



Technical FAQ

EF-Series Flash Array Technical FAQ

December 2014

DESCRIPTION

This document provides information and answers to common questions regarding the EF-Series flash array. It is intended for NetApp direct sales teams only.

TABLE OF CONTENTS

1	OVERVIEW	5
	What is the EF-Series flash array?	5
	Which workload characteristics and profiles will benefit from using the EF-Series?	5
	How does the EF-Series perform?	5
	What are the key differences between the EF550 and the new EF560?	6
	What's new with the SANtricity Storage Manager 11.20 and the SANtricity OS 8.20 release?	7
	When do I position an EF560 over an EF550?	8
	Does an EF550 customer get a performance boost with the latest SANtricity release?	8
	What are my configuration options for the EF-Series?	8
	Is the EF-Series a block or file device?	8
	Can I expand beyond the initial 12 or 24 drives? And is it nondisruptive?	8
	Do we support both FC and iSCSI in a single system?	8
	What's included in the EF560?	9
	What's included in the EF550?	9
	What's included in the fixed EF540 configuration?	9
	Does the EF-Series include SANtricity software?	10
	Why is synchronous mirroring not included?	10
	What is the raw and usable capacity for the EF-Series?	10
	Are there any performance differences between the different supported interfaces?	11
	Which warranty options are available?	12
	Which professional services are available?	12
	Does the EF array fit in a standard 19" rack cabinet?	12
	Does the EF-Series offer NetApp AutoSupport?	12
	Is there a roadmap for the EF-Series?	12
2	SANTRICITY SOFTWARE	12
	Which storage operating system does the EF-Series run?	12
	Does the EF-Series require any special hardware to manage the array?	12
	What has been done in SANtricity to optimize it for the EF-Series?	13
	Do we configure an EF-Series just as we do a FAS?	13
	What is the volume (LUN) size limit?	13
	What is the maximum number of volumes that can be configured on the EF-Series?	14
	What is the limit on the number of volumes (LUNs) that can be assigned to one host?	14
	What are the various limits to SANtricity software?	14
	Are there any performance impacts when using the advanced features?	14
	Does SANtricity have snapshot, clone, or remote replication features?	15
	Can SANtricity dynamically expand volumes (physically, logically, plus restripe)?	15
	Does SANtricity offer compression or deduplication functionality?	15
	Does SANtricity offer alerting capabilities?	15
	Does SANtricity have a performance monitor?	16
	Which metrics are collected by the performance monitor?	16
	What is the maximum number of days I can collect performance data?	17
3	HIGH AVAILABILITY AND FAULT TOLERANCE	17
	Is the EF-Series a fully redundant, highly available array?	17
	Does the EF-Series support batteries to protect data in cache?	17
	Which RAID levels are supported?	17

Does the EF-Series support Dynamic Disk Pools?	17
Does the EF-Series support hot spares?	17
Is the controller firmware nondisruptive to attached hosts?	18
Does the firmware upgrade procedure affect performance?	18
4 EF-SERIES ENVIRONMENTALS.....	18
What are the EF-Series' power and cooling requirements?	18
What are the EF-Series' dimensions and weight?.....	18
What are the temperature and humidity specifications?.....	19
5 SOLID-STATE DISK POSITIONING AND OVERVIEW	20
Which drives are supported in the EF-Series?	20
What is the predicted wear-out period for the SSDs?	20
What is the SSD's endurance?	20
What kind of monitoring is available on the EF array to track the SSDs' wear life?	22
What is the difference between SLC, MLC, eMLC, and cMLC technologies?.....	22
Does the EF array support T10 protection information and data assurance?	22
Does the EF array support secure erase functionality?.....	22
What is TRIM/UNMAP and does the EF array support it?	23
What is write amplification?.....	23
How does the EF array manage wear-leveling and garbage collection?.....	23
For what period of time is a support agreement available for SSDs?.....	24
Once it's up and operational, what happens if an SSD drive fails?	24
6 PLATFORM POSITIONING	24
Is there a common management system for the EF-Series and FAS?	24
Is the EF array supported behind FlexArray (V-Series)?.....	24
Does the EF array support Flash Cache and Flash Pool?	24
Is there a way to move data between the EF array and FAS?	24
Are we developing reference architectures that include the EF array and FAS?	24
Are there any other reference architectures utilizing EF-Series?	25
Does the EF array support SANtricity advanced replication features?.....	25
7 RESOURCES	25
Which resources are available?	25
Is documentation available for the EF array and SANtricity software?.....	25
8 COMPETITION	25
What is our key advantage over our competitors in the all-flash-array market?	25
How do I compete against other all-flash arrays that claim 1 million IOPS?	26
What about competitors who say that the EF array isn't a ground-up flash design?	26
9 TRAINING	27
What training is available on the EF-Series array?	27

LIST OF TABLES

Table 1) Performance highlights	5
Table 2) Key similarities and differences	6
Table 3) End of life dates for EF540	9
Table 4) Raw capacity based on drives	10
Table 5) Capacity examples	10
Table 6) Interface Performance	11
Table 7) SANtricity software limits	14
Table 8) 12 drive power/cooling requirements	18
Table 9) 24 drive power/cooling requirements	18
Table 10) 48 drive power/cooling requirements	18
Table 11) Dimensions and weight	19
Table 12) Temperature and humidity information	19

1 Overview

What is the EF-Series flash array?

Answer: The NetApp® EF-Series is an all-flash storage array that brings together extreme performance and enterprise-grade reliability to create a system optimized for latency-sensitive workloads. Building on a heritage of performance leadership, its core architecture has been proven in the world's most demanding and complex computing environments. Its field-proven design is the culmination of 20 years of industry knowledge focused on designing enterprise-class storage. Leveraging the experience gained from over 850,000 systems shipped, architects have enabled the fully redundant EF-Series all-flash array to provide the highest levels of reliability, availability, and data protection.

The EF-Series runs on the enterprise-proven NetApp SANtricity® software platform. SANtricity is designed to combine sustainable, low-latency performance with enterprise-class availability and protection. Its streamlined firmware is well suited for the extreme demands of latency-sensitive workloads. And SANtricity helps keep I/O flowing with automated path failover, online administration, advanced data protection, proactive monitoring and repair, nondisruptive upgrades, and extensive diagnostic capabilities.

Which workload characteristics and profiles will benefit from using the EF-Series?

Answer: In general, any database-driven application will benefit from the EF array's ultrahigh IOPS and submillisecond latency. Examples include real-time analytics, online transaction processing (OLTP), and business intelligence.

How does the EF-Series perform?

Answer: Although we can give you and our customers “brochure” performance numbers, what's most important is how the EF-Series performs in the customer's environment. Focused development and years of experience supporting tier 1 database environments enable the EF-Series to deliver consistent, sustainable, low-latency performance under varying workloads and conditions.

The Field Portal contains two technical reports—TR-4116i, on the EF540; and TR-4294, on the EF550; a TR on the EF560 is in process—with detailed information on performance testing and results:

<https://fieldportal.netapp.com/ef-series.aspx>.

Here are a few highlights.

Table 1) Performance highlights.

