



Technical Report

NetApp EF300 with Splunk Enterprise

Mitch Blackburn and Kyler Mick, NetApp
July 2021 | TR-4903

Abstract

This technical report describes the integrated architecture of the NetApp® EF300 all-flash array and Splunk design. Optimized for node storage balance, reliability, performance, and density, this design employs the Splunk clustered index node model, with higher scalability and lower TCO. Decoupling storage from compute provides the ability to scale each separately, saving the cost of overprovisioning one or the other. In addition, this document summarizes the performance test results obtained from a Splunk machine log event simulation tool.

TABLE OF CONTENTS

Executive summary 4

About Splunk Enterprise..... 4
 Primary use cases4

About NetApp EF300 5

About SANtricity OS 5
 SANtricity security features.....6
 SANtricity reliability features6
 SANtricity data protection features7
 Dynamic Disk Pools.....7

Decoupling storage from compute 8
 Sizing.....9
 Other considerations 11

Splunk Enterprise edition and NetApp EF300 testing 12
 Overview of Splunk cluster testing used for E-Series 13
 Eventgen data 13
 Cluster replication and searchable copies factor 14
 E-Series with DDP baseline test setup 14
 Baseline test results for E-Series..... 15

Conclusion 17

Appendix A: Splunk App for NetApp E-Series and EF-Series 18

Where to find additional information 20

Version history..... 20

LIST OF TABLES

Table 1) SANtricity features for long-term reliability.6

Table 2) Sizing example for a nonclustered environment.....9

Table 3) Increased capacity needs for clustering of e-commerce logs 10

Table 4) Splunk cluster server hardware 12

Table 5) SSD hot/warm tier indexing results. 16

Table 6) SSD hot/warm tier searching results. 16

LIST OF FIGURES

Figure 1) Sample of storage and compute separation.....8

Figure 2) Splunk logical configuration.	13
Figure 3) Splunk cluster with E-Series.....	15
Figure 4) Job inspector output for first matching result.	16
Figure 5) Job inspector output for first 1,000 matching results.	16
Figure 6) Failed drive impact.	17
Figure 7) Configuration tab of the SANtricity Performance App for Splunk Enterprise	18
Figure 8) Major Event Log (MEL) information for an array needing attention	18
Figure 9) Performance tab of the SANtricity Performance App for Splunk Enterprise.....	19
Figure 10) Events tab of the SANtricity Performance App for Splunk Enterprise	19

Executive summary

NetApp EF-Series enables Splunk environments to maintain the highest levels of performance and uptime for workloads by providing advanced fault recovery features and easy in-service growth capabilities to meet ever-changing business requirements. By decoupling storage from compute, you gain the ability to scale capacity and compute separately, saving the cost of overprovisioning one or the other.

The EF-Series is designed to handle the most extreme application workloads with very low latency. Typical use cases include application acceleration; improving the response time of latency-sensitive applications; and improving the power, environmental, and capacity efficiency of overprovisioned environments. EF300 storage systems leverage the latest NVMe technologies to provide NVMe over fabric (NVMe-oF) host interfaces as well as NVMe SSD drives.

NetApp EF-Series also protects your previous investments by providing SCSI-based host interfaces until you are ready to move to NVMe-oF. Attaching SAS expansion shelves to the EF300 array can provide superior business value by adding large capacity NL-SAS HDD drives to act as a cold tier for Splunk.

Splunk is the leading operational intelligence software that enables you to monitor, report, and analyze live streaming and historical machine-generated data, whether it is located on premise or in the cloud. An organization's IT data is a definitive source of intelligence because it is a categorical record of activity and behavior, including user transactions, customer behavior, machine behavior, security threats, and fraudulent activity.

This technical report describes the integrated architecture of the NetApp EF300 all-flash array with expansion and Splunk Enterprise.

About Splunk Enterprise

All your IT applications, systems, and technology infrastructure generate data every millisecond of every day. This machine data is one of the fastest growing and most complex areas of big data. Splunk Enterprise employs machine learning powered analytics to turn this machine data into actionable insights. It allows you to leverage artificial intelligence (AI) and machine learning (ML) for predictive and proactive business decisions.

For more information about Splunk Enterprise, visit <https://www.splunk.com/>.

Primary use cases

Splunk can be deployed for use in a wide variety of use cases, and it provides creative ways for users to gain intelligence from data.

Application delivery

Gain end-to-end visibility across distributed infrastructures, troubleshoot application environments, monitor performance for degradation, and monitor transactions across distributed systems and infrastructure.

Security, compliance, and fraud

Enable rapid incident response, real-time correlation, and in-depth monitoring across data sources. Conduct statistical analysis for advanced pattern detection and threat defense.

Infrastructure and operations management

Proactively monitor across IT silos to enable uptime, rapidly pinpoint and resolve problems, identify infrastructure service relationships, establish baselines, and create analytics to report on SLAs or track service provider SLAs.

Business analytics

Provide visibility and intelligence related to customers, services, and transactions. Recognize trends and patterns in real time and provide valuable understanding of new product features' impact on back-end services. Gain valuable understanding of the user experience for greater user satisfaction and prevent drop-offs, improve conversions, and boost online revenues.

About NetApp EF300

In one powerful all-flash array package, the EF300 array delivers optimal performance for both random workloads and large sequential workloads. The array can deliver consistent response times for up to 670,000 4KB random read IOPS at 250µsec with as few as 24 NVMe SSDs. The same configuration can deliver up to 20GBps large sequential read throughput and about 7GBps cache-mirrored large sequential write throughput. When your workload meets the criteria of the built-in full stripe write acceleration feature, you can accelerate write performance up to 12GBps.

The EF300 can also provide SAS expansion shelves for additional SAS SSDs or large capacity NL-SAS HDDs. This allows one system to provide a hot tier and cold tier for Splunk.

SAS expansion supports up to 240 NL-SAS HDDs or 96 SAS SSDs in various shelf combinations.

