



Technical Report

NetApp ONTAP and Splunk Enterprise

ONTAP Performance and Reliability in a Splunk Enterprise Environment

Karthikeyan Nagalingam, NetApp
January 2019 | TR-4650

Abstract

This document presents the performance and reliability validation test results for NetApp® ONTAP® in a Splunk Enterprise environment. It also includes storage efficiency test results for Splunk indexer data. This document also includes test procedures and results for verifying healthy Splunk responses while restoring deleted data from Snapshot copy and when recovering from controller and disk failures.

TABLE OF CONTENTS

1	Big Data Performance and Reliability Challenges	5
1.1	Machine Data	5
1.2	Challenges	5
2	Solution	5
2.1	Solution Architecture Details	7
2.2	Solution Validation	7
2.3	Solution Positioning	8
3	Technology Overview.....	8
3.1	NetApp ONTAP 9.....	8
3.2	Splunk Enterprise.....	9
3.3	NetApp AFF Storage Arrays	9
3.4	NetApp ONTAP Snapshot and SnapRestore Technologies	9
3.5	NetApp FlexClone Volumes	9
3.6	NetApp Storage Efficiency	9
3.7	End-to-End NVMe Support with NetApp ONTAP	9
3.8	Splunk App for NetApp ONTAP	10
4	Test Plan Summary	10
4.1	Test Plan Procedure	10
5	Test Results Summary	11
5.1	ONTAP Features Validation with Splunk	11
5.2	Splunk Indexing Rate Performance Comparison Summary (Using AFF A800)	12
5.3	Splunk Search Performance for AFF A800 versus DAS Summary	12
5.4	NetApp ONTAP Data Storage Efficiency with Splunk Enterprise	13
6	Test Configuration Details.....	14
6.1	Configuration Used for ONTAP Features and Resilience Testing.....	14
6.2	Indexer Storage Configuration with AFF A700.....	15
6.3	Configuration Used for Performance Testing and Storage Efficiency Validation (AFF A800)	16
6.4	Storage Provisioning for iSCSI.....	17
6.5	Storage Provisioning for NVMe/FC	18
6.6	Storage Configuration for Internal DAS Deployment.....	19
7	Test Procedure and Detailed Results	20
7.1	AFF A700 – FC SAN.....	20
7.2	AFF A800 – iSCSI, FC SAN, NVMe/FC	20

7.3 Backup and Restore Test – AFF A700	20
7.4 Storage System Resilience Test (Controller Failure) – AFF A700	24
7.5 Disk Failure and Reconstruct Test – AFF A700	27
7.6 Splunk Indexer Failure Test – AFF A700	30
7.7 Creation of Data Copies with FlexClone Technology Test – AFF A700	31
7.8 Index Performance Testing	34
7.9 Performance Baseline Test: Splunk Search Performance	38
7.10 ONTAP Storage Efficiency Test	40
8 Recommendations.....	44
8.1 Splunk Settings	44
8.2 Network Configuration	44
8.3 NFS for Splunk Storage	44
8.4 Server-Side Configuration	44
Appendix A: Splunk Use Cases	44
Appendix B: Test Configuration Details.....	45
Appendix C: Server Hardware	46
Servers Used with AFF A700 FC SAN Configuration	46
Servers Used with AFF A800 and Internal DAS Configuration	46
Appendix D: Server OS Details – AFF A800 Configuration	46
OS Kernel Settings	46
Logical Volumes and File System Configuration, Mount Options	47
9 Conclusion	47
Acknowledgments	49
Where to Find Additional Information	49
Version History	49

LIST OF TABLES

Table 1) Tested solution architecture base components	7
Table 2) ONTAP and Splunk Enterprise features and validation test results summary (AFF A700).	11
Table 3) Splunk data ingest performance for AFF A800 versus DAS.....	12
Table 4) Search types with definitions and application in the tests that were performed.....	12
Table 5) Splunk search performance for AFF A800 versus DAS.	13
Table 6) Space savings from ONTAP storage efficiency.....	13
Table 7) LUN provisioning details.....	15

Table 8) Storage provisioning for iSCSI and FC SAN.	17
Table 9) Storage provisioning for NVMe/FC.....	18
Table 10) Splunk indexing rate performance comparison for 1TB data ingest.	35
Table 11) Average indexer CPU utilization during data ingest.	37
Table 12) Splunk search performance: AFF A700 versus DAS.....	38
Table 13) Storage efficiency for Splunk Indexer data summary.	43

LIST OF FIGURES

Figure 1) NetApp ONTAP and Splunk Enterprise solution.	6
Figure 2) Network topology of configuration used for ONTAP features and resilience testing.	14
Figure 3) Network topology used for AFF A800 performance testing and storage efficiency validation.	17
Figure 4) Network topology used for Splunk internal DAS performance.	19
Figure 5) Internal DAS configuration for Splunk performance testing.....	19
Figure 6) Average indexing rate comparison: AFF A800 versus DAS.	36
Figure 7) Peak indexing rate comparison: AFF A800 versus DAS.	36
Figure 8) Graphical example of peak and average index performance.	37
Figure 9) Dense search during data ingest completion time comparison.	39
Figure 10) Dense search without data ingest completion time comparison.	39
Figure 11) Sparse search during data ingest completion time comparison.	40
Figure 12) Sparse search without data ingest completion time comparison.....	40
Figure 13) Savings from storage efficiency as shown in OnCommand System Manager (only data).	41
Figure 14) ONTAP savings from storage efficiency details.	42
Figure 15) Savings as shown in OnCommand System Manager (data + 1 Snapshot copy).	42
Figure 16) Savings as shown in OnCommand System Manager (data + 1 Snapshot copy + 1 FlexClone).	43

1 Big Data Performance and Reliability Challenges

It's been estimated that data will grow exponentially well beyond the year 2020. Some studies predict that by the year 2025, the amount of data in existence will have grown by a factor of 10 times the amount in 2017. This data comes in many forms, ranging from semistructured to totally unstructured. Sources include the Internet of Things (IoT), business applications, social media, customer behavior, and machine sensors, to name just a few.

1.1 Machine Data

In this age of digital transformation, *machine data* is one of the key drivers that is fueling that growth. Machine data is generated by technology infrastructures, security systems, and business applications. It is one of the fastest growing types of data, and it's also one of the most complex categories of big data.

The format of machine data is unpredictable, coming from so many different sources, at such a high rate, and in such great volumes, that it's often referred to as *digital exhaust*. It's constantly being generated by servers, server infrastructures, applications, sensors, electronics, buildings, security systems, and all the elements that make up the IoT. Machine data is tremendously valuable, in that it contains records of customer behavior, transactions, diagnostics from critical mechanical systems, message queues, change events—and the list goes on. But it is difficult to unlock the value in this data due to its high volume and lack of structure.

Enterprise organizations depend on machine data to help run their businesses, to meet competitive demands in the marketplace, and to avoid costly infrastructure downtime. The ability to collect and analyze this data is key to transportation safety, machine reliability, fraud detection, and security. For example, in the healthcare industry, critical medical devices are monitored in real time and seconds can mean the difference between life and death. As another example, real-time analytics can help prevent financial fraud and attacks on sensitive computer systems. If a security breach has already occurred, or in case of a cyberattack, analytics can help identify the source and limit the damage. For an enterprise guarding against security breaches and intrusions, microseconds matter.

1.2 Challenges

Splunk Enterprise provides the tools and capabilities that allow an enterprise organization to collect that data and extract high value from it. Those tools and capabilities include visualization, fast data ingest, real-time analytics, a rich set of APIs, notification capabilities, and extreme scalability.