	EF550	EF560
Burst I/O rate – cache reads (512B)	900,000 IOPS	900,000 IOPS
RAID 5, 4K blocks Sustained I/O rate – random disk reads	450,000 to 500,000 IOPS <1ms	650,000 IOPS @ 0.78ms 640,000 IOPS @ 0.40ms 500,000 IOPS @ 0.26ms
RAID 5, 8K blocks Sustained I/O rate – random disk reads	450,000 to 500,000 IOPS <1ms	645,000 IOPS @ 0.78ms 628,000 IOPS @ 0.41ms 467,000 IOPS @ 0.27ms

RAID 5, 4K blocks Sustained I/O rate – random disk writes	75,000 IOPS @ 1ms	120,000 IOPS @ 2.13ms 112,000 IOPS @ 0.85ms 107,000 IOPS @ 0.44ms
RAID 5, 8K blocks Sustained I/O rate – random disk writes	75,000 IOPS @ 1ms	117,000 IOPS @ 2.18ms 111,000 IOPS @ 0.85ms 101,000 IOPS @ 0.47ms
RAID 5, 8K blocks Sustained I/O rate – random 50/50 read/write	75,000 IOPS @ 1ms	197,000 IOPS @ 0.97ms 182,000 IOPS @ 0.51ms 142,000 IOPS @ 0.33ms
Sustained throughput – sequential disk reads (512K)	12GB/s	12GB/s
Sustained throughput – sequential disk write (CMD 512K)	10GB/s	10GB/s
Sustained throughput – sequential disk write	6GB/s	6GB/s

All of these results were achieved in a fully redundant, highly available RAID 5 configuration running Iometer. There were no tricks and no benchmark specials. These are just real-world numbers designed to be achievable in a custom environment.

What are the key differences between the EF550 and the new EF560?

Answer: Key similarities and differences effective as of December 16, 2014, are outlined in the following table

Table 2) Key similarities and differences.

Dual-Controller Attributes	EF550	EF560
Base models	6 400GB SSDs (2.4TB) Can start with any capacity point	6 400GB SSDs (2.4TB) Can start with any capacity point
Expandable	Single or multiple SSD increments Up to 120 drives	Single or multiple SSD increments Up to 120 drives
SSDs	400GB, 800GB, 1.6TB SSDs	400GB, 800GB, 1.6TB SSDs

Dual-Controller Attributes	EF550	EF560
I/O interface options	Options: 8 16Gb FC 8 6Gb SAS 8 12Gb SAS 8 10Gb iSCSI 4 40Gb InfiniBand 4 56Gb InfiniBand	Options: 8 16Gb FC 8 12Gb SAS 8 10Gb iSCSI 4 56Gb InfiniBand
Backend drive ports	4 6Gb SAS	4 6Gb SAS
System cache	24GB (12GB per controller)	24GB (12GB per controller)
Storage software	SANtricity Storage Manager 11.20 NetApp SANtricity OS 8.20	SANtricity Storage Manager 11.20 NetApp SANtricity OS 8.20

The EF560 is a new controller with a six-core Intel® Xeon® E5-1428Lv2 2.2GHz CPU; the EF550 has a four-core Xeon E5-2418L 2.0 GHz CPU. With the SANtricity 8.20 (that is the code version on the controllers as opposed to the SANtricity 11.20 management software version) we move from using only a single core to using multiple cores. The performance increase is largely due to the additional core as well as to the increased power of the CPU.

In addition to the CPU, the EF560 will also have the 12G SAS and 56G IB connections (the EF550 gets these as well) but will not carry forward the slower 6G SAS and 40G IB protocols.

What's new with the SANtricity Storage Manager 11.20 and the SANtricity OS 8.20 release?

Answer:

- Multicore support
 - EF550/EF560
 - Free performance boost for existing customers
- RAS features
 - Online drive firmware upgrades
 - Drive evacuator
 - Proactive drive monitoring that evacuates data from suspect drives before they fail—any required reconstruction picks up where the drive evacuation functionality stopped
 - Dynamic Disk Pools (DDP) drawer loss protection
 - DDP critical segment rebuild
- Other features
 - Support for online drive firmware upgrades—Caution: Consult NetApp Customer Success Services (CSS) before upgrading drive firmware
 - Support for changing the host I/O functionality between FC or iSCSI on all controller models in the field

When do I position an EF560 over an EF550?

Answer: In most cases you will want to position the EF550 over the EF560. Here are some important facts to understand when making this decision.

- There is no price delta between the platforms.
 - For same configurations
 - Higher-speed host interface card (HIC) configurations will have higher pricing.
 - Example: EF550 40G IB versus EF560 56G IB
- Both systems run the same OS, but the EF550 is capable of staying on the previous version of code.
- The EF560 has higher performance.
- The EF550 EOL is planned for December 2015.
- Higher port speeds in the EF560 will autonegotiate down to match network speeds.
- Customers unable to support the iSER protocol over 56Gb IB will need to use the EF550 with 40Gb IB to retain SRP protocol support.
- Moving to 12Gb SAS host connections will require replacement of 6Gb SAS cabling. Cabling for IB may also need to be replaced when moving from 40Gb to 56Gb.

Likely cases for positioning the EF550 include:

- The customer has qualified the EF550 and is not interested in going through another qualification process, either hardware or software.
- You do not want to put the account at risk by forcing a new qualification process.
- The customer is running EF550 and wants to stay consistent in platform deployments.

Does an EF550 customer get a performance boost with the latest SANtricity release?

Answer: Yes. Just load up the SANtricity 11.20 bundle (which includes the 8.20 SANtricity OS release) and the customer has additional performance. But the customer environment must be capable of pushing additional IOPS for any effect to be seen.

What are my configuration options for the EF-Series?

Answer: Beginning with the May 2014 Quote Tool release, the EF550 and the EF560 support configure-to-order options. The EF-Series can have as few as 6 and as many as 24 drives in the initial shelf (along with the controllers). Supported drives are 400GB, 800GB, and 1.6TB SSDs. Additionally there is a “full disk encryption” (FDE) option for the 800GB SSD. Additional 24-SSD expansion shelves can be added to any initial configuration.

The EF540 is offered only in two fixed configurations based on 800GB SSDs in a 2U/24 shelf: 9.6TB (12 drives) and 19.2TB (24 drives). Additional 24-SSD expansion shelves can be added to any initial configuration.

Is the EF-Series a block or file device?

Answer: The EF is a block-level storage system.

Can I expand beyond the initial 12 or 24 drives? And is it nondisruptive?

Answer: Yes and yes. EF-Series systems support a 24-drive expansion shelf. Up to four shelves can be added nondisruptively. Systems can be expanded up to 120 drives.

Do we support both FC and iSCSI in a single system?

Answer: This option is available only with the EF540, which has entered EOL except for the option that supports both FC and iSCSI. Starting in November 2014, the following E-Series platforms can no longer be ordered: E26xx, E54xx, and EF54x series systems. Newer platforms are currently available and shipping.

At the request of the field, the availability of E5400 and EF540 systems with 10Gb iSCSI host interface cards (HICs) is being extended. These parts will be available through the restricted Quote Tool. The base FC port on the E5400/EF540 controllers combined with the iSCSI HIC provide the requested dual-protocol functionality. The product (E5400/EF540 + iSCSI HIC) is available in the restricted Quote Tool only through CQ4-2015 and must be ordered through a NetApp SSR. Additional details can be found in the EOA bulletin and ordering guide (see below). These publications are also available on the Field Portal.