The EF300 array is used for storage solutions that require the depth of enterprise-grade SAN storage and that consistently deliver response times in the sub-250µsec range. The array supports the SCSI over FC protocol and the NVMe/FC protocol on the 32Gb FC host interface card (HIC). The iSCSI protocol is supported on the 25Gb iSCSI HIC. NVMe/IB, NVMe/RoCE, SRP/IB, and iSER/IB are supported on the 100Gb HIC.

EF300 arrays use the web-based SANtricity System Manager UI to manage individual arrays, and SANtricity Unified Manager enables you to organize and manage multiple new-generation E-Series and EF-Series arrays from a central management application. The built-in web services API integration or the management client-based web services package makes the EF-Series product line easier than ever to integrate with your standard API-driven environment.

The EF300 array provides the following benefits:

- Support for wide-ranging workloads and performance requirements
- Fully redundant I/O paths, advanced protection features, and proactive support monitoring and services for high levels of availability, integrity, and security
- A level of performance, density, and economics that leads the industry

For in-depth information on the EF300, see [TR-4877: Introduction to NetApp EF300 array](#).

About SANtricity OS

The NetApp EF300 controller and SANtricity OS use the on-box, browser-based management interface, SANtricity System Manager.

EF300 storage systems are shipped preloaded with SANtricity OS, which includes SANtricity System Manager. To discover multiple EF300 storage systems running SANtricity OS from a central view,

download SANtricity Unified Manager (which includes the Web Services Proxy) from the NetApp Support site. Then load it on a management server that has IP access to the storage systems.

For further information about the SANtricity Unified Manager and the SANtricity System Manager, see the [E-Series and SANtricity documentation resources page](#).

SANtricity security features

The new-generation EF-Series arrays running the latest SANtricity OS are Common Criteria certified (NDcPP v2 certification). SANtricity security features include LDAP, RBAC, and SSL certificates. For complete details and workflow examples, see the following resources:

- [TR-4712: NetApp SANtricity Management Security Feature Details and Configuration Guide](#),
- [TR-4855: Security Hardening Guide for NetApp SANtricity](#)
- [TR-4813: Managing Certificates for NetApp E-Series Storage Systems](#)

SANtricity drive security technology provides comprehensive security for data at rest without sacrificing system performance or ease of use and supports both internal and external key management.

For more information about disk encryption, see [TR-4474: SANtricity drive security](#).

SANtricity reliability features

Table 1 provides a list of SANtricity reliability features and a brief explanation of each with references to additional information.

Table 1) SANtricity features for long-term reliability.

Reliability features with SANtricity

Proactive drive monitor and data evacuator. Nonresponsive drives are automatically power-cycled to see if the fault condition can be cleared. If the condition cannot be cleared, the drive is flagged as failed. For predictive failure events, the evacuator feature starts to remove data from the affected drive to move the data before the drive fails. If the drive fails, rebuild resumes where the evacuator was disrupted, reducing the rebuild time.

Automatic drive fault detection, failover, and rebuild. You can perform these tasks by using global hot spare drives for standard RAID and spare pool capacity for DDP.

SSD wear-life tracking and reporting. This metric is found in the Hardware tab's Drive Settings dialog box. It indicates the wear life of SSDs and replaces two SSD wear-life metrics (average erase count and spare blocks remaining) that were in previous versions of SANtricity. The metric is Percent Endurance Used; to access it, select a drive from the hardware view and then select Settings.

Online drive firmware upgrade. This feature upgrades one drive at a time and tracks writes to the affected drives during the upgrade window; it should be used only during low write I/O periods. Parallel drive firmware upgrades are supported offline to upgrade multiple drives more quickly during a maintenance window.

Automatic load balancing. This feature provides automated I/O workload balancing and confirms that incoming I/O traffic from hosts is dynamically managed and balanced across both controllers. The workload of each controller is continually monitored and analyzed in the background. When I/O on one controller significantly exceeds the I/O on the other controller for a prolonged, predictable period, SANtricity can change volume ownership from the busy controller to the less busy controller. The feature does not react to short-term changes in I/O patterns. However, when a change of ownership is needed, SANtricity interacts with the affected host multipath driver to initiate an implicit path failover. Most current server operating systems and associated multipath drivers support implicit failover. For more information, search for "What is automatic load balancing?" in the System Manager online help.

Embedded SNMP agent. For the EF300 controller, SNMP is supported natively. The embedded SNMP agent complies with the SNMP V2C standard and RFC 1213 (MIB-II). For more information, search for "manage SNMP alerts" in the System Manager online help.

Reliability features with SANtricity

Automatic alerts. This feature sends email alerts to notify data center support staff about events on the storage array.

Event Monitor and system log. The SANtricity Storage Manager Event Monitor automatically records events that occur on the storage array. Syslog enables a second level of activity tracking that allows you to connect events with associated changes recorded in the system log.

AutoSupport. E-Series products have supported AutoSupport for several releases.

Ability to enable or disable AutoSupport maintenance window. AutoSupport includes an option for enabling or suppressing automatic ticket creation on error events. Under normal operation mode, the storage array uses AutoSupport to open a support case if there is an issue. To enable or disable the AutoSupport maintenance window, select Support > Access Management > AutoSupport.

SANtricity data protection features

EF-Series has a reputation for reliability and availability. Many of the data protection features in EF-Series systems can be beneficial in a Splunk environment.

Background media scan and data assurance (T10 PI)

Media scan is a background process that is performed by the controllers to provide error detection on the drive media. The main purpose of the feature is to detect and repair media errors on disk drives that are infrequently read by user applications and where data loss might occur if other drives in the volume group fail. A secondary purpose is to detect redundancy errors such as data/parity mismatches. A background media scan can find media errors before they disrupt normal drive reads and writes.