However, traditional Splunk deployments with direct-attached storage (DAS) are subject to server sprawl. The Splunk best practice recommendation to configure server storage with mirroring means that only half of a server's storage capacity is available for data. Adding storage capacity requires servers to be added, even if additional compute capacity is not needed. This type of configuration leads to poor storage efficiency and a perpetual imbalance between compute and storage, because they cannot be scaled independently.

Splunk also relies on traditional methods for data protection which are slow and consume valuable storage and compute resources. Backups can't be created often enough to meet low recovery point objectives (RPOs), and recovery takes too long to meet recovery time objectives (RTOs).

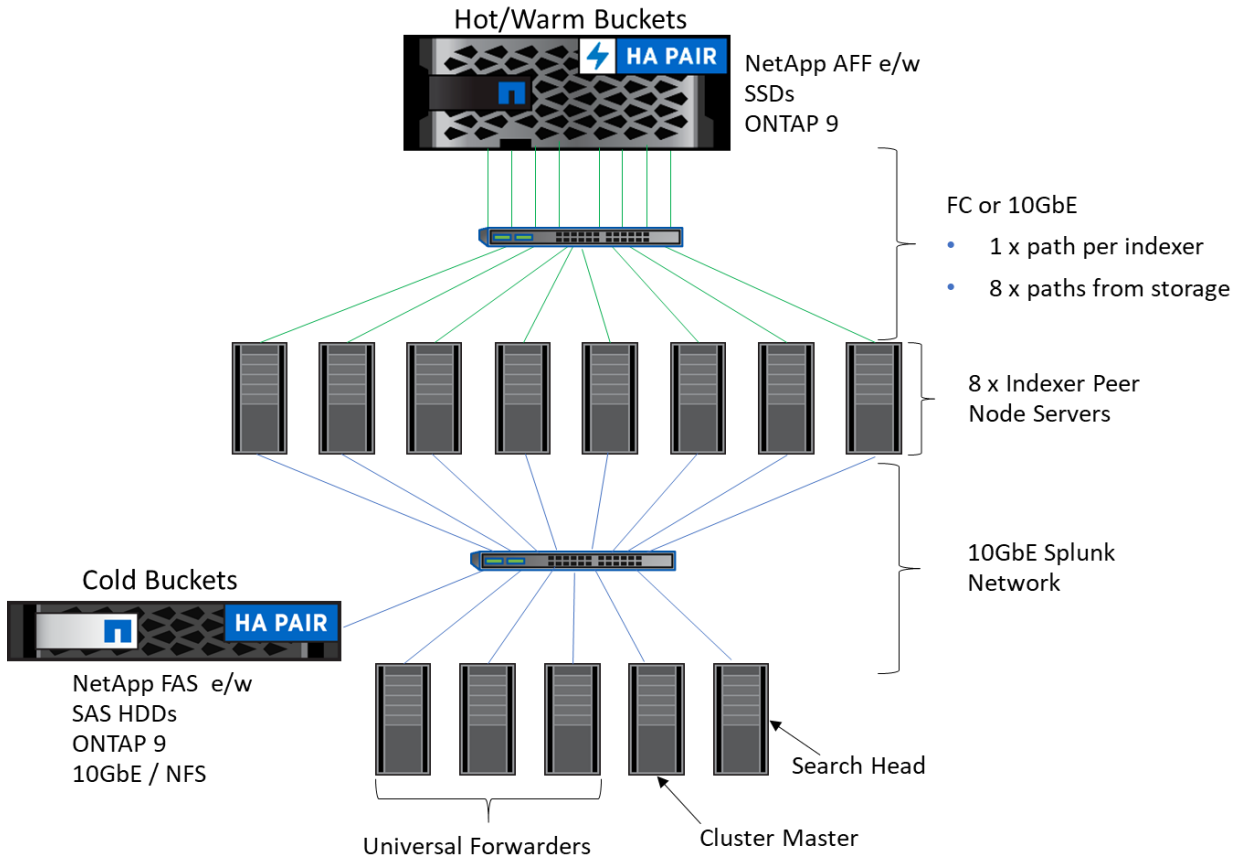
2 Solution

The basic NetApp ONTAP and Splunk Enterprise solution, shown in Figure 1, is an optimal data protection and data management platform for enterprise-class organizations to capture and analyze machine data.

Storing Splunk data on NetApp FAS and AFF powered by ONTAP 9 software instead of on internal storage media, decouples compute and data storage resources, enabling enterprises to create efficient

configurations that meet their needs today and tomorrow. Servers need to be added only when additional compute resources are required and storage can be scaled out independently of compute. This solution conserves valuable rack space, floor space, and energy. Figure 1 gives a graphical representation of the configuration tested.

Figure 1) NetApp ONTAP and Splunk Enterprise solution.



In addition to independent scaling of resources, ONTAP software provides the data protection, data governance, storage efficiency, and copy management features that are needed to meet the requirements of enterprise organizations that use Splunk.

The NetApp ONTAP and Splunk Enterprise solution provides the following enterprise requirements:

- 100% uptime and data availability
- Robust data protection to meet low RPOs and RTOs, with zero loss of data
- Full support for DR
- Cost-effective support for DevOps
- Data security
- SEC-compliant backups
- High performance
- Seamless integration with public and private cloud for scale and agility
- Full scalability of data storage
- Storage efficiency to meet cost objectives

2.1 Solution Architecture Details

The basic architecture for the NetApp ONTAP and Splunk Enterprise solution includes network infrastructure, Splunk server functionality, storage provisioning details, and Linux host-side storage configuration.

Note: See section 3 for detailed information about the technologies used in this solution.

Figure 1 depicts the tested solution architecture and Table 1 lists its base components.

Table 1) Tested solution architecture base components.

Solution Components	Details
Splunk versions 6.6.3 and 7.1.1	<ul style="list-style-type: none">• Eight indexer servers• Three universal forwarders• One cluster master• One search head• Splunk search factor: 2• Splunk replication factor: 2
Linux (RHEL 7.2 and SLES 12.3)	All servers
NetApp AFF storage array HA pair for hot, warm buckets	<ul style="list-style-type: none">• ONTAP 9• 24 x 960GB SSDs• FC, iSCSI, or NVMe over Fibre Channel (NVMe/FC) protocols• 16 (used with AFF A700 or 32Gb (used with AFF A800) FC, 8 ports• --Or --• 10GbE, 8 ports
NetApp FAS2552 storage array HA pair for cold buckets	<ul style="list-style-type: none">• ONTAP 9• 32 x 450GB SAS HDDs• 10GbE, 8 ports used• NFS protocol
13 Fujitsu PRIMERGY RX2540 servers	Each equipped with: <ul style="list-style-type: none">• 2 CPUs, 16 physical cores total• Intel Xeon• 256GB physical memory• SAN booted• 16Gbps or 32Gbps FC, dual controller• 10GbE dual port

For testing ONTAP resilience, data protection, and other features, Splunk 6.6.3 was used with RHEL 7.2 and AFF700. An AFF800 and Splunk 7.1.1 with SLES 12.3 were used for performance testing.

2.2 Solution Validation

Four variations of the solution architecture (Figure 1) were tested to show that a NetApp AFF with ONTAP 9 offers an optimal data protection and data management platform for enterprise-class organizations that are Splunk to capture and analyze machine data.

This technical report describes the solution and covers the following information:

- Variation 1: Splunk deployed with AFF A700 FC SAN
 - Reliability data for NetApp ONTAP in a Splunk Enterprise environment
 - Test results for verifying healthy Splunk responses when restoring deleted data from Snapshot copies
 - Test results for recovering from a controller and disk failure
- Variations 2-4: Splunk deployed with AFF A800 using FC SAN, NVMe/FC, and iSCSI
 - Splunk log data ingestion performance
 - Splunk search performance
 - Comparison between the AFF A800 configurations and a typical Splunk configuration that uses commodity servers and internal storage media

Section 4 of this document contains the test plan summary. A summary of the test results and can be found in Section 5. Details of the test configuration, test procedures, and test results are in Sections 6 and 7.