- The updated EOA bulletin is available on the Field Portal.
- The E5400/EF540 ordering guide is also available on the Field Portal.

The reason for the extension is because some customers require the ability to run both FC and iSCSI protocols in the same system. This capability exists in the E5400/EF540 platform with the combination of FC base ports and an iSCSI HIC. Currently shipping systems, the E5500/EF550, do not have any base ports and cannot support this capability. The ability to run multiple protocols is planned for a later release of SANtricity.

Table 3) End-of-life dates for EF540.

Activity	Key Date	Activity	Key Date
All E2600 and most E5400 and EF540 systems, end of availability announcement	14-Jul-2014	E5400 and EF540 systems with iSCSI HIC only, end of availability announcement	10-Nov-2014
Last order	10-Nov-2014	Last order	14-Dec-2015
Last ship	31-Dec-2014	Last ship	15-Jan-2016
End of SW support*	31-Dec-2019	End of SW support*	15-Jan-2021
End of HW support	31-Dec-2019	End of HW support	15-Jan-2021

*SANtricity 11.20 will be the last supported release for E2600, E5400, and EF540 series systems.

What's included in the EF560?

Answer: The EF560 includes a 2U/24-drive shelf, 2 EF560 controllers, 24GB of controller cache (12GB per controller), 6 to 24 SSDs, and SANtricity Storage Manager software. You must specify a host interface at the time of order. The choices are 16Gb FC (four ports per controller), 10Gb iSCSI (four ports per controller), 12Gb SAS (four ports per controller), and 56Gb InfiniBand (four ports per controller). There is an additional charge for country-specific power cords.

What's included in the EF550?

Answer: The EF550 includes a 2U/24-drive shelf, 2 EF550 controllers, 24GB of controller cache (12GB per controller), 6 to 24 SSDs, and SANtricity Storage Manager software. You must specify a host interface at the time of order. The choices are 16Gb FC (four ports per controller), 10Gb iSCSI (four ports per controller), 12Gb or 6Gb SAS (four ports per controller), and 56Gb or 40Gb InfiniBand (four ports per controller). There is an additional charge for country-specific power cords.

What's included in the fixed EF540 configuration?

Answer: The EF540 includes a 2U/24-drive shelf, 2 EF540 controllers, 24GB of controller cache (12GB per controller), 8 8Gb FC I/O ports (4 ports per controller), 12 or 24 800GB SSDs, and SANtricity Storage Manager Software. There is an additional charge for country-specific power cords.

Does the EF-Series include SANtricity software?

Answer: Yes. The EF-Series flash array includes SANtricity software, which is both the controller OS and the management client. Most “premium features” (also known as PFKs) are enabled by default. This includes NetApp Snapshot[®] copies, volume copy, and asynchronous remote mirroring. Synchronous mirroring is available at no charge through PVR. Full drive encryption (FDE) is also available for an additional fee.

Why is synchronous mirroring not included?

Answer: Synchronous mirroring was not included because the primary use case for an EF array is to deliver high IOPS with ultralow latency. Adding a synchronous mirroring operation to the environment would likely eliminate any advantage of implementing an EF array.

Given this use case, the synchronous mirroring feature is not automatically included so that the environment and customer expectations can be properly set and communicated to avoid encountering any issues when using the feature. If the customer and environment are satisfied with the limitations, the feature can be included through a PVR at no additional charge.

What is the raw and usable capacity for the EF-Series?

Answer: Raw capacity is calculated by multiplying the number of drives by the drive size.

Table 4) Raw capacity based on drives.

# SSDs	SSD Capacity (TB)		
	400GB	800GB	1.6TB
6	2.4	4.8	9.6
12	4.8	9.6	19.2
18	7.2	14.4	28.8
24	9.6	19.2	38.4
36	14.4	28.8	57.6
48	19.2	38.4	76.8
72	28.8	57.6	115.2
96	38.4	76.8	153.6
120	48.0	96.0	192.0

Usable capacity can vary greatly based on RAID level and drive per RAID set. The simplest way to determine usable capacity is to multiply the number of *data* drives by 744.5GB, which is the usable capacity of a configured 800GB SSD. For example, a configuration with 10 data drives (10+1) would be 7.27TB (10 x 744.5) GB = 7445GB / 1024 = 7.27TB.

Below are a few additional examples based on the 800GB SSD.

Table 5) Capacity examples.

System	RAID Level	RAID Set(s)	Hot Spares	Usable Capacity
12-drive	5	5+1 (two)	0	7.27TB
12-drive	5	10+1	1	7.27TB
12-drive	6	9+2	1	6.55TB
12-drive	DDP	1 DDP	1	6.76TB

System	RAID Level	RAID Set(s)	Hot Spares	Usable Capacity
24-drive	5	5+1 (four)	0	14.54TB
24-drive	5	7+1 (three)	0	15.27TB
24-drive	5	8+1 (two) & 4+1	1	14.54TB
24-drive	6	6+2 (three)	0	13.09TB
24-drive	6	6+2 (two) & 5+2	1	12.36TB
24-drive	DDP	1 DDP	1	14.15TB
48-drive	5	7+1 (six)	0	30.54TB
48-drive	5	7+1 (five) & 6+1	1	29.82TB
48-drive	6	6+2 (six)	0	26.18TB
48-drive	6	8+2 (four) & 5+2	1	26.91TB
48-drive	DDP	1 DDP	2	28.90TB

To see usable capacity for a specific configuration, a capacity sizing tool is available on the Field Portal at <https://fieldportal.netapp.com/Core/DownloadDoc.aspx?documentID=88880&contentID=117029>.

Are there any performance differences between the different supported interfaces?

Answer: Yes. The different interfaces have various operational differences, which impact performance. The best way to size and understand these differences is to use the standard sizing tool on the web. Depending on the customer's requirements you may meet or surpass the customer's expectations with any of the interfaces. The chart below is a general guideline.

Table 6) Interface performance.

Controller	Metric	SAS 12Gbps	FC 16Gbps	IB 56Gbps	iSCSI 10Gbps
EF560	Sustained I/O disk read	650,000 <1ms	650,000 <1ms	530,000 <1ms	475,000 >1ms
	Sustained I/O disk write	100,000	100,000	90,000	90,000
EF550	Sustained I/O disk read	500,000	500,000	425,000	375,000
	Sustained I/O disk write	75,000	75,000	70,000	75,000
Both	Sustained throughput disk write (CMD 512KB)	9,000	9,000	9,000	5,500
	Sustained throughput disk write (CME 512KB)	6,000	6,000	6,000	4,500
	Burst I/O rate cache reads (512KB)	900,000	900,000	900,000	520,000

Which warranty options are available?

Answer: All standard NetApp warranty options are available.

Which professional services are available?

Answer: All standard professional services are available.

Does the EF array fit in a standard 19" rack cabinet?

Answer: Yes. The EF array is designed for a standard 19" rack and requires just two rack units. The 24-drive expansion shelf is also designed for a standard 19" rack and requires just two rack units.

Does the EF-Series offer NetApp AutoSupport?

Answer: Yes.

Is there a roadmap for the EF-Series?