The data assurance feature provides controller-to-drive data integrity protection through the SCSI direct-access block device protection information model. This model protects user data by appending protection information to each block of user data. The protection model is sometimes referred to as data integrity field protection or T10 PI. This model makes sure that an I/O has completed without any bad blocks written to or read from disk. It protects against displacement errors, data corruption resulting from hardware or software errors, bit flips, and silent drive errors, such as when the drive delivers the wrong data on a read request or writes to the wrong location.

You need both data assurance and media scan. They work complementarily to protect your data.

Unreadable sector management

This feature provides a controller-based mechanism for handling unreadable sectors detected both during normal I/O operation of the controller and during long-lived operations such as reconstructions. The feature is transparent to the user and requires no special configuration.

Dynamic Disk Pools

With Dynamic Disk Pools (DDP) technology, NetApp SANtricity OS and management software allows you to create pools in addition to traditional volume groups (generally referred to as RAID groups). A pool can range in size from a minimum of 11 drives to as large as all the drives in a storage system, which is up to 240 drives in the NetApp EF300 system. Pools can consist of either HDDs or SSDs. In addition, pools and volume groups can coexist in the same system.

For more information about DDP, see [TR-4652: SANtricity OS Dynamic Disk Pools](#).

Decoupling storage from compute

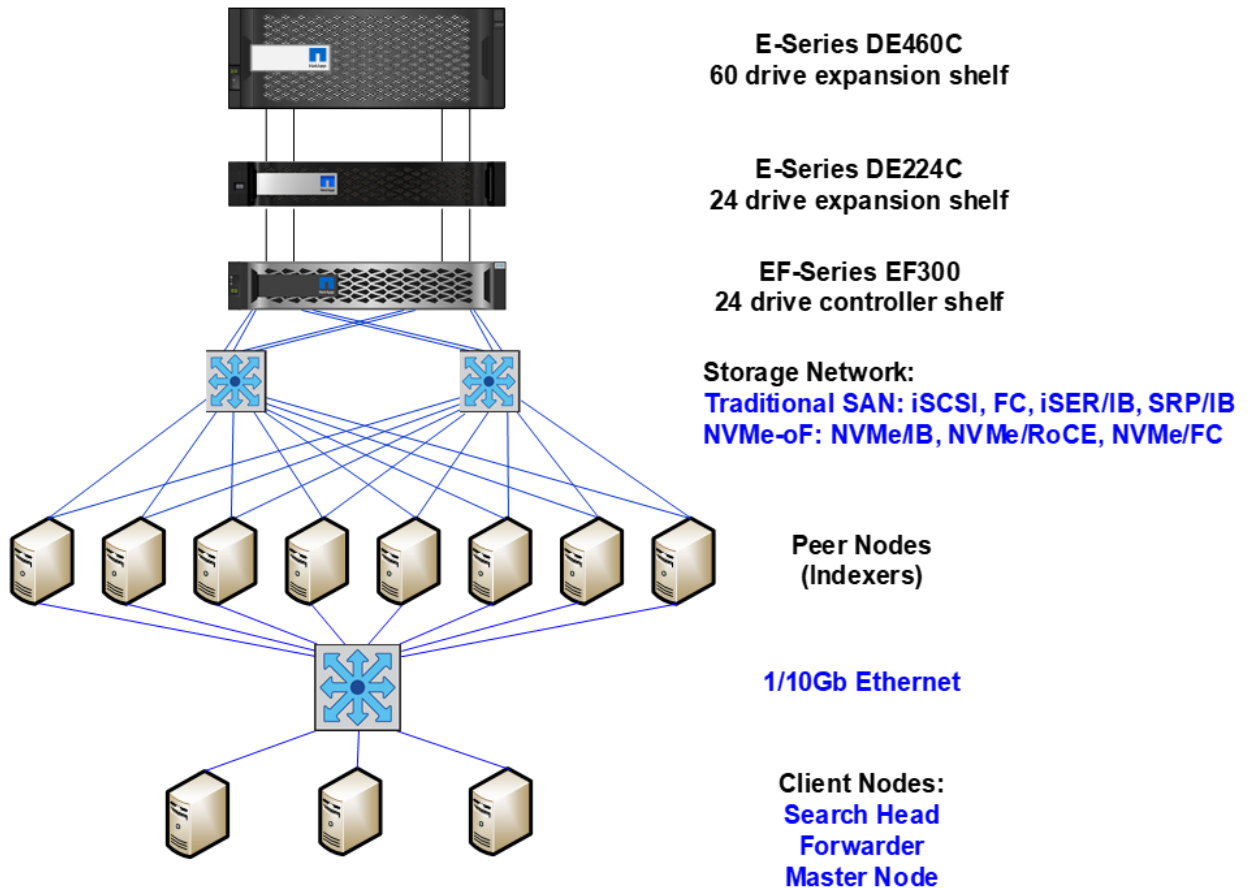
With the advent of larger and larger SSDs, up to 15.3TB SSDs now available, and the expanded use of specialized compute for data analytics such as GPUs, the ability to decouple storage and compute separately with Splunk is becoming an economic necessity.

Figure 1 shows a sample Splunk architecture where an EF300 array is connected over iSCSI to eight cluster nodes, and then client nodes are connected over Ethernet to the cluster nodes. With NetApp EF-Series, it's possible to begin with 24 NVMe SSDs and add additional SAS SSDs up to 120 SSDs without needing to add more cluster HDDs in an all-SSD system. If HDDs are being deployed for the cold tier, it is possible to add up to 240 HDDs by adding additional SAS drive shelves to the system shown. For more information about the EF300 configuration, see [TR-4877: Introduction to NetApp EF300 array](#).

The advantages of decoupling storage from compute include:

- Ability to scale capacity and compute separately, saving the cost of overprovisioning one or the other.
- Flexibility to use excess top-of-rack switches bandwidth for the storage network, use a wholly different storage network such as Fibre Channel, or connect the array as direct-attached storage (DAS).

Figure 1) Sample of storage and compute separation.