2.3 Solution Positioning

The solution described in this document is not the only NetApp solution for Splunk. There is also a solution based on the NetApp E5700 storage array, which is documented in [TR-4623: NetApp E-Series E5700 and Splunk Enterprise](#). The E-Series solution provides excellent performance and flexibility at a competitive cost. The AFF A700 and AFF A800 solutions described in this technical report are designed for environments that require enterprise-class data management features. Those features include fast backup and restore capabilities with storage-efficient ONTAP Snapshot technology, storage efficiency with deduplication and compression, Data Fabric enablement, and DevOps support with NetApp FlexClone® volumes.

Although the solutions described in this document use the AFF A800, AFF A700, and FAS 2552, other NetApp AFF arrays can be used for Splunk hot/warm buckets, and other FAS systems can be used for Splunk cold bucket storage, depending on workload requirements.

3 Technology Overview

This section describes the technology used in the NetApp ONTAP and Splunk Enterprise solution.

3.1 NetApp ONTAP 9

NetApp ONTAP 9 data management software is an optimal solution for a big data platform. ONTAP is the enterprise data management software that powers the NetApp AFF and FAS systems, and software-only Cloud Volumes ONTAP. ONTAP is designed to meet the needs of the entire enterprise. The ONTAP software is capable of hybrid web-scale deployments, is highly available and resilient, and provides the data management capabilities that enterprises require.

NetApp is the data authority for hybrid cloud, data protection, data availability, and copy management. NetApp empowers customers to simplify and integrate data management across cloud and on-premises environments to accelerate digital transformation. Together with its partners, NetApp offers a full range of hybrid cloud data services to help global organizations unleash the full potential of their data to expand customer touchpoints, foster greater innovation, and optimize their operations.

3.2 Splunk Enterprise

Splunk Enterprise provides the platform for collecting, indexing, and analyzing machine data from any source to deliver operational intelligence, which can be used to optimize IT, security, and business performance. It provides powerful search, analysis, and visualization capabilities that can be accessed across organizations, and it is available as on-premises software or as a cloud service.

Splunk uses a distributed search framework that scales linearly. Its implementation of MapReduce enables large-scale search, reporting, and alerting. The Splunk platform is open and has software development kits (SDKs) and APIs, including a REST API and SDKs for Python, Java, JavaScript, PHP, Ruby, and C#. It also provides node failover and workload balancing across components.

For a discussion of Splunk use cases, see appendix A.

3.3 NetApp AFF Storage Arrays

NetApp AFF systems address enterprise storage requirements with high performance, flexibility, and best-in-class data management. Built on ONTAP software and designed specifically for flash, AFF systems deliver industry-leading performance, capacity density, scalability, security, and network connectivity. NetApp AFF storage provides both 40GbE and 32Gbps FC connectivity. AFF systems are Data Fabric ready, with proven cloud connectivity, enabling you to move workloads where they run best and data where it's needed.

With NetApp RAID DP® technology, AFF systems provide industry-leading, in-place data protection. Also, with the NetApp flash-optimized WAFL® system and enhanced built-in quality of service (QoS), consistent high performance at 1ms latencies and lower is achieved.

AFF systems are scalable and highly available. For more information about the enterprise-grade data management features of AFF systems, see the [NetApp AFF Datasheet](#).

3.4 NetApp ONTAP Snapshot and SnapRestore Technologies

With ONTAP Snapshot technology, you can create point-in-time data copies with no impact on performance and with minimal consumption of storage space. You can create these Snapshot copies almost instantaneously and use them with NetApp SnapRestore® software to recover entire file systems or data volumes in seconds.

3.5 NetApp FlexClone Volumes

FlexClone volumes are space-efficient, writable data copies that can be created almost instantly, anywhere in the Data Fabric or hybrid cloud where a Snapshot copy exists. FC, iSCSI, and NFS are supported protocols, and they are perfect for DevOps. FlexClone technology makes it possible for developers, QA engineers, and software testers to work with real production data. They can be easily created and deleted, and any changes made to them have zero effect on the parent production data. If necessary, they can even be split from the parent production data and promoted to production. FlexClone volumes can be considered as free data copies.

3.6 NetApp Storage Efficiency

ONTAP 9 offers inline deduplication, compression, and compaction. Whether written to on-premises or cloud storage, the data occupies less space, which translates to lower data storage costs.

3.7 End-to-End NVMe Support with NetApp ONTAP

SSDs deliver I/O only as fast as the protocols used to access them. Traditional storage protocols, such as SAS, create a bottleneck, preventing the user from reaping all the benefits of flash storage. NVMe (non-volatile memory express), which was created exclusively for flash, removes that bottleneck. With NVMe, the transfer protocol no longer gets in the way of the low latency and high throughput of flash storage.

NetApp is the first data management company to offer end-to-end NVMe. The SSDs themselves are configured to use NVMe to connect with the ONTAP AFF controller, and the controller front-end is connected to the storage network using NVMe. Implementation is nondisruptive in that it uses existing FC storage infrastructure. In other words, NVMe/FC uses the same FC SAN as FC-SCSI (FCP), and does not require hardware, cable, or switch changes. Also, enabling NVMe/FC does not disrupt the operation of existing storage network protocols. NVMe/FC and FCP can coexist on the same FC SAN without any problems.

Note that NVMe over FC is often referred to as NVMe/FC and FC-NVMe. Those abbreviations are interchangeable; however, the current standard and trademarked term NVMe/FC.

For more information about NVMe/FC, see [TR-4684: Implementing and Configuring Modern SANs with NVMe/FC](#).

3.8 Splunk App for NetApp ONTAP

The Splunk App for NetApp ONTAP enables you to visualize the configuration, performance, and syslog events for all ONTAP storage arrays in your Splunk deployment. With the Splunk App for NetApp ONTAP, you can do the following from a single pane of glass:

- Reduce problem investigation and resolution times.
- Gain real-time insights into key performance metrics, anomalies, and outliers across all your storage systems and configured subsystems.
- Improve your storage monitoring efficiency and proactively plan storage capacity allocations with more than 30 out-of-the-box, customizable reports.
- Correlate data from all NetApp systems with data from operating systems, applications, networks, and virtual and physical infrastructure for enterprise-wide 360-degree visibility.
- Get central proactive monitoring of ONTAP systems, including real-time notification of important ONTAP events.

For more information, see [Splunk App for NetApp ONTAP](#).

4 Test Plan Summary

We executed a test plan to prove that ONTAP is the enterprise-class data management and protection platform for Splunk. We demonstrated most of the enterprise requirements listed in section 2. Some requirements were omitted because they are considered more generic to enterprise-class storage.

4.1 Test Plan Procedure

We used a custom script provided by a third-party consultant to create the log files used for workflow generation throughout all tests. This section describes the validation test plan.

ONTAP Features Testing and Validation Using an ONTAP AFF A700 Deployment

The following tests were performed to test and validate AFF A700 deployment with Splunk:

1. Backup and restore operations by using NetApp Snapshot and SnapRestore technologies.
2. Storage resilience tests:
 - a. Failure of a storage node by generating a system panic:
 - Takeover
 - Giveback
 - b. Disk failure and reconstruction.
 - c. Capture Splunk indexer failure: four indexer failures, one after the other.

- d. Create a cloned Splunk environment by using NetApp FlexClone technology.