Answer: Absolutely there is. NetApp is committed to the EF platform and multiple enhancements are coming. Highlights include:

- Support for more cores
- Support for multiprotocol within the controllers
- Cloud vaulting support

2 SANtricity Software

Which storage operating system does the EF-Series run?

Answer: The EF runs enterprise-proven SANtricity software. SANtricity's streamlined firmware and extensive tuning capabilities are optimized for high performance, making it a great match for the EF array.

Two parts are needed to run the system:

- SANtricity Storage Management. This is the management interface, currently at version 11.20, and normally what is referred to when the software version is discussed. The higher-level version can communicate with lower-level SANtricity OS versions.
- The SANtricity OS. This is the software loaded on the actual array controllers and is currently at version 8.20. The management software must be at least at this level of the controllers to communicate with a storage array.

Does the EF-Series require any special hardware to manage the array?

Answer: No. SANtricity is a Java[®]-based management framework that can communicate with the system either in-band over the data transfer connections and/or out-of-band over Ethernet. SANtricity is a lightweight management client and can be installed on most Windows[®]-based PC/laptops or Linux[®] desktop environments.

What has been done in SANtricity to optimize it for the EF-Series?

Answer: In a sense, SANtricity has been building toward flash for years. From the beginning, the design goal was to get I/O to and from disk as quickly and efficiently as possible. The result is streamlined firmware and extensive configuration flexibility that enable our high-performance platform to perform exceptionally well across a wide range of workloads—from the world's largest databases to the world's largest supercomputers and everything in between.

With that said, we made several improvements specifically for SSDs to improve both the performance and the system reliability:

- Dual-active ports. Most enterprise-class drives are dual-ported; however, those ports are configured as active-passive, meaning that one is the primary port and the other is there for redundancy/failover. On the EF array, the SSD ports are active-active and SANtricity deliberately balances the I/O load across both SSD ports, significantly improving the performance (throughput and latency) of each individual drive.
- Increased queue depth. SANtricity supports four times greater queue depth for SSDs than for SAS drives (15K, 10K, and 7.2K). This increase in queue depth allows more concurrent access operations to/from the SSD without sacrificing latency, which typically results in better LUN performance.
- SSD wear-life monitoring. Flash has a limited wear life before individual memory locations can no longer reliably persist data. To protect against data loss, SANtricity proactively monitors SSD wear life and can advise when a drive is reaching the end of its useful life and should be replaced. Upon receiving this alert, NetApp will replace a drive under warranty prior to any failure. The current status of each SSD can be seen at any time with SANtricity.

We continually look for ways to further optimize the overall system architecture (firmware, management, hardware).

Do we configure an EF-Series just as we do a FAS?

Answer: No. SANtricity does not have the concept of aggregates and it is rarely a best practice to put all SSDs in the EF array into a single RAID set. The key concepts within SANtricity are:

Volume group. A set of drives that is logically grouped and assigned a RAID level. Supported RAID levels are 0, 1, 3, 5, 6, 10, and Dynamic Disk Pools (DDP). Each volume group created provides the overall capacity needed to create one or more volumes.

Volume. The logical component created for the host to access storage on the EF array. A volume is created from the capacity available on a volume group. It can be all or part of the volume group capacity.

Additional insight into SANtricity can be found in the "SANtricity Concepts Guide," as well as in two great courses from NetAppU:

- "Configuring E-Series: What FAS Experts Need to Know":
<https://www.brainshark.com/netapp/vu?pi=zHPzYwy7tzB2XBz0&intk=510403940>
- "The Best of Both Technologies: Comparing Data ONTAP and E-Series"
<http://learningcenter.netapp.com/LC?ObjectType=WBT&ObjectID=00230069>

What is the volume (LUN) size limit?

Answer: A volume/LUN can be anything from 10MB (minimum) to the full size of the containing volume group. The theoretical maximum volume size is 2 exabytes.

What is the maximum number of volumes that can be configured on the EF-Series?

Answer: The maximum number of supported volumes/LUNs on an EF array is 2,048.

What is the limit on the number of volumes (LUNs) that can be assigned to one host?

Answer: The maximum number of volumes/LUNs that can be assigned to one host is 256.

What are the various limits to SANtricity software?

Answer: Below is a table listing SANtricity software limits.

Table 7) SANtricity software limits.

	Attribute	Maximum
System	Volumes	2,048
	Thin volumes	2,048
	Partitions	512
	Disk pools	10
	Consistency groups	16
	Volumes per consistency group	64
Snapshot Copies	Snapshot copies	2,048
	Snapshot copies per volume	128
	Snapshot volumes per Snapshot copy	4
	Snapshot groups	1,024
	Snapshot copies per Snapshot group	32
	Snapshot groups per volume	4
Async Mirrors	Mirrors	128
	Mirrors per volume	1
	Mirrors per async mirror group	64
	Async mirror groups	4

Are there any performance impacts when using the advanced features?

Answer: As in most cases, the answer is “It depends,” but generally there is some performance impact when using any of the advanced features. SANtricity is a real-time operating system designed for high performance and ultralow latency operations. It was not designed to accommodate advanced features

and maintain the exceptional performance it's known for. If no storage management features are used, that part of the code is never executed and has no impact.

There can be a significant effect on performance when using the storage management features; the impact will depend on the workload. Following are three examples.

1. Snapshot copies, for example, have an impact on write operations. The heavier the write operations in the environment, the larger the impact. This can be as small as 5% or as large as 20%, and perhaps even higher depending on the overall load on the system. Reads that depend on a write taking place can also be affected.
2. Thin provisioning is another feature that can have a significant performance impact. When a volume is extended the system must first initialize the new space. We have seen 10–30% performance drops and increased latency in these environments. Again, this will vary by workload because read operations do not cause a volume expansion operation.
3. Asynchronous mirroring is a function that has minimal effect. Write operations are tracked separately from the primary volume and therefore do not have the same impact as a snapshot operation. The performance impact for an asynchronous mirroring operation is typically between 5% and 10%. There is a strong use case for this because the EF can be protected by asynchronously mirroring to an E-Series platform with hard drives, providing great cost savings and a DR option.

Does SANtricity have snapshot, clone, or remote replication features?

Answer: Yes. SANtricity software supports both a snapshot capability as well as volume copy (that is, cloning) within a given system. For remote replication, the EF array is capable of mirroring to another EF array or E-Series system.

Can SANtricity dynamically expand volumes (physically, logically, plus restripe)?

Answer: Yes. SANtricity's Dynamic Volume Expansion (DVE) feature increases the capacity of a standard volume or a snapshot repository volume. DVE uses the free capacity available on the volume group of the standard volume or the snapshot repository volume.

In addition, capacity with a volume group (that is, a RAID set) also can be dynamically (nondisruptively) expanded by using the SANtricity Dynamic Capacity Expansion (DCE) feature. This modification operation allows you to add unassigned drives to a volume group. Adding unassigned drives increases the free capacity in the volume group. You can use this free capacity for DVE or to create additional volumes.

Both Dynamic Volume Expansion and Dynamic Capacity Expansion are considered to be dynamic because you have the ability to continually access data in the volume group throughout the entire operation. Additionally, both operations automatically restripe data across the new containing space.

Does SANtricity offer compression or deduplication functionality?