This type of decoupling, for example, can allow a company with 100 nodes to reduce its number of nodes substantially, if 100 nodes of compute aren't required. This change provides a significant reduction in the rack space required and the associated cooling and power requirements. In contrast, if the need is more compute, than less expensive servers can be purchased that don't require space for additional storage and have a smaller footprint.

After the decision is made to use a separate storage array, sizing and configuring it are straightforward.

Sizing

Estimate your storage requirements

Before sizing, it is necessary to estimate your daily ingest rate. When Splunk Enterprise indexes your data, it creates two main types of files: the rawdata file that contains the original data in compressed form and the index files that point to this data. (It also creates a few metadata files, which don't consume much space.) With some experimentation, you can estimate how much index disk space you will need for a given amount of incoming data. A general rule used by Splunk for syslog-type data is that after it has been compressed and indexed in Splunk, it occupies approximately 50% of its original size:

- 15% for the rawdata file
- 35% for associated index files

As an example, assume the current environment is split between two logging formats: syslog and Windows. The logs generated by the Windows servers and Active Directory are logged as Windows Events. For antivirus software, we have Symantec Endpoint Protection; the Symantec Endpoint Protection Manager runs on a Windows server. The logs ingested into Splunk include the firewall logs from the Cisco ASA devices, Linux logs, webserver logs, Cisco WSA proxy log, Window logs, Symantec logs, and e-commerce transaction logs (these are our data sources). Retention periods are as shown, assume 30 days of the retention period for the hot/warm tier.

At a high level, Splunk calculates total disk storage as follows:

$$(\text{Daily average indexing rate}) \times (\text{retention policy}) \times \frac{1}{2}$$

Table 2 shows what the results of our sizing to this point might look like.

Table 2) Sizing example for a nonclustered environment.

Data source	GB per day	Raw comp rate	Base size of raw	Index comp rate	Base Size of index files	Retention days	Estimated size on disk	Hot/warm	Cold
Cisco ASA firewall logs	100	0.15	1,350	0.35	3,150	90	4,500	1,500	3,000
Badge reader log	60	0.15	1,620	0.35	3,780	180	5,400	900	4,500
Cisco WSA Proxy	120	0.15	540	0.35	1,260	30	1,800	1,800	0
Linux logs	80	0.15	1,080	0.35	2,520	90	3,600	1,200	2,400
Windows logs	140	0.15	1,890	0.35	4,410	90	6,300	2,100	4,200
Symantec logs	20	0.15	270	0.35	630	90	900	300	600
Web logs	440	0.15	5,940	0.35	13,860	90	19,800	6,600	13,200
e-commerce logs	300	0.15	16,425	0.35	38,325	365	54,750	4,500	50,250
Total	1260						97,050	18,900	78,150

The best way to get an idea of your space needs is to experiment by indexing a representative sample of your data, and then checking the sizes of the resulting directories in `$(SPLUNK_HOME)/var/lib/splunk/defaultdb`. For information about estimating storage requirements, see the [Capacity Planning Manual](#) of the online Splunk documentation; specifically, the [Estimate Your Storage Requirements](#). Splunk also provides an online application to aid in sizing, [Splunk Storage Sizing](#).

At this point, we have a good idea of our daily indexing volume. In this example, 1260GB per day. We can use Splunk's [Summary of Performance Recommendations](#) in the Capacity Planning Module to estimate the number of reference machines required for indexing and searching. Because the daily indexing volume is just over 1TB per day, let's assume we need one Search Head and eight Indexers.

The last step is to decide if you are clustering any of the indexes and how they will affect the capacity requirements. Let's assume that the e-commerce data is critical, and that the Sales Support team needs to be able to always access this data.

You need to determine the following factors:

- **Replication factor (RF).** Specifies how many total copies of rawdata the cluster should maintain. This sets the total failure tolerance level.
- **Search factor (SF).** Specifies how many copies are searchable (searchable buckets have both rawdata and index files). This determines how quickly you can recover the search capability.

This is where the power of decoupling storage from compute with an E-Series array comes in. Because high availability is built into the array, it is only necessary to have an RF=2, and because we want to recover the search capability immediately, we set the SF=2. So, it is not necessary to add compute to add storage or install additional unneeded storage in all nodes. For the e-commerce logs, the amount of capacity required will double then, as shown in Table 3.

Table 3) Increased capacity needs for clustering of e-commerce logs.

Data source	GB per day	Raw comp rate	Base size of raw	Index comp rate	Base size of index files	Retention days	Estimated size on disk	Hot/warm	Cold
Cisco ASA firewall logs	100	0.15	1,350	0.35	3,150	90	4,500	1,500	3,000
Badge reader log	60	0.15	1,620	0.35	3,780	180	5,400	900	4,500
Cisco WSA Proxy	120	0.15	540	0.35	1,260	30	1,800	1,800	0
Linux logs	80	0.15	1,080	0.35	2,520	90	3,600	1,200	2,400
Windows logs	140	0.15	1,890	0.35	4,410	90	6,300	2,100	4,200
Symantec logs	20	0.15	270	0.35	630	90	900	300	600
Web logs	440	0.15	5,940	0.35	13,860	90	19,800	6,600	13,200
e-commerce logs	300	0.15	32,850	0.35	76,650	365	109,500	9,000	100,500
Total	1260						151,800	23,400	128,400

You now have an estimate of the number of servers and the amount of storage required for our implementation.

Configure your E-Series storage based on estimate

You can now configure the EF-Series storage. Use the DDPs from which you will create the volumes (LUNs). You need one volume/indexer for the hot/warm tier (eight total) and one volume/indexer for the cold tier (eight total).

For the hot/warm tier, the usable storage capacity required is 23400GB, which can be satisfied by using high-performance NVMe SSD drives.

For the cold tier, the usable capacity required is 128,400GB. In this instance, NetApp recommends using high-capacity NL-SAS drives, such as 18TB drives.