Performance Testing Using an ONTAP AFF A800 Deployment

1. The following tests were performed on the AFF A800 deployment with Splunk to compare it with the internal DAS deployment.
1. Determine the baseline Splunk performance of the internal DAS Splunk deployment.
 - a. Use the custom script mentioned previously to generate 1TB of log data, spread across three universal forwarders.
 - b. Ingest the logs from the forwarders into the indexer cluster and capture the performance results.
 - c. Baseline search performance with the index created in the previous step (1b), and record the time required for completion of each search. Search types to be tested include dense, sparse, super sparse, and rare.
2. Test Splunk performance on NetApp AFF A800 deployments that use SAN protocols, including NVMe/FC, FC SAN, and iSCSI. Capture and record the following metrics for each deployment:
 - a. Data ingest performance results from Splunk.
 - b. ONTAP performance metrics for each data ingest operation. This data is used for validation of Splunk reported results and to help in determining tuning best practices.
 - c. Completion times for all searches listed in step 1c.
 - d. ONTAP storage efficiency metrics for Splunk data stored on the AFF A800.
 - i) For data only, no Snapshot copies or FlexClone volumes.
 - ii) For data with one Snapshot copy per storage volume, no FlexClone volumes.
 - iii) For data with one Snapshot copy and one FlexClone for each storage volume.

5 Test Results Summary

Each test case completed successfully and fully supported ONTAP as the enterprise-class data management and protection platform for Splunk. This section contains a summary of all test results. For detailed results, see section 7.

5.1 ONTAP Features Validation with Splunk

This section summarizes the ONTAP features validation tests that were performed with Splunk Enterprise, as summarized in section 4.1.

Table 2 shows a summary of the results from our ONTAP features testing and validation performed with Splunk Enterprise and a NetApp AFF A700 storage array. Each test ran without any problems and the results were as expected.

Table 2) ONTAP and Splunk Enterprise features and validation test results summary (AFF A700).

Test Description	Results Summary
Snapshot copy creation	No visible effect on performance.
SnapRestore	Successful. After the data restore operation, the Splunk search result was the same as before the data loss.
Storage controller panic	No observed impact on Splunk performance; search completed successfully.
Disk failure and reconstruct	No observed effect on Splunk performance; search completed successfully.

Test Description	Results Summary
Splunk indexer failure	Minimal effect on Splunk performance; minor spikes during failure event.
FlexClone volume creation	No observed effect on Splunk performance; clone configuration for Splunk completed.

Note: For the Snapshot copy creation test, we created a regular Snapshot copy. However, NetApp strongly recommends creating a consistency group (CG) Snapshot copy.

5.2 Splunk Indexing Rate Performance Comparison Summary (Using AFF A800)

Table 3 lists the performance results from the average and peak indexing rates observed during the benchmark phase of the tests.

The AFF A800 outperformed internal DAS for each protocol tested. As expected, the peak indexing rate was considerably higher than the average rate, meaning that there was variation in the indexing rate during the indexing performance test, which is not unusual.

Table 3) Splunk data ingest performance for AFF A800 versus DAS.

Average Indexing Rate for 1TB (AFF A800 versus Internal DAS Configuration)		
Configuration Summary	Avg Rate Observed Versus DAS	Peak Rate Observed Versus DAS
FC SAN with eight indexer peer nodes	5% faster than DAS	7% faster than DAS
iSCSI with eight indexer peer nodes	5% faster than DAS	7% faster than DAS
NVMe/FC (32Gbps FC) with eight indexer peer nodes	6% faster than DAS	7% faster than DAS

5.3 Splunk Search Performance for AFF A800 versus DAS Summary

In addition to the indexing rate, we also tested the search performance. We used indexes created during our indexing performance tests to support the four basic types of searches, described in Table 4.

Table 4) Search types with definitions and application in the tests that were performed.

Search Type	% of Occurrences of Search Term	% Matched in Tested Searches
Dense	Greater or equal to 1%	1%
Sparse	Between 0.01% and 1% of events matched	0.1%
Super Sparse	Between 0.0001% and 0.01%	0.0001%
Rare	Between 0.00001% and 0.0001%	0.00001%

The middle column defines the search type in terms of % of occurrences found in the data searched. For example, a dense search, by definition, will find several occurrences equal to or greater than 1% of relevant events in the dataset that is searched.

The third column of Table 4 shows the actual occurrence rate of each search type executed as part of this test plan.

The metric of interest in this set of tests is the search completion time in seconds. Table 5 shows a comparison of search performance for each AFF A800 protocol tested and compares the results as to an identical Splunk deployment that uses internal direct attached SSDs and disks (DAS). For details about these results, see section 5. Note that the AFF A800 configurations performed better than internal DAS for all dense and sparse searches executed. Performance was the same for super sparse and rare searches with all Splunk deployments.

Table 5) Splunk search performance for AFF A800 versus DAS.

Search Time Comparison – AFF A800 versus Internal DAS								
	With Ingest				With No Ingest			
	Dense	Sparse	Super Sparse	Rare	Dense	Sparse	Super Sparse	Rare
FC SAN	18% faster	17% faster	Same as DAS	Same as DAS	17% faster	17% faster	Same as DAS	Same as DAS
NVMe/FC	14% faster	17% faster	Same as DAS	Same as DAS	12% faster	16% faster	Same as DAS	Same as DAS
iSCSI	12% faster	17% faster	Same as DAS	Same as DAS	10% faster	17% faster	Same as DAS	Same as DAS

5.4 NetApp ONTAP Data Storage Efficiency with Splunk Enterprise

Another benefit offered by the NetApp AFF storage system is storage efficiency. As described previously, NetApp AFF storage arrays, by default, provide inline compression, and deduplication to reduce the footprint of data stored. This is important to reduce data sprawl. Table 6 shows the storage efficiencies achieved for the following:

1. Data only, which is the data reduction ratio of Splunk data without Snapshot copies or FlexClone volumes.
2. Data plus one Snapshot copy for each storage volume.
3. Data plus one Snapshot copy and one FlexClone for each storage volume.

Table 6) Space savings from ONTAP storage efficiency.

Savings from Storage Efficiency Summary	
Description	Storage Efficiency Ratio
Data Only	1.3:1
Data + one Snapshot copy per volume	2.6:1
Data + one Snapshot copy + one FlexClone per volume	3.88:1

For data only, savings came from a combination of volume deduplication, volume compression, aggregate-level deduplication, and aggregate-level compaction.

Inline storage efficiency is enabled automatically for NetApp AFF storage arrays, and has negligible, if any, effect on performance.

The multiplying effect on storage efficiency of using ONTAP Snapshot and FlexClone technologies should be noted. Creating one Snapshot copy backup of each storage volume doubles the storage efficiency achieved for data alone. Creating one Snapshot copy and one FlexClone volume of each storage volume nearly triple the storage-efficiency ratio.