Answer: It does not within the controller, but this can be added through third-party offerings. We currently have a partnership agreement with Permabit to provide this capability. The EF typically relies on the server OS to provide this capability, for example, through Oracle®.

Does SANtricity offer alerting capabilities?

Answer: Yes. SANtricity can be configured to send e-mail notifications or SNMP notifications when an event occurs in a storage array. The notification summarizes the event and details about the affected storage array, including these items:

- The name of the affected storage array
- The host IP address (for an in-band managed storage array)
- The host name and ID (shown as out-of-band if the storage array is managed through the Ethernet connection of each controller)
- The event error type related to an Event Log entry
- The date and the time when the event occurred
- A brief description of the event

Does SANtricity have a performance monitor?

Answer: Yes. The performance monitor within SANtricity provides visibility into performance activity across your monitored storage devices. You can use the performance monitor to perform these tasks:

- View in real time the values of the data for a monitored device. This capability helps you to determine if the device is experiencing any problems.
- See a historical view of a monitored device to identify when a problem started or what caused a problem.
- Specify various reporting attributes, such as time increments and filtering criteria, to examine performance trends and to pinpoint the cause of availability and performance issues.
- Show data in tabular format (actual values of the collected metrics) or graphical format (primarily as line graphs) or export the data to a file.
- Show a dashboard where you can monitor six different performance graphs at once.
- Monitor I/O latency for drives, volumes, and throughput for the entire storage array and volumes.
- Start textual performance monitoring or background performance monitoring.

Which metrics are collected by the performance monitor?

Answer: Metrics are measurements of the data that the performance monitor collects from the storage devices that you monitor. Metrics help to pinpoint problems and define their cause. Metrics define the types of data that you collect as well as the type of data source from which you collect the data. With the SANtricity performance monitor you can collect the following metric data:

- Total I/Os—The total I/Os performed by this device since the beginning of the polling session.
- Read percentage—The percentage of total I/Os that are read operations for this device. Write percentage can be calculated as 100 minus this value.
- Cache hit percentage—The percentage of total I/Os that are processed with data from the cache rather than requiring a read from the drive.
- Current or maximum KBs per second—The current or maximum transfer rate during the current polling interval. The transfer rate is the amount of data in kilobytes (Table view) or megabytes (Graphical view) that can be moved through the I/O data connection in a second (also called throughput).
- Current or maximum I/O per second—The current or maximum number of I/O requests serviced per second during the current polling interval (also called an I/O request rate).
- Throughput performance—The performance measure for the entire storage array and for individual volumes (max-average-min).
- I/O latency—The I/O performed.

Metric sources are the aspects of a storage array or a controller that provide data. You can configure the performance monitor to report data from the following metric sources:

- Volume, volume group, disk pools, controller, and storage array

What is the maximum number of days I can collect performance data?

Answer: Within SANtricity, the maximum length of historical performance collection is seven days. This data can be saved to alternative storage for an additional length of time if required.

3 High Availability and Fault Tolerance

Is the EF-Series a fully redundant, highly available array?

Answer: Yes. When properly configured, the EF array is a fully redundant, highly available array. It includes dual controllers; dual power supplies; dual fans; RAID protection; hot-swappable components; nondisruptive configuration, administration, and expansion/service; and support for multipath drivers and failover.

Does the EF-Series support batteries to protect data in cache?

Answer: Yes. Each controller provides a set of batteries to protect the data in cache. After a complete loss of power, the controller batteries provide power to the chassis for a limited time. During this time, all “dirty” write data in the controller cache is flushed to a nonvolatile storage medium housed within the controller canister. After this action has completed, the controller or controllers gracefully shut down.

Which RAID levels are supported?

Answer: SANtricity offers RAID levels 0, 1, 10, 5, and 6. These may all be used concurrently and each is allowed with different numbers of drives:

- RAID 0—1 drive to all drives in the system
- RAID 1 and 10—Minimum of 2 drives to all drives in the system
- RAID 5—Minimum of 3 drives to a maximum of 30 drives
- RAID 6—Minimum of 5 drives to a maximum of 30 drives

In SANtricity, a set of drives that is logically grouped and assigned a RAID level is called a volume group.

Does the EF-Series support Dynamic Disk Pools?

Answer: Yes. Starting with SANtricity 11.10, support of Dynamic Disk Pools technology, or DDP, is included. This feature allows storage administrators to simplify RAID management, improve data protection, and maintain predictable performance under all conditions. DDP evenly distributes data, protection information, and spare capacity across the EF-Series' entire pool of drives, simplifying setup and maximizing utilization. Its next-generation technology eliminates hot spots, minimizes the performance impact of a drive failure, and can return the system to optimal condition up to eight times more quickly than traditional RAID. With shorter rebuild times and patented prioritize reconstruction technology, DDP offers a level of data protection that simply can't be achieved with traditional RAID.

Does the EF-Series support hot spares?

Answer: Yes. Hot spare SSDs are assigned by the user. NetApp generally recommends at least one, but there is no specific requirement for any to be defined. When a failed drive has been physically replaced, a copyback operation automatically occurs from the hot spare drive to the replacement drive.

Is the controller firmware nondisruptive to attached hosts?

Answer: Yes. Controller firmware on the EF array may be completed nondisruptively to attached hosts with the assumption that there are connections to both controllers and the proper multipathing driver is configured and functional.

Does the firmware upgrade procedure affect performance?

Answer: There will be minimal performance impact on the attached hosts. Essentially during a firmware upgrade each controller is independently rebooted, thereby shifting ownership of volumes to the alternate controller for the duration of the reboot/upgrade cycle. NetApp expects that this process will require a minimal amount of time to complete.

4 EF-Series Environmentals

What are the EF-Series' power and cooling requirements?

Answer: Specifications are based on 220V input power and base controllers (no host cards).

Table 8) Twelve-drive power/cooling requirements.

12-Drive System	VA	WATTS	BTU/HR	AMPS
Total System Power (Normalized)	360.69	357.08	1218.41	1.64
Typical Power	375.72	371.96	1,269.18	1.71

Table 9) Twenty-four-drive power/cooling requirements.

24-Drive System	VA	WATTS	BTU/HR	AMPS
Total System Power (Normalized)	402.42	398.39	1,359.36	1.83
Typical Power	419.19	414.99	1,416.00	1.91

Table 10) Forty-eight-drive power/cooling requirements.

48-Drive System	VA	WATTS	BTU/HR	AMPS
Total System Power (Normalized)	634.61	628.25	2,143.67	2.89
Typical Power	661.05	654.43	2,233.00	3.01

Note: The EF array site-preparation guide has the following maximum specs for a 24-drive system. These likely result from some sort of failure scenario with additional headroom factored in.

- 567 VA, 567 watts, 1,935 BTU/hr

What are the EF-Series' dimensions and weight?

Answer: Both the 12- and 24-drive models, as well as the 24-drive expansion shelf, have the following dimensions:

Table 11) Dimensions and weight.

Dimensions and Weight	NetApp EF Array
Height	3.47" (8.81cm)
Width	19" (48.26cm)
Depth	19.6" (49.78cm)
Weight	12 SSDs: 59 lb (26.78 kg) 24 SSDs: 64 lb (29.07 kg)

What are the temperature and humidity specifications?