Configure your array by using the EF300 with a SAS expansion shelf, which is a 2U shelf containing 12 drives. Use an all-SSD drive hot/warm tier, which requires approximately 20x 1.9TB drives. This provides a usable capacity of 27.3TB, which leaves some room for growth. For the cold tier, use 11x 18TB NL-SAS drives. This configuration provides a usable capacity of 141.06TB, which again allows some room for growth.

You can now create two DDPs, one for the hot/warm tier using the SSD drives and the other for the cold tier using the NL-SAS drives. Each of these DDPs is then cut into eight volumes (LUNs) that are presented to the eight indexers. Therefore, each indexer has two volumes presented to it, one for the hot/warm tier and one for the cold tier.

Other considerations

EF-Series

To prepare your server for storage access, see [SANtricity Software Express Configuration for Linux](#).

This document guides you through the following procedures:

- Installing SANtricity host-side applications
- Configuring multipath
- Installing NetApp Host Utilities
- Using the `iscsiadm open-iscsi` utility with E-Series products (if using iSCSI)

The NetApp Interoperability Matrix Tool (IMT) has approximately 85,000 entries that not only connect to any SAN but also support it. To verify whether your configuration is supported and to check for any changes that might be required for correct functioning of your E-Series, see the [Interoperability Matrix Tool](#).

Linux configuration

All servers in the Splunk cluster were tested with SLES 15 SP2 with default kernel settings.

For persistent deployments, administrators should consider the following flags when adding a mount into `/etc/fstab`:

- `nobarrier`: Allows data to sit in cache instead of being flushed. There is a large performance gain on particular workloads by allowing `nobarrier`. This option should only be used for E-Series storage, because internal disks might not have battery backup.
- `noatime`: Forces file reads to not record their access times to disk, which can increase I/O dramatically on heavy read loads. Setting the `noatime` flag is only recommended for file systems or dependent applications where a record of the last access time of a file for reading is unnecessary.
- `_netdev`: Required for configurations using iSCSI and iSER network protocols. The `_netdev` option forces the mount to wait until the network is up before trying to mount. Without this option, the OS

attempts to mount the disk prior to the network being completely available, and it could lead to various timeouts or the OS entering recovery mode.

- `discard`: If the storage volume is thinly provisioned, providing the `discard` flag allows the file system to reclaim space. This flag can cause performance degradation. Administrators who want to control when discards take place (for example, nightly) should consider using `fstrim` or an equivalent command for the OS.

Note: The use of thin-provisioned volumes is not recommended with Splunk installations.

To increase performance, jumbo frames should be set on the network. Setting jumbo frames for the storage is explained in the E-Series documentation. On the server, they are configured by adding an entry of `MTU=9000` to the interface file in the `/etc/sysconfig/network-scripts` directory and restarting the interface. To validate that jumbo frames have been set, use the `ip link show` command:

```
[root@ictk0103r720-4 ~]# ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT link/loopback
00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP mode DEFAULT qlen 1000
link/ether b0:83:fe:d5:ae:62 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc mq state UP mode DEFAULT qlen 1000
link/ether b0:83:fe:d5:ae:64 brd ff:ff:ff:ff:ff:ff
```

Splunk

To make sure that your Linux environment is set up correctly, see the deployment information for Splunk Enterprise Edition 8: [Splunk Enterprise Installation Instructions](#).

Splunk Enterprise edition and NetApp EF300 testing

NetApp tested a simulated Splunk cluster environment with the E-Series EF300 and commodity servers configured for the index peer node disk for hot/warm and cold data buckets. The server hardware was selected by following the recommendations from the Splunk reference architecture system requirements. Table 4 lists the Splunk cluster server hardware that was used in the test environment.

Table 4) Splunk cluster server hardware.

Splunk component	Qty.	CPU	Sockets	Cores per socket	Logical processors	Speed	RAM
Indexer Peer Node	12	E5-2640 v3	2	8	32	2.60GHz	256GB
Search Head	2	E5-2640 v3	2	8	32	2.60GHz	256GB
Cluster Master	1	E5-2640 v4	2	8	32	2.60GHz	256GB
License Master, Monitoring Console	1	E5-2640 v5	2	8	32	2.60GHz	256GB
Forwarder	4	E5-2640 v6	2	8	32	2.60GHz	256GB

The ingest machine log data was created by using the Splunk workload tool `eventgen`. The cluster had 12 index peer nodes to handle ingesting 170GB of simulated machine syslog data per indexer, for a total of 2TB per day for the entire cluster. The ingest was then scaled to 340GB of data per indexer, for a new total of 4TB per day.

Overview of Splunk cluster testing used for E-Series

The Splunk cluster configuration includes the following components:

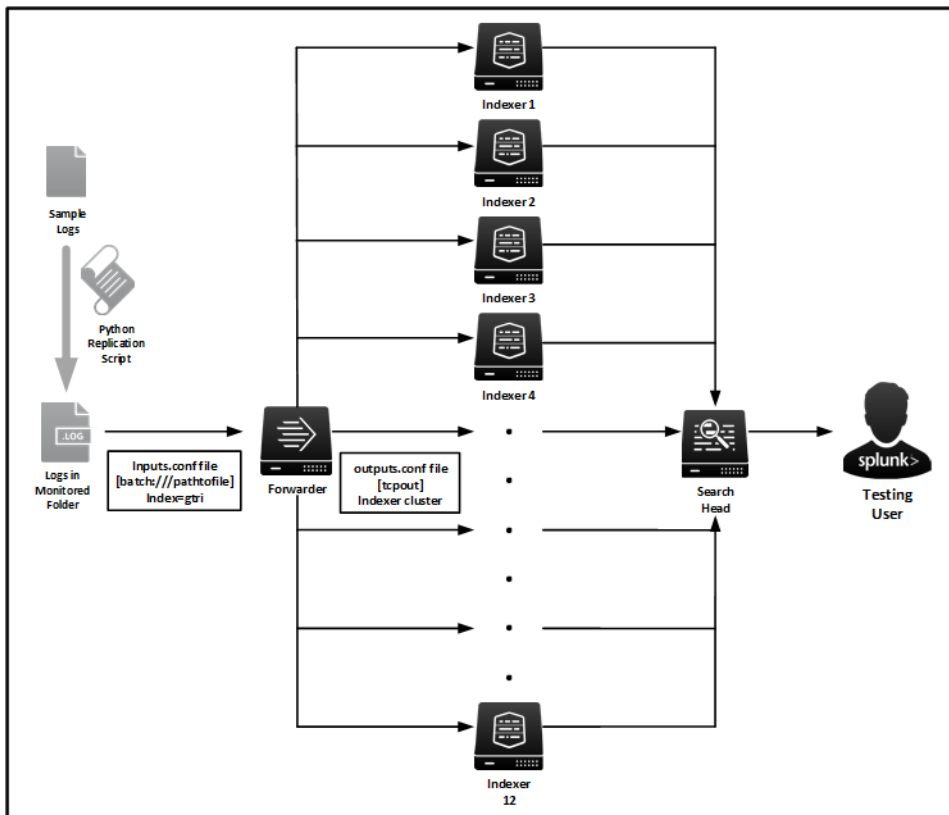
- **Forwarder.** Ingest 170GB of machine log data files into the cluster of index node peers.
- **Index peer nodes.** Index the ingested machine syslog data and replicate data copies in the cluster.
- **Search head.** Execute custom searches for dense, very dense, sparse, and very sparse data from the cluster of index peer nodes.
- **Master.** Monitor and push configuration management changes for the cluster.

Eventgen data

The machine log dataset was created with Splunk's event generator, the eventgen. The Splunk event generator is a downloadable Splunk app that is available from the Splunk website. Splunk eventgen allows users to load samples of log files or exported `.CSV` files as an event template. The templates can then be used to create artificial log events with simulated timestamps. A user can modify the field values and configure the random variance while preserving the structure of the events. The data templates can be looped to provide a continuous stream of real-time data. For more eventgen information, visit [Splunk eventgen app](#).

For our testing, the eventgen was loaded into the cluster and was configured to produce a 170GB simulated syslog type file for the Splunk forwarder instance. The file was then split into smaller syslog files on each of 12 individual Splunk heavy forwarder instances, each ingesting data in a one-to-one data path to one of the 12 index peer nodes. The total ingested data was 2TB per day loaded into the cluster for each simulated daily index. The ingested data was then scaled to 4TB per day during scale testing. The logical configuration is shown in Figure 2.

Figure 2) Splunk logical configuration.



Cluster replication and searchable copies factor

The search factor determines the number of searchable copies of data that the indexer cluster maintains for the number of searchable copies of each bucket. The default value for the search factor is two, meaning that the cluster maintains two searchable copies of all the data. The search factor must be less than or equal to the replication factor. The replication factor is the number of copies of data that you want the cluster to maintain. Peer nodes store incoming data in buckets, and the cluster maintains multiple copies of each bucket. The cluster stores each bucket copy on a separate peer node. The number of copies of each bucket that the cluster maintains is the replication factor. The default replication value for a cluster is three.

The E-Series test was configured with a replication factor of two and searchable copies with a factor of two. The E-Series provides additional redundancy with additional copies of indexed data located in the DDP volumes for each indexer. This additional redundancy allows the replication factor of two to seamlessly provide fewer copies of index data in the Splunk cluster for performance and data storage benefits.

E-Series with DDP baseline test setup


The EF300 configuration for the baseline test was configured with DDP LUNs by using the following components:

- 24x 1.8TB SSDs with a pool preservation capacity of 2 drives and 14% SSD optimization capacity, offering ~31TB of usable capacity for the Splunk cluster hot/warm data buckets
- 12x 10TB NL-SAS drives with a pool preservation capacity of 2 drives, offering ~85TB of usable capacity for the Splunk cluster cold data buckets
- 10Gb Ethernet (10GbE) private network for index peer nodes
- 32Gb NVMe/FC SAN for E-Series and index peer nodes

The DDP LUNs were configured into 12 volumes—one each for the 12 index peer node hosts. The mounted volumes were configured as XFS file systems on the SLES 15 OS of each indexer.

Figure 3 illustrates the E-Series baseline configuration.

Figure 3) Splunk cluster with E-Series.



Baseline test results for E-Series

To ensure consistency, the same data was loaded through the same scripts and Splunk configurations. The testing pattern was to ingest 2TB of data into the storage configuration by using the EF300 configurations. After the script used to transfer data and manipulate the data ran, the data began to be copied to the `splunk_forward` directory where the forwarder began sending the files to be indexed across the indexers.

At this time, measurements were taken of the indexing rate for each of the indexers based on the following Splunk query:

- Average and Peak Indexing Rates: `index="_internal" source=*metrics.log* series=gtri per_index_thruput | timechart span=5m max(kbps) avg(kbps)`

Next, static searches were provided on data that was indexed and record the times. For the purpose of this testing, we were provided with a dataset that consisted of strings that were interspersed at given intervals and searched for these strings in order to determine the performance of a search. These strings accounted for ~0.01% of the dataset. Search performance was measured by the time required to return the first matching result and the first 1,000 matching results.

For all searches, the Splunk Job Inspector was used as a barometer of how quickly the search had run. The Splunk Job Inspector can be accessed through the Web UI below the search bar by selecting Job > Inspect Job.

Hot/warm tier SSD

The first testing scenario used 2.5TB of hot/warm storage per indexer using NVMe SSD drives. In addition, a 6.5TB cold storage was added using a NL-SAS 7.2K HDD drives with no read cache on it. The sample logs provided were forwarded to all indexers for improved performance (as would be best practice for this architecture). Table 5 shows the indexing results and Table 6 shows search results. Figure 4 and Figure 5 show the Splunk Job Inspector output for the searches.