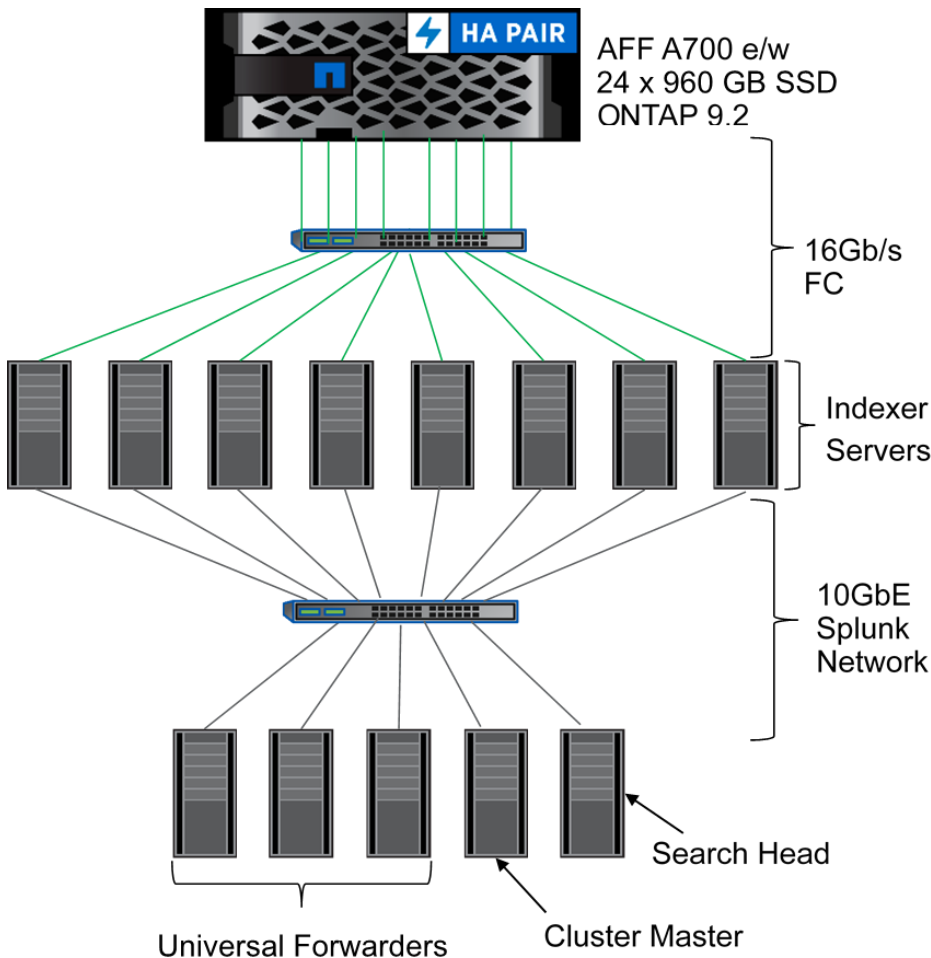
6 Test Configuration Details

This section describes the tested configurations: the network infrastructure, Splunk server functionality, storage provisioning details, and Linux host-side storage configuration.

6.1 Configuration Used for ONTAP Features and Resilience Testing

To test the ONTAP features and resilience, an AFF A700 running ONTAP 9.2 was used. The test configuration, shown in Figure 2, includes eight Splunk indexer servers connected to the NetApp AFF A700 storage array over 16Gb FC. We connected the indexer servers to a Splunk 10GbE network of three forwarders, one cluster master (controller monitor), and search head RHEL 7.2 servers. Figure 2 contains a graphical representation of this configuration. For these tests, we used Fujitsu PRIMERGY RX2540 M1 servers running RHEL 7.2. For information about the components used in the architecture, see Table 1. In this test, Splunk version 6.6.3 was used.

Figure 2) Network topology of configuration used for ONTAP features and resilience testing.



The components, shown in Figure 2 and listed in Table 1, provide the following functions:

- The indexer servers were configured as a cluster. Each of the eight servers was configured as an indexer peer node. Incoming data is balanced across the peer nodes. Each peer indexes the data it receives and sends or receives data replicas per the search factor and replication factor. Searches are directed to the appropriate peers and the results are aggregated by the search head.
- The forwarders consume data from external sources and forward it to the indexer servers (or cluster).

- The search head manages searches across the cluster of indexer servers. It distributes the search queries to the indexer servers and consolidates the results.

Note: All searches are run from the search head. Each indexer cluster must have at least one search head.

- The cluster master (controller monitor) manages the Splunk indexer cluster. It coordinates replication activities of the indexer peer node servers and communicates with the search head for information about where to locate data for searches. It also remediates activities if an indexer server goes offline.

Note: Each cluster has only one controller monitor.

- As the indexer peer nodes receive data, they extract the relevant events and create an index of those events. The event data itself is written to a file referred to as the raw data file. As those events are written, an index file is also created. The raw data file is automatically compressed by Splunk. The index file is not compressed. An index is made up of a compressed raw data file and an index file. This is also referred to as a searchable copy of the data. As events are written to the raw data file, they are also directed to other peer nodes to support replication. The replication factor determines how many copies of the raw data are created. For indexer servers using server-based internal disks or SSDs, the recommended replication factor is 3 or more. In the same manner, one or more copies of the index file are replicated across the indexer cluster. The search factor determines how many searchable copies of the index will be created. Splunk recommends a search factor of at least 2, to support indexer peer node failover. Since the search factor determines the number of raw data copies and the number of index file copies, the replication factor must be greater than or equal to the search factor. Splunk's best practice of setting the replication factor of 3 is to provide additional data protection, beyond the search factor. Because NetApp storage systems provide the required level of data protection (resilience), NetApp recommends a replication factor equal to the search factor. That being the case, NetApp recommends a replication factor and search factor of 2.

6.2 Indexer Storage Configuration with AFF A700

We assigned one FC LUN to each indexer. The LUNs were configured with one LUN per storage volume, with four volumes in aggregate 1 (aggr1) and the other four volumes in aggregate 2 (aggr2). Aggr1 was configured on the first storage controller and aggr2 was configured on the second controller. The LUN size was 550GB, with the corresponding volumes being 600GB in size. Table 7 lists the provisioning details for the LUN used by the indexers.

Table 7) LUN provisioning details.

Controller	Aggregate	Aggregate Size	Volume	LUN	Volume Size	LUN Size
Controller 1	Aggr1	6.91TB	Vol1	LUN1	600GB	550GB
			Vol2	LUN2	600GB	550GB
			Vol3	LUN3	600GB	550GB
			Vol4	LUN4	600GB	550GB
Controller 2	Aggr2	6.91TB	Vol5	LUN5	600GB	550GB
			Vol6	LUN6	600GB	550GB
			Vol7	LUN7	600GB	550GB
			Vol8	LUN8	600GB	550GB

We created both aggregates (aggr1 and aggr2) with an SSD count of 23.

After the storage was provisioned and mapped to the indexers, we completed the following tasks as part of the test:

1. Create a single logical volume group on each LUN, one volume group per indexer.
2. Create one logical volume on each volume group.
3. Create an XFS file system on each logical volume.
4. Create a mount point for the XFS file system named `/splunk` on each indexer.
5. Mount the new file system under the `/splunk` mount point by using the following options:
 - `noatime`
 - `nobarrier`
 - `nodiratime`

6.3 Configuration Used for Performance Testing and Storage Efficiency Validation (AFF A800)

For performance testing, we used the following four different Splunk configurations, using the same servers and the same Splunk 10GbE network:

1. Indexers that use iSCSI over 10GbE for storage
2. Indexers that use FC SAN over 32Gbps FC
3. Indexers that use NVMe/FC (32Gbps FC)
4. Indexers that use internal DAS (SSD and HDD) for index storage, for comparison to AFF A800 performance.