Answer: The table below contains temperature and humidity information

Table 12) Temperature and humidity information.

	Parameter	Requirement
Temperature	- Operating range - Maximum rate of change	50°F to 95°F (10°C to 35°C) 18°F (10°C) per hour
	- Storage range - Maximum rate of change	14°F to 122°F (-10°C to 50°C) 27°F (15°C) per hour
	- Transit range - Maximum rate of change	-40°F to 140°F (-40°C to 60°C) 36°F (20°C) per hour
Relative Humidity	- Operating range - Operating gradient	20% to 80% 50°F (10°C) per hour maximum
	- Storage range - Storage gradient	10% to 90% 50°F (10°C) per hour maximum
	- Transit range - Transit gradient	5% to 90% 68°F (20°C) per hour maximum
	- Maximum dew point - Maximum gradient	79°F (26°C) 10% per hour

Note: If you plan to operate a system at an altitude between 1,000 m and 3,000 m (3,280 ft to 9,842 ft) above sea level, lower the environmental temperature 1.7°C (3.1°F) for every 1,000 m (3,280 ft) above sea level.

5 Solid-State Disk Positioning and Overview

Which drives are supported in the EF-Series?

Answer: The EF-Series currently supports many SSDs. All SSDs are a medium-endurance SSD rated at 10 drive writes per day (DWPD) and warranted for up to 5 years. The SSDs are designed to offer solid I/O cost-to-performance benefits with an advanced innovative controller that uses a parallel processing architecture to deliver proven and tested superior reliability for storage and server environments in which quick access to data is vital. The supported SSDs have been specifically designed for demanding, mission-critical, 24/7 applications.

When a drive is selected it must meet certain expectations, both for endurance/lifespan and performance. Over the past several years there have been significant improvements in SSD drives. The drive technology used will most definitely change over time, but what will remain consistent is the commitment to delivering a reliable product delivering the fastest performance.

What is the predicted wear-out period for the SSDs?

Answer: The chart below shows various workloads and the lifetime expectancy, with those workloads having an 800GB SSD. Because there will likely be various workload patterns, we estimate that the 800GB SSD will last more than five years.

			Read Performance (MiB/s)	Write Performance (MiB/s)	Hourly P/E Growth	WAF	Lifetime (Years)	TBW in Lifetime
Drive Model	LB 806M	1MB Wr 70% 4KB Rd 30%	125.58	62.05	0.78	3.48	21.81	44757.57
Physical Density (MiB)	990,660	1MB Wr 50% 4KB Rd 50%	57.59	60.58	0.76	3.47	22.44	44960.68
Endurance (P/E)	150,000	1MB Wr 30% 4KB Rd 70%	39.73	95.09	1.21	3.51	14.12	44387.41
OverProvisioning (%)		1MB Wr 100%	0.00	181.26	2.33	3.54	7.35	44069.20
		2MB Wr 70% 4KB Rd 30%	121.67	56.02	0.70	3.46	24.29	44993.17
		2MB Wr 50% 4KB Rd 50%	74.20	82.20	1.05	3.52	16.29	44283.96
Physical Density (MB)	1,038,782	2MB Wr 30% 4KB Rd 70%	47.45	121.44	1.57	3.55	10.93	43909.82
		2MB Wr 100%	0.00	181.99	2.36	3.56	7.27	43739.29
		4MB Wr 70% 4KB Rd 30%	125.14	60.64	0.77	3.48	22.33	44767.28
		4MB Wr 50% 4KB Rd 50%	79.18	90.01	1.16	3.54	14.79	44013.22
		4MB Wr 30% 4KB Rd 70%	46.71	118.54	1.53	3.56	11.18	43812.38
		4MB Wr 100%	0.00	183.19	2.38	3.58	7.19	43582.29
		4kb_rnd_wr_qd16	0.00	3.92	0.33	23.35	51.42	6672.67
		8kb_rnd_wr_qd16	0.00	79.22	3.85	13.37	4.45	11654.71
		64kb_seq_wr_qd16	0.00	164.57	2.21	3.69	7.75	42183.80
		4kb_rnd_7030_qd16	53.37	22.86	2.00	24.05	8.57	6478.70
		8kb_rnd_7030_qd16	82.21	35.22	1.72	13.42	9.97	11614.47

- Read Performance: Observed read performance with given profile
- Write Performance: Observed write performance with given profile
- Hourly P/E Growth: Hourly Performance/Endurance Growth rating
- WAF: Write Amplification Factor (see below for more information)
- Lifetime: Years SSD will last with given profile
- TBW in Lifetime: TB written during the lifetime of the SSD

What is the SSD's endurance?

Answer: Let's begin with some terms and an overview of how an SSD increases its endurance.

Endurance is the amount of data that can be written to an SSD during its lifetime and is typically characterized by P/E (program-erase) cycles, TBW (total bytes written), or DWPD (drive writes per day on a single cell). Endurance essentially tells how long your SSD will last.

Until today, SSD drive endurance was a function of the component-level endurance of the flash technology used in the drive. High endurance was achieved by using SLC NAND flash, but SSDs based on this flash technology were cost prohibitive for most enterprise applications. Lower cost was achieved through the use of MLC NAND flash, but SSDs based on this technology could not be used for mainstream write-intensive workloads because they quickly wear out. To combat this, SSD manufacturers introduced the concept of overprovisioning into their SSD technology.

Solid-state disks reserve a portion of their total flash address space for overprovisioning, a percentage of the total physical memory reserved by the SSD and not part of the device's logical address space. The level of overprovisioning affects both write performance and endurance (operational lifetime). The higher the level, the better it is. As cells wear out, the SSD uses part of the overprovisioned space to replace the worn-out cell.

The HDD equivalent to this is called “short stroking” and the implications of the change are identical. The exact same device running the exact same firmware simply accesses a larger or smaller range of the physical address space.

SSD manufacturers can control the level of overprovisioning, which can be designed/targeted to address a range of workloads with which they communicate in the form of drive writes per day (DWPD).

There are three general classifications of endurance for SSDs:

- Write intensive: Implies 25+ DWPD
- Medium endurance: 10 DWPD
- Read intensive: 1–3 DWPD

The SSDs used in the EF-Series are medium endurance devices rated at 10 DWPD. This rating is measured with a Joint Electron Devices Engineering Council–specified workload that approximates SPC1. It is a mix of reads and writes and provides real-life “dwell” between consecutive program/erase cycles and also implies three months of powered-off data retention at the end of life. Your mileage may vary (more or fewer DWPD) if the client application's write pattern is more or less flash friendly (larger sequential aligned overwrites are better, and smaller random ones are worse). An array-level write amplification and tax also come with the RAID protection level selected.

The number of times data is written to the drive is finite and, with time, the flash begins to wear. As it does, more and more errors on the flash cause time-consuming error-recovery mechanisms, which in turn cause performance degradation.

The differences in performance and endurance that result from the changes in overprovisioning are all a “natural” result of the amount of reserved physical address space. There are no algorithmic changes; the drive just runs more efficiently when overprovisioning is higher.