Table 5) SSD hot/warm tier indexing results.

Data ingest volume (GB per day)	Average indexing rate observed
2000GB per day	24.890MBps
4000GB per day	49.357MBps

Table 6) SSD hot/warm tier searching results.

Search performed	Total scanned events	Average search time
1 matching result	8,754	0.718 seconds
1,000 matching results	15,627,224	88.5 seconds

Figure 4) Job inspector output for first matching result.

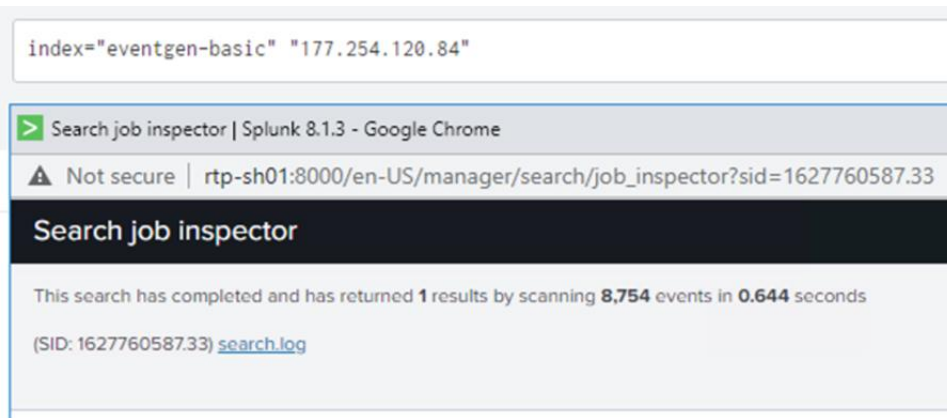
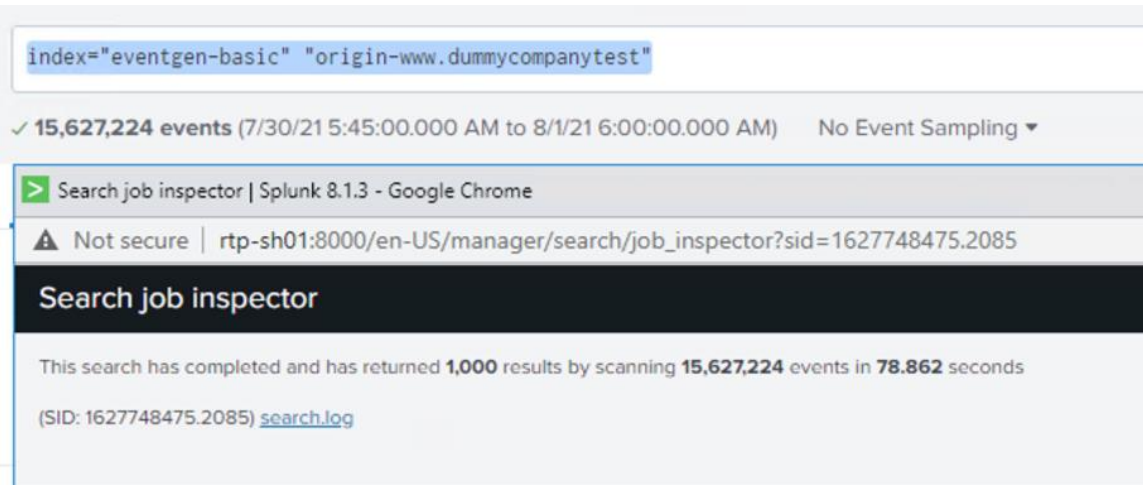


Figure 5) Job inspector output for first 1,000 matching results.




To improve search times for an HDD hot/warm tier, an SSD read cache could be used. This is dependent on the types of search queries being run against the indexers. If the queries are regularly querying the same data, an SSD read cache would be appropriate for improved performance. This would also be true if the cold tier is regularly queried.

System drive failure simulation

In order to validate the recommendation of DDPs, a drive failure of the all SSD hot/warm tier was tested to observe the overall impact on Splunk indexing. For the simulation, a drive failure was induced while the overall index rate of the array under heavy ingest rates was monitored. After drive failure, the array recovered and migrated mirrored data to additional SSDs within the pool without noticeably affecting Splunk performance. Figure 6 shows the Splunk ingestion chart with a failed drive and the normal performance fluctuations.

Figure 6) Failed drive impact.



While there was a slight dip in the indexing speed, within approximately eight minutes, the indexing rate was restored to normal.

In a traditional DAS model where all drives are in the indexer, a drive failure can lead to the indexer being taken out of service until the drive is replaced and the data recovered or rebuilt. The effect on performance while this replace and recovery operation is taking place can range from hours to days for large capacity HDD drives.

Conclusion

The NetApp EF300 all-flash array provides several significant advantages over internal DAS for Splunk deployments. These advantages include exceptional storage management capabilities, dramatically improved reliability, high availability, and limited performance degradation because of failure conditions such as disk or controller failures. By decoupling storage from compute, you gain the ability to scale capacity and compute separately, saving the cost of overprovisioning one or the other.

The advantages also include excellent performance handling ingest of machine log data and excellent search capabilities at very low latency with an EF300 configuration of all-flash NVMe SSDs for hot and warm Splunk buckets or a hybrid configuration employing HDDs in SAS expansion shelves providing excellent performance and reliability for the Splunk cold data bucket tiers.

Organizations that use Splunk often use traditional server-based storage with inefficient, hard-to-scale internal DAS. The NetApp reference design employs the managed DAS model, with higher scalability and performance. The reduction of the Splunk cluster replication factor available in EF-Series storage reduces the amount of data indexed. The reduction also prevents unnecessary purchases of compute nodes for storage-intensive workloads for Splunk environments that need to grow to meet organizational requirements.