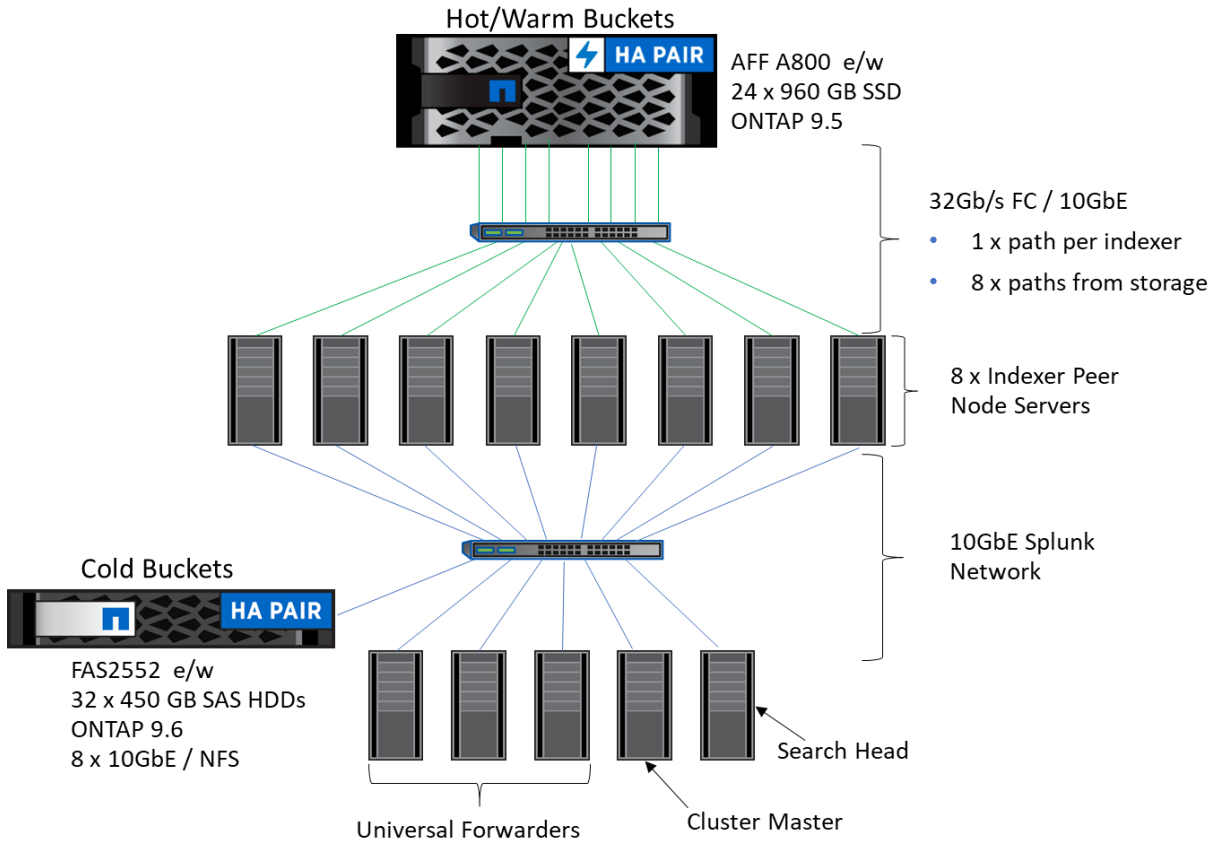
Each of these configurations is described in the sections that follow.

NetApp AFF800 Configuration

For Splunk performance testing and ONTAP storage-efficiency validation, an AFF A800 running ONTAP 9.5 was used for hot and warm buckets. A FAS2552 running ONTAP 9.6, with spinning disks (HDDs) was used for cold buckets (connected using NFS over 10GbE). The configuration tested, shown in Figure 2, includes eight Splunk indexer servers connected to the NetApp AFF A800 storage array over 32Gb FC for FC SAN and 10GbE for iSCSI, and a FAS2552 with HDDs for cold buckets. We connected the indexer servers to a Splunk 10GbE network of three forwarders, one cluster master, and search head SLES 12.3 servers. Figure 3 shows the details of that configuration. For this configuration, we used Fujitsu PRIMERGY RX2540 M4 servers running SLES 12.3. For information about the components used in the architecture, see Table 1.

Splunk version 7.1.1 was used.

Figure 3) Network topology used for AFF A800 performance testing and storage efficiency validation.



6.4 Storage Provisioning for iSCSI

For the iSCSI test, each indexer peer node accessed the AFF A800 over the 10GbE network for iSCSI and 32Gb for FC SAN described previously. Each of the peer nodes was assigned four 155GB LUNs. The Linux logical volume manager was used to configure a single logical volume from each set of four LUNs. A partition and an XFS file system were then created on the logical volume. Logical volume and XFS configuration details are included in appendix B.

Table 8 details the storage provisioning of the AFF A800.

Table 8) Storage provisioning for iSCSI and FC SAN.

Controller	Aggregate	Aggregate Size	Volume	# LUNs / Volume	Volume Size	LUN Size (Each)
Controller 1	Aggr1	5.36TB	Vol1	4	700GB	155GB
			Vol2	4	700GB	155GB
			Vol3	4	700GB	155GB
			Vol4	4	700GB	155GB
Controller 2	Aggr2	5.36TB	Vol5	4	700GB	155GB
			Vol6	4	700GB	155GB

Controller	Aggregate	Aggregate Size	Volume	# LUNs / Volume	Volume Size	LUN Size (Each)
			Vol7	4	700GB	155GB
			Vol5	4	700GB	155GB

One 700GB volume was thin provisioned for each of the eight indexer peer nodes. Four 155GB LUNs were created in each volume and mapped to a single peer node. In other words, each peer node was assigned four 155GB LUNs. The LUNs for each peer node were created in a single 700GB volume.

6.5 Storage Provisioning for NVMe/FC

For NVMe/FC each indexer peer node accessed the AFF A800 over the 32Gbps FC network described previously and shown in Figure 3. Each indexer peer node accessed the AFF A800 over 32Gbps FC. Each of the peer nodes was assigned a single NVMe namespace. Functionally, an NVMe namespace is equivalent to a LUN. Each namespace was mapped to an NVMe subsystem, which in terms of function is much like an FC igroup. On the indexer server, a partition and an XFS file system were then created on each NVMe namespace. Because only one namespace was created for each indexer server, logical volume manager was not used. The configuration details of the XFS are included in appendix B. Table 9 lists the AFF A800 storage provisioning details.

Table 9) Storage provisioning for NVMe/FC.

Controller	Aggregate	Aggr Size	SVM	Volume	Volume Size	NVMe Namespace	Namespace Size	NVMe Sub-system
Controller 1	Aggr1	5.36 TB	NVMF01	Vol1	700GB	NVMF01_ns	700GB	NVMF01_sub
			NVMF02	Vol2	700GB	NVMF02_ns	700GB	NVMF02_sub
			NVMF03	Vol3	700GB	NVMF03_ns	700GB	NVMF03_sub
			NVMF04	Vol4	700GB	NVMF04_ns	700GB	NVMF04_sub
Controller 2	Aggr2	5.36 TB	NVMF05	Vol5	700GB	NVMF05_ns	700GB	NVMF05_sub
			NVMF06	Vol6	700GB	NVMF06_ns	700GB	NVMF06_sub
			NVMF07	Vol7	700GB	NVMF07_ns	700GB	NVMF07_sub
			NVMF08	Vol8	700GB	NVMF08_ns	700GB	NVMF08_sub