The bottom line is that solid-state technology is changing: Drives are moving from 32nm to 24nm to 19nm. SSD drive vendors are moving away from the SLC and MLC language and moving toward endurance because it's more about how long the media will last. Additionally, the controllers add a buffering layer that further increases the life of an SSD. As the product evolves, more efficiencies will be included, allowing NetApp to continue to deliver high performance and low latency while using different, lower endurance devices and still offering outstanding warranties.

What kind of monitoring is available on the EF array to track the SSDs' wear life?

Answer: A flash-based SSD has a limited wear life before individual memory locations can no longer reliably persist data. The drive continuously monitors itself and reports its wear life status to the controller and can be found in SANtricity.

Two mechanisms exist to alert you that an SSD is nearing the end of its useful life: average erase count and spare blocks remaining. You can find these two pieces of information in the drive properties, which you can see in the storage management software by selecting a drive on the Hardware tab. These mechanisms indicate how much life remains in your SSDs.

Wear life monitoring:	Enabled
Average erase count:	0%
Spare blocks remaining:	100%

The average erase count is reported as a percentage of the rated lifetime. When the average erase count reaches 80%, schedule the replacement of the SSD. When the average erase count reaches 90%, a Needs Attention condition occurs. At this time, replace the SSD as soon as possible.

The spare blocks remaining are reported as a percentage of the total blocks. When the number of spare blocks remaining falls below 20%, schedule the replacement of the SSD. When the number of spare blocks remaining falls below 10%, a Needs Attention condition occurs. At this time, replace the SSD as soon as possible.

What is the difference between SLC, MLC, eMLC, and cMLC technologies?

Answer: SLC (Single Level Cell) SSDs only store bits of information on a single level per cell. This decreases power consumption and allows faster transfer speeds. This technology is typically reserved for higher-end or enterprise memory cards for which speed and reliability are more important than cost.

MLC (Multi-Level Cell) SSDs are a memory technology that stores bits of information in multiple levels in a cell. Because of this, MLC drives have a higher storage density and the per-MB manufacturing cost is less. However, there is a higher chance of error on the drive, which increases error checking process time.

eMLC and cMLC technologies are newer categories of the MLC SSD. The “e” is for enterprise and the “c” is for consumer. These letters simply indicate another level of grading within the MLC category. Any drive chosen by NetApp will meet our performance and endurance expectations prior to being released and supported.

Does the EF array support T10 protection information and data assurance?

Answer: Yes, we do support T10 protection information and data assurance on the SSD drives when an FC, SAS, or IB interface is used.

Does the EF array support secure erase functionality?

Answer: Yes, we have an 800GB SSD that is FDE capable.

What is TRIM/UNMAP and does the EF array support it?

Answer: TRIM (ATA command set) and UNMAP (SCSI command set) are commands intended for storage space reclamation primarily targeted at consumer flash-based devices, which do not typically support complex “garbage collection” routines (as opposed to enterprise-class SSDs, which do). These commands essentially communicate to consumer flash-based devices which pages/blocks can be freed up to allow potentially faster write performance and potentially increased wear life. This is also known as garbage collection.

The UNMAP (or TRIM) command is typically issued by a host operating system (for example, Windows 7) directly to the flash-based device. The host OS has direct visibility and knowledge of where the file system data is on the devices so those cells can be freed up.

In the case of the EF array, our storage controller sits between the OS and the SSDs and abstracts the data location through the RAID data layout. That and the fact that different SSDs behave, perform, and respond differently to the commands, as well as uncertainty as to the specific performance benefits in our SSD use case are reasons that the EF array does not currently support UNMAP.

TRIM/UNMAP with a conventional block-based RAID array is not practical, because trimming data out of the stripe incurs additional writes to the parity disk (to fix the parity). To perform a TRIM operation, you need to deallocate or trim an entire stripe's worth of data, something that isn't really under the control of a block-based array without an additional layer of indirection.

Additionally, in an EF array, we assume that the capacity will be fully utilized by the application. In practice, then, trimming is not practical or possible and designs that demand extra host free space to operate properly in effect apply a “hidden tax.”

Features and techniques are being explored for future releases to complement and enhance the enterprise-class SSD garbage collection routines (and improve upon the benefits of the consumer-based UNMAP command) but they have not been committed to at this time.

What is write amplification?

Answer: Write amplification is a byproduct associated with flash memory and solid-state drives for which the actual amount of physical information written is a multiple of the logical amount intended to be written. Because flash memory must be erased before it can be rewritten, the process to perform these operations results in moving (or rewriting) user data and metadata more than once. This multiplying effect increases the number of writes required over the life of the SSD, which shortens the time it can reliably operate. The increased writes also consume bandwidth to the flash memory, which mainly reduces random write performance to the SSD. Many factors affect the write amplification of an SSD; some can be controlled by the user and some are a direct result of the data written to and usage of the SSD.

All systems will have some write amplification associated with protecting the data on the SSDs. NetApp calculated this for our arrays, the expected lifespans of the drives, and performance expectations.

How does the EF array manage wear-leveling and garbage collection?

Answer: With the EF array, wear-leveling and garbage collection are managed within the SSD firmware.

The SSDs employ dynamic and static wear-leveling techniques that consider both wear and free space when selecting an erase block for garbage collection. Most of that work happens in the foreground as part of the initial write, based on an ordered list of erase blocks with the most free space. While this approach increases the per-command overhead slightly, it leads to very deterministic and sustainable write performance that does not vary with the recent history of writes (some designs that do background

garbage collection can write very fast if left idle for a while, but performance drops dramatically under sustained write loads that deplete the available free space).

There is also a background scanner running all the time that tries to find idle time for its work. Its main function is to look for erase blocks that are colder than average (fewer erase and write cycles) and to swap them for blocks that are hotter than average. Firmware thresholds are calculated to avoid too much "churn" and extra writes.

For what period of time is a support agreement available for SSDs?

Answer: Our standard support agreement for SSDs is for a three-year period. This can be extended to four or five years.

Once it's up and operational, what happens if an SSD drive fails?

Answer: The array supports many different RAID levels. Assuming that the disks were not configured with RAID 0 there will be no impact on the data. The user can simply replace the failed drive and the system will automatically rebuild the data on the new drive. If the system was configured with a hot spare, the system will automatically rebuild the failed drive using the hot spare.

A best practice is to configure the array with either RAID 5 or DDP, which will provide protection against a single- or dual-drive failure. Additionally, in a RAID configuration we suggest using a hot spare.

6 Platform Positioning

Is there a common management system for the EF-Series and FAS?

Answer: The two platforms have different management systems at this time: SANtricity software for the EF array (and E-Series) and NetApp OnCommand® System Manager for FAS systems.

Is the EF array supported behind FlexArray (V-Series)?

Answer: Yes.

Does the EF array support Flash Cache and Flash Pool?

Answer: No. NetApp Flash Cache™ and NetApp Flash Pool™ intelligent caching offer Virtual Storage Tiering functionality on FAS and FlexArray systems.

Is there a way to move data between the EF array and FAS?

Answer: Data movement between an EF array and FAS can be accomplished with the NetApp DTA2800 or third-party tools. Additionally, many databases provide data movement functionality as part of their management suite. In the case of database acceleration, we expect the database application to manage data movement between the EF array tier and the FAS tier.

Are we developing reference architectures that include the EF array and FAS?