The NetApp reference architecture for Splunk is optimized for node storage balance, reliability, performance, storage capacity, and density. From an administrative standpoint, EF-Series offers simplified storage management with a browser-based UI. This solution allows new volumes and DDPs to be created easily and provisioned immediately for use by Splunk cluster servers. In addition, existing volumes and DDPs can all be increased in size dynamically to provide additional capacity and performance as required for the Splunk indexer cluster environment.

Appendix A: Splunk App for NetApp E-Series and EF-Series

The SANtricity Performance App for Splunk Enterprise makes it easy to monitor the health and performance of NetApp E-Series and EF-Series storage systems from within the Splunk environment.

The Configuration tab, shown in Figure 7, provides a basic overview of each actively monitored array. A user can scroll down to view individual volume groups, volumes, and drives by selecting an individual array. If an array needs attention, a user can click “(more info)” to view Major Event Log (MEL) data specific to the issue, as in Figure 8.

Figure 7) Configuration tab of the SANtricity Performance App for Splunk Enterprise

Array	Array Name	arrayid	Needs Attention	Fixup Activity	Config Generation	Volume Groups/Pools Count	Volume Count	Drive Count	Boot Time	Serial Number
1	ICTM0802501C1	66bc2e92-c17e-4b28-a6af-b1776bfc69e	true	false	12690	1	6	180	2017-07-07 13:06:05	021711026794
2	Site_A_05	524ca6ad-7d5e-4219-8cc9-53fc5318b7e4	true	false	90002	2	3	12	2017-05-28 17:31:08	031541000839
3	Site_B_02	b373c106-4fe4-4cab-9d8f-ce73c267845d	true	true	40101	3	6	24	2017-06-28 12:53:11	711417000086

Figure 8) Major Event Log (MEL) information for an array needing attention

Failure Type	Object Type	Object Name	Object Reference	Object Data	Extra Data
Failed Drive	drive	99.6.6	010000005000CCA251ACBC9000000000000000	null	null
netNtpQueryFail	controller		0700000000000000000000000000000001	null	null
netNtpQueryFail	controller		0700000000000000000000000000000002	null	null
netNtpServiceUnavailable	controller		0700000000000000000000000000000002	null	null
netNtpServiceUnavailable	controller		0700000000000000000000000000000001	null	null
diskPoolReconstructionDriveCountBelowThreshold	pool	Pool_1	04000000600A098000A4B279000010C2597098F8	null	null
diskPoolDriveFailure	pool	Pool_1	04000000600A098000A4B279000010C2597098F8	null	null
degradedVolume	pool	Pool_1	04000000600A098000A4B279000010C2597098F8	null	null


The Performance tab, shown in Figure 9, displays real-time graphical information about read/write operations, read/write latency, and read/write throughput. The view defaults to a single array, but a user can choose to view multiple arrays or to view performance by controller, volume group, volume, or individual drive.

Figure 9) Performance tab of the SANtricity Performance App for Splunk Enterprise



The Events tab, shown in Figure 10, provides information about all MEL events reported by an array. Events can be sorted by a variety of fields and filtered by priority as needed.

Figure 10) Events tab of the SANtricity Performance App for Splunk Enterprise



The SANtricity Performance App for Splunk Enterprise requires:

- Splunk 6.1 or later running on Linux
- NetApp E-Series/EF-Series storage arrays running firmware 7.84 or later
- [NetApp SANtricity Web Services Proxy](#) 1.3 or later (2.0 or later for best feature set) running on Windows or Linux

- [NetApp SANtricity Performance App for Splunk Enterprise](#) (available from SplunkBase)
- [Technology Add-On for NetApp SANtricity](#) (available from SplunkBase)

To use the app, first, install NetApp SANtricity Web Services Proxy on any server with network access to the monitored arrays. Then, upload and install NetApp SANtricity Performance App for Splunk Enterprise and Technology Add-On for NetApp SANtricity from within the Splunk environment. Additional configuration is required to add each array to NetApp SANtricity Web Services Proxy and to Splunk. For more information, see the README included with each app or the [Details](#) page on SplunkBase.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

NetApp documentation

- NetApp E-Series SANtricity Software datasheet
<https://www.netapp.com/pdf.html?item=/media/19775-ds-3171-66862.pdf>
- SANtricity Release 11.70 Documentation Center
<http://docs.netapp.com/ess-11/index.jsp?topic=%2Fcom.netapp.doc.ssm-sam-117%2Fhome.html>
- TR-4652: SANtricity Dynamic Disk Pools Best Practice Guide
<https://www.netapp.com/media/12421-tr4652.pdf?v=113202090653P>
- TR-4474: SANtricity drive security
<https://www.netapp.com/pdf.html?item=/media/17162-tr4474.pdf>
- Interoperability Matrix Tool
<https://mysupport.netapp.com/matrix/#welcome>
- NetApp SANtricity Web Services Proxy
<https://mysupport.netapp.com/site/products/all/details/eseries-webservices/downloads-tab>

Splunk documentation

- Splunk>docs
<https://docs.splunk.com/Documentation>
- Splunk>answers
<https://community.splunk.com/t5/Community/ct-p/en-us>
- Splunk Enterprise Installation Manual
<https://docs.splunk.com/Documentation/Splunk/latest/Installation/Whatsinthismanual>
- Splunk Enterprise Capacity Planning Manual
<https://docs.splunk.com/Documentation/Splunk/8.2.0/Capacity/IntroductiontocapacityplanningforSplunkEnterprise>
- Splunk Enterprise Managing Indexers and Clusters of Indexers
<https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Aboutmanagingindexes>
- NetApp SANtricity Performance App for Splunk Enterprise
<https://splunkbase.splunk.com/app/1932/>
- Technology Add-On for NetApp SANtricity
<https://splunkbase.splunk.com/app/1933/>

Version history

Version	Date	Document version history
Version 1.0	July 2021	Initial release.

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 © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4903-0721