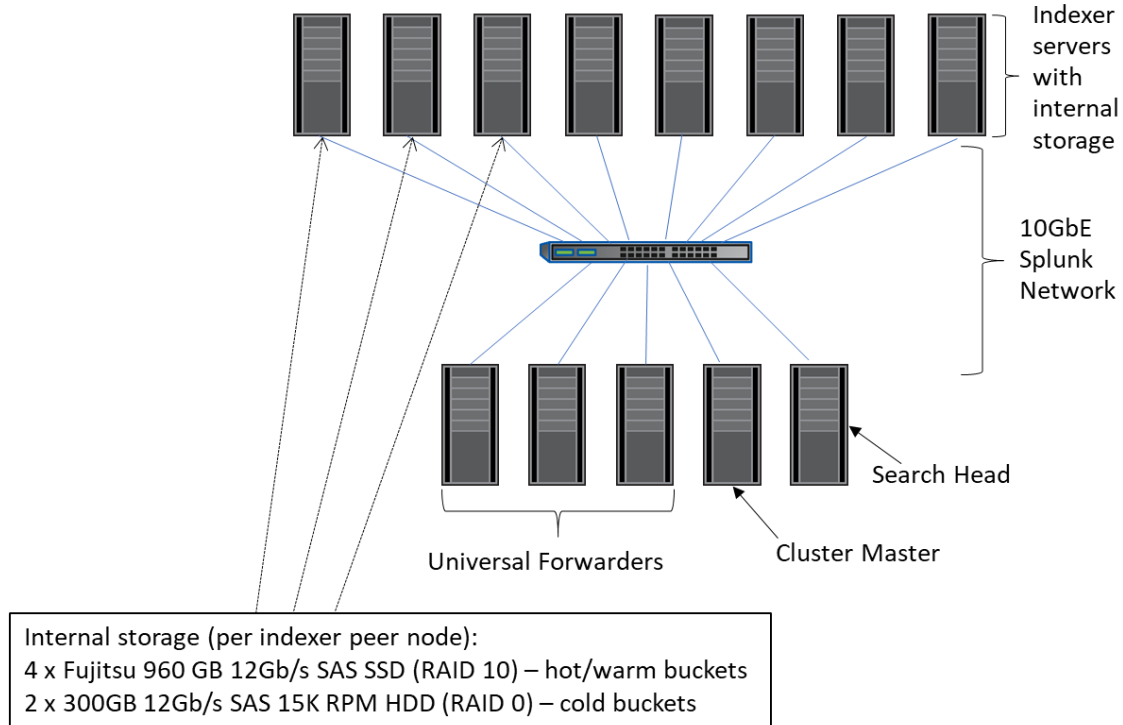
Just to clarify, one storage virtual machine (SVM) was created for each of eight NVMe namespaces. A single thin provisioned flexible volume was created in each SVM, and a single NVMe namespace was created in each volume. Each namespace was mapped to a single NVMe subsystem, with each subsystem mapped to a single indexer peer node.

For more information about NVMe/FC, see [TR-4684: Implementing and Configuring Modern SANs with NVMe/FC](#).

6.6 Storage Configuration for Internal DAS Deployment

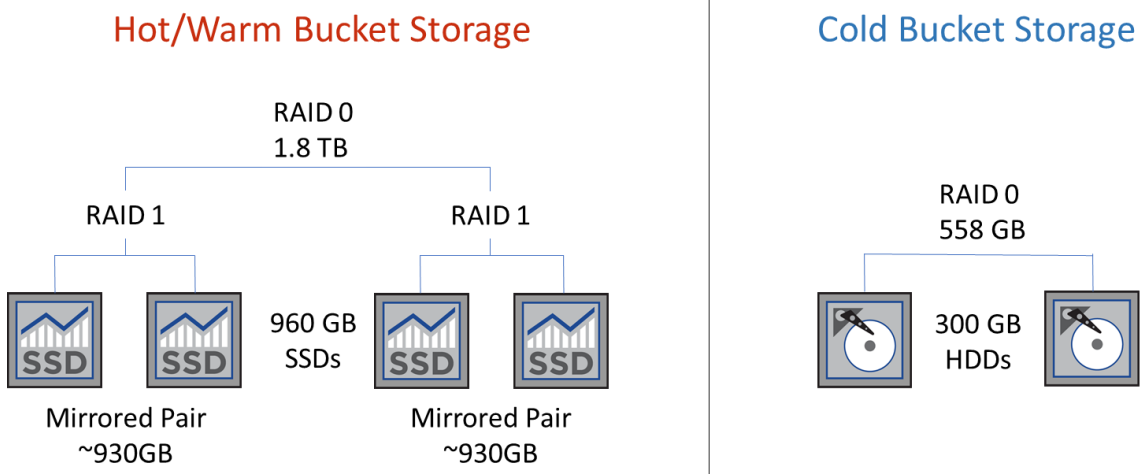
For performance comparison, eight indexer peer nodes were configured. Figure 4 illustrates this configuration.

Figure 4) Network topology used for Splunk internal DAS performance.



RAID configuration was performed at the hardware level by using the Fujitsu MegaRAID configuration utility during the hardware RAID controller BIOS startup. Figure 5 illustrates the internal DAS configuration in each indexer server for use with Splunk.

Figure 5) Internal DAS configuration for Splunk performance testing.



The 1.8TB RAID 10 device was partitioned and an XFS file system was created using the entire capacity.

For cold buckets, the 558GB RAID 0 device (depicted in Figure 5) was partitioned and an XFS file system was created using the entire capacity.

Details of file system configuration and the mount options can be found in appendix B.

7 Test Procedure and Detailed Results

Resilience testing was performed using an AFF A700 for hot/warm buckets. Performance tests were performed using an AFF A800 for hot/warm buckets, and a FAS2552 with spinning disks for cold buckets. Performance tests were repeated with indexer servers using internal DAS for comparison.

7.1 AFF A700 – FC SAN

The following describe the AFF Splunk configuration test plans:

1. Snapshot copy creation; assess its effect on performance.
2. Backup and restore using SnapRestore.
 - Test the capability of SnapRestore to successfully restore data from a Snapshot copy following catastrophic data loss.
3. Storage system resilience.
 - Test the impact of storage controller panic from running the Splunk query.
4. Disk failure and reconstruct.
 - Test the impact of running the Splunk query.
5. Splunk indexer failure.
6. Creation of data copies by using FlexClone.
 - Test the impact of running the Splunk query.
 - Test the ability to configure indexers to use clones.

7.2 AFF A800 – iSCSI, FC SAN, NVMe/FC

1. Index performance: AFF A800 versus internal DAS
2. Search performance: AFF A800 versus internal DAS
3. ONTAP storage efficiency validation for Splunk indexer data

The following sections describe each test case in detail.

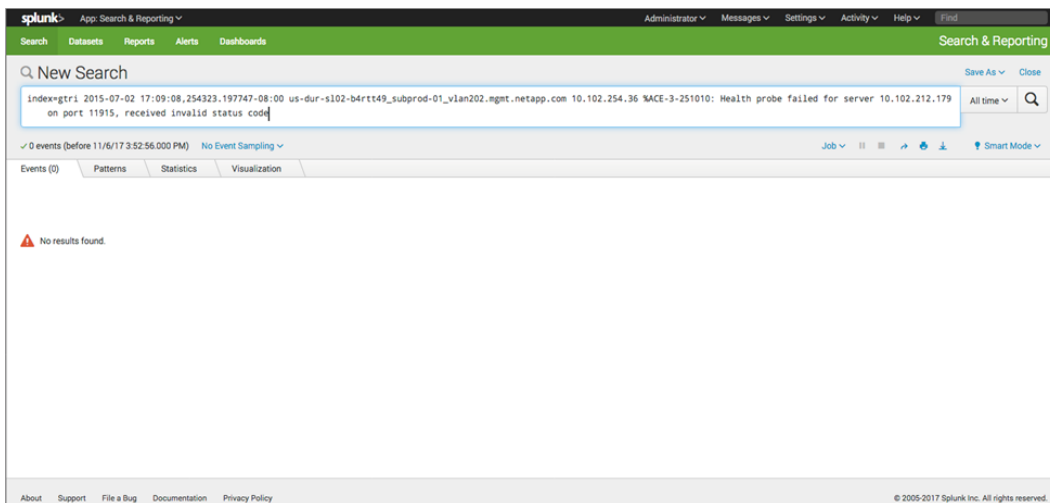
7.3 Backup and Restore Test – AFF A700

Test Details and Results

Backup Test

The following procedure was used for the backup test. The results are listed after each corresponding step:

1. Use OnCommand System Manager to create Snapshot copy backups of all eight data volumes.
2. Create the Snapshot copies while running the query.