Answer: Yes. NetApp is developing reference architectures for key business applications for which a combination of an EF array and FAS is used to enable a solution that delivers extreme high performance

through the EF array and comprehensive data management and storage efficiency benefits through FAS and Data ONTAP. Two architectures are available today, described in the following technical reports:

- TR-4126: Sybase Adaptive Server Enterprise on NetApp EF540 and FAS Storage
<https://fieldportal.netapp.com/Core/DownloadDoc.aspx?documentID=88039&contentID=115504>
- TR-4145: NetApp Extreme Analytics Solution for Oracle
<https://fieldportal.netapp.com/Core/DownloadDoc.aspx?documentID=91095&contentID=121344>

Additional reference architectures will come out over the next few months.

Are there any other reference architectures utilizing EF-Series?

Answer: Yes. We have a white paper that documents the NetApp and Atlantis joint storage architecture designed to meet the extreme VDI IOPS requirements utilizing EF-Series and Atlantis ILIO. The paper provides a detailed look at the performance and validation testing, the test results, as well as configuration and best practices for cost-effectively meeting the extreme IOPS use case for VDI. The white paper, FAQ, and presentation regarding this joint architecture are available on the Field Portal at <https://fieldportal.netapp.com/Modules/FieldPortal/Binders/Content.aspx?contentID=166371>.

We also have a NetApp Validated Architecture paper with EF in a NetApp FlexPod® Select solution in which the system exceeds 1 million IOPS. This paper is available on the Field Portal.

Does the EF array support SANtricity advanced replication features?

Answer: Yes. The EF array supports SANtricity Snapshot copies, volume copy, and asynchronous mirroring. A great use case is one in which the ability to replicate data asynchronously from an EF array to an E-Series array with HDDs is needed. This capability provides the customer with a very affordable data protection solution.

7 Resources

Which resources are available?

Answer: The Field Portal has a wide range of EF-Series material, including presentations, technical reports, analyst papers, and sizing guides. And there's more each week. Visit the Flash, EF-Series page to see what's available: <https://fieldportal.netapp.com/ef-series.aspx>.

Is documentation available for the EF array and SANtricity software?

Answer: Yes. Documentation for the EF array and SANtricity Storage Manager software is available on the NetApp Support site at <http://support.netapp.com>.

8 Competition

What is our key advantage over our competitors in the all-flash-array market?

Answer: Building an all-flash array for the enterprise is more than just designing something that can go fast. Real life is different from the engineering or benchmarking lab. As a result, the strengths of the EF-Series are based on its long-standing heritage of performance leadership and enterprise-class RAS.

We've had 20 years of experience designing high-performance storage for production environments, including some of the most demanding tier 1 database environments. The result is a resilient and highly efficient design that's ideally suited for the extreme demands of latency-sensitive workloads. The EF-Series architecture is field-proven to deliver consistent, sustainable, low-latency performance.

Because the applications that an all-flash array is designed to support are often at the heart of a corporation's business processes, the EF-Series has been built with enterprise robustness in mind. Leveraging our extensive industry knowledge focused on designing enterprise-class storage, we architected the EF-Series flash array to provide the highest levels of reliability, availability, and serviceability with features such as a fully redundant I/O path with automated failover and hot-swappable components, advanced monitoring and diagnostic features, and nondisruptive serviceability. And complementing our expertise is the comfort you get knowing that the EF-Series is backed by NetApp's worldwide support organization.

Finally, SANtricity's management interface is both intuitive and powerful, enabling storage administrators to achieve maximum performance and utilization of their EF-Series. And SANtricity offers plug-ins and providers for tight integration with VMware®, Oracle, and Microsoft® applications

Add it all up and you have a collection of features that very, very few competitors can match.

Competitive one-pagers on Violin Memory, IBM TMS, and Nimbus Data can be found on the Field Portal at <https://fieldportal.netapp.com/ef-series.aspx>.

How do I compete against other all-flash arrays that claim 1 million IOPS?

Answer: Several competitors tout 1 million IOPS, and they say we can't match it. That number does look great on collateral, but there's more to performance than brochure numbers. We stayed away from any sort of benchmark special and focused on providing our customers with a real-world number that was achieved in a fully redundant, highly available configuration and sustainable from media (SSD) at less than 1ms latency. As a result, we have a performance number that we can promote with confidence because it can be repeated and achieved by our customers.

I like to think of other flash arrays as dragsters—race cars designed to go very fast on a short, straight, smooth road. This is how they get such great performance numbers for their brochure. But production environments aren't like that. There are curves and bumps, and a dragster wouldn't do very well on a road with curves and bumps. The EF-Series is more like a Formula-1 car, designed to take curves at speed. We've had 20 years of experience designing high-performance storage for production environments, including some of the most demanding tier 1 database environments. The result is a resilient and highly efficient design that's ideally suited for the extreme demands of latency-sensitive workloads. The EF-Series architecture is field-proven to deliver consistent, sustainable, low-latency performance.

Ultimately, we need to make sure that we understand our customers' environment and determine both what they really need and what is realistic to achieve from their configuration.

What about competitors who say that the EF array isn't a ground-up flash design?

Answer: Although there are merits to a ground-up flash design, most competitors today tout this because they want to get credit for all the work they had to do as a result of using flash chips instead of solid-state drives. The reality is that our systems do all these things, too. We're simply leveraging the best-in-class flash technology from our SSD vendor and adding our own technology on top. We do wear-leveling and garbage collection and all the other buzz-word tasks. We just do them at the drive level. The SSDs are incredibly efficient at these tasks. What's more, SANtricity monitors SSD wear life and issues an alert if certain thresholds are reached. The customer can then replace that drive under warranty before ever

experiencing any sort of problem. And since the drives are hot-swappable, there's no interruption in service.

You may also hear competitors say that our legacy architecture limits our performance. The fact is, our legacy is performance leadership, scalability, and enterprise reliability. The EF-Series architecture is the culmination of 20 years of innovation. This technology has been tested and proven in the world's most demanding and complex computing environments. Add to that a legacy of enterprise-class reliability, availability, scalability, manageability, and data protection and it's easy to see that we have a legacy and a knowledge base that others simply can't match.

At the end of the day, we want to make sure that we don't talk about brochure performance and flash buzz words. We want to talk about the customer's requirements and our system's complete capabilities.

9 Training

What training is available on the EF-Series array?

Answer: Multiple web-based training classes are available for the EF array.

- ABCs of EF-Series: <http://learningcenter.netapp.com/LC?ObjectType=WBT&ObjectID=00260062>
- Technical Overview of EF-Series:
<http://learningcenter.netapp.com/LC?ObjectType=WBT&ObjectID=00258361>
- Flash Learning Map:
https://learningcenter.netapp.com/content/secure1/production/learning_maps/flash/lm_flash_t1.html
- Configuring E-Series: What FAS Experts Need to Know:
<https://www.brainshark.com/netapp/vu?pi=zHPzYwy7tzB2XBz0&intk=510403940>
- The Best of Both Technologies: Comparing Data ONTAP and E-Series:
<http://learningcenter.netapp.com/LC?ObjectType=WBT&ObjectID=00230069>
- There is also a learning map for all NetApp flash technologies:
http://learningcenter.netapp.com/content/public/production/learning_maps/flash/lm_flash_t1.htm
!

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States and/or other countries. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.