in real time, without the need for policy engines or scheduled data migration. The minimum cache is a single SSD drive; the maximum is 5TB per E-Series array. SSD Cache can be shared by any or all volumes on an E-Series system.

SSD Cache has several design optimizations to provide greater flexibility for high-performance applications:

- Cache block sizes are adjustable from 2K to 8K. Internal tests have shown that a properly tuned E-Series cache will be populated up to 500% faster. This is particularly important in applications in which the working set changes frequently such as data analytics.
- New writes can be cached immediately on SSD or written to HDD only. Some applications tend to read data back after writing. Others write data and don't read it back until much later, in which case caching writes are a waste.

Optimizing the cache based on the read/write characteristics of an application maximizes available space on SSD.

Workloads tested with SANtricity SSD Cache show up to a 700% increase in IOPS over the same arrays without cache.

## Conclusion

---

The new E5500 and the rest of the E-Series line are ideal for big data, analytics, and HPC applications that require dedicated storage with maximum bandwidth and minimum footprint. The E5500 delivers double the performance of the previous generation and sets a new bar for price/performance, performance per disk spindle, and more. The E-Series line offers proven reliability, availability, and serviceability along with simplified management, tuning, and advanced data protection features. Dynamic Disk Pools and SSD Cache can help you optimize performance for transaction-oriented workloads that generate random I/O.

---



**By Michael Jastad, Sr. Technical Marketing Engineer**

Michael began his career in the military working as a NATO cryptologist stationed in Naples, which led him to pursue a degree in computer science. After college, he filled engineering roles at companies including Fujitsu, Pyramid Technologies, Sequent Computers, and IBM, where he worked on the Tivoli product line. He joined LSI as a developer, HPC technologist, and solutions architect and came to NetApp with the E-Series line. His current responsibilities include developing best practices to help customers and partners and creating white papers on E-Series solutions.

---

### *Quick Links*

- › [Tech OnTap Community](#)
  - › [Archive](#)
  - › [User Groups](#)
  - › [PDF](#)
-