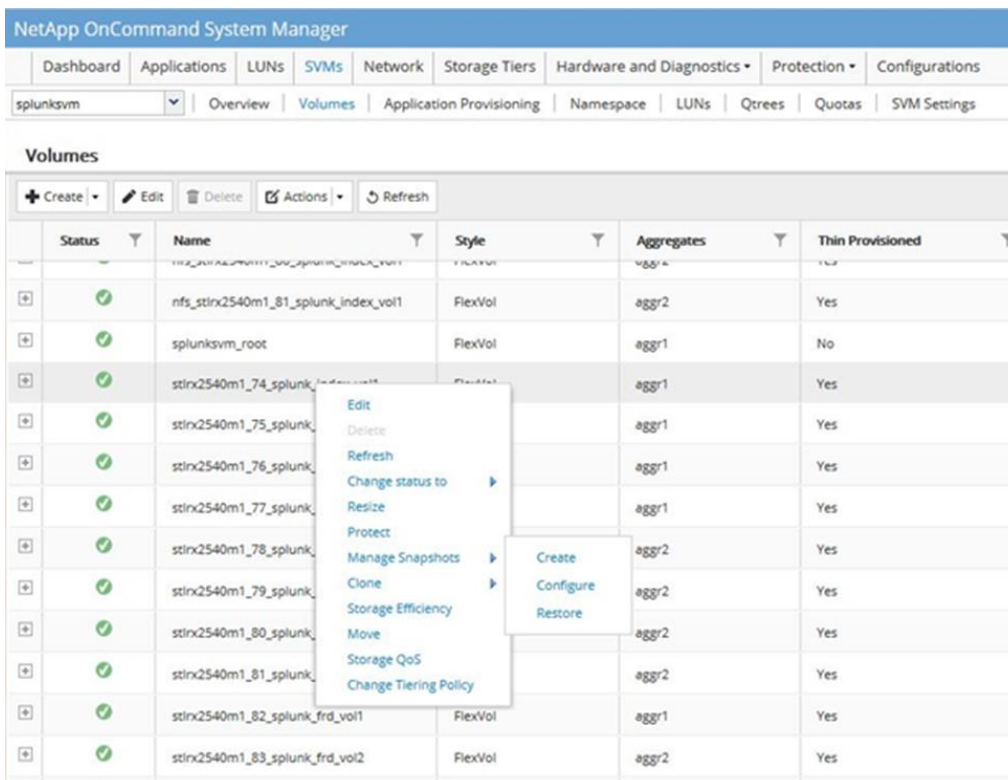
Test Result

The query failed with an error.

Restore Test

The following procedure was used to test the restore operation:

1. Stop all Splunk processes.
2. Unmount all data file systems on the indexers.
3. Restore all eight data volumes used by the indexers from previously created Snapshot copies by using SnapRestore, as shown by the following screenshots.



NetApp OnCommand System Manager

Dashboard Applications LUNs **SVMs** Network Storage Tiers Hardware and Diagnostics Protection Configurations

splunksvm Overview **Volumes** Application Provisioning Namespace LUNs Qtrees Quotas SVM Settings

Volumes

+ Create Edit Delete Actions Refresh

Status	Name	Style	Aggregates	Thin Provisioned
+	nfs_stlrx2540m1_81_splunk_index_vol1	FlexVol		
+	splunksvm_root	FlexVol		
+	stlrx2540m1_74_splunk_index_vol1	FlexVol		
+	stlrx2540m1_75_splunk_index_vol2	FlexVol		
+	stlrx2540m1_76_splunk_index_vol3	FlexVol		
+	stlrx2540m1_77_splunk_index_vol4	FlexVol		
+	stlrx2540m1_78_splunk_index_vol5	FlexVol		
+	stlrx2540m1_79_splunk_index_vol6	FlexVol		
+	stlrx2540m1_80_splunk_index_vol7	FlexVol		
+	stlrx2540m1_81_splunk_index_vol8	FlexVol		
+	stlrx2540m1_82_splunk_frd_vol1	FlexVol		
+	stlrx2540m1_83_splunk_frd_vol2	FlexVol		
+	stlrx2540m1_84_splunk_frd_vol3	FlexVol	aggr2	Yes

Restore Volume from Snapshot copy

Name	Created On	Application Dep...
backup2	Nov/06/2017 19:06:35	None
golden_ss	Nov/06/2017 18:32:06	None

Restore Cancel

NetApp OnCommand System Manager

Dashboard Applications LUNs **SVMs** Network Storage Tiers Hardware and Diagnostics Protection Configurations

splunksvm Overview **Volumes** Application Provisioning Namespace LUNs Qtrees Quotas SVM Settings

Volumes

+ Create Edit Delete Actions Refresh

Status	Name	Style	Aggregates	Thin Provisioned
+	nfs_stlrx2540m1_81_splunk_index_vol1	FlexVol	aggr2	Yes
+	splunksvm_root	FlexVol	aggr1	No
+	stlrx2540m1_74_splunk_index_vol1	FlexVol		
+	stlrx2540m1_75_splunk_index_vol2	FlexVol		
+	stlrx2540m1_76_splunk_index_vol3	FlexVol		
+	stlrx2540m1_77_splunk_index_vol4	FlexVol		
+	stlrx2540m1_78_splunk_index_vol5	FlexVol		
+	stlrx2540m1_79_splunk_index_vol6	FlexVol	aggr2	Yes
+	stlrx2540m1_80_splunk_index_vol7	FlexVol	aggr2	Yes

Restore Volume

⚠ Volume 'stlrx2540m1_74_splunk_index_vol1' will be restored using the Snapshot copy 'backup2'?

All changes made after this Snapshot copy was created will be lost.

☒ Restore volume from this Snapshot copy.

Restore Cancel

4. Mount the restored file systems on the indexers and restart Splunk.

The screenshot shows the Splunk Search & Reporting interface. The search bar contains the query: `index=ctrl 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code`. The results show 12 events. The table below represents the data shown in the results pane.

Time	Event
11/6/17 2:01:46.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder2 source = /splunk/splunk_ingest/log2ar-20171106-020146-274-000.log sourcetype = access_combined
9/5/17 4:37:40.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder1 source = /splunk/splunk_ingest/log1ar-20170905-043724-827-000.log sourcetype = access_combined
9/5/17 4:36:41.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder3 source = /splunk/splunk_ingest/log3ar-20170905-043635-954-000.log sourcetype = access_combined
9/5/17 4:34:23.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder2 source = /splunk/splunk_ingest/log2ar-20170905-043417-532-000.log sourcetype = access_combined
9/1/17 2:03:31.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder1 source = /splunk/splunk_ingest/log1ar-20170901-020327-559-002.log sourcetype = access_combined
9/1/17 2:03:27.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder1 source = /splunk/splunk_ingest/log1ar-20170901-020322-090-001.log sourcetype = access_combined
9/1/17 2:03:22.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder1 source = /splunk/splunk_ingest/log1ar-20170901-020303-804-000.log sourcetype = access_combined
9/1/17 1:57:54.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder3 source = /splunk/splunk_ingest/log3ar-20170901-015750-115-001.log sourcetype = access_combined
9/1/17 1:57:50.000 PM	<3.3> 2015-07-02 17:09:08,254323.197747-08:00 us-dur-s102-b4rtt49_subprod-01_vlan202.mgmt.netapp.com 10.102.254.36 %ACE-3-251010: Health probe failed for server 10.102.212.179 on port 11915, received invalid status code host = forwarder3 source = /splunk/splunk_ingest/log3ar-20170901-015745-695-000.log sourcetype = access_combined

Test Results

The same query was executed again and it produced identical results as before when all data was destroyed. The SnapRestore operation completed in a matter of seconds. The data was restored and fully available in less than five minutes following a catastrophic data loss.

7.4 Storage System Resilience Test (Controller Failure) – AFF A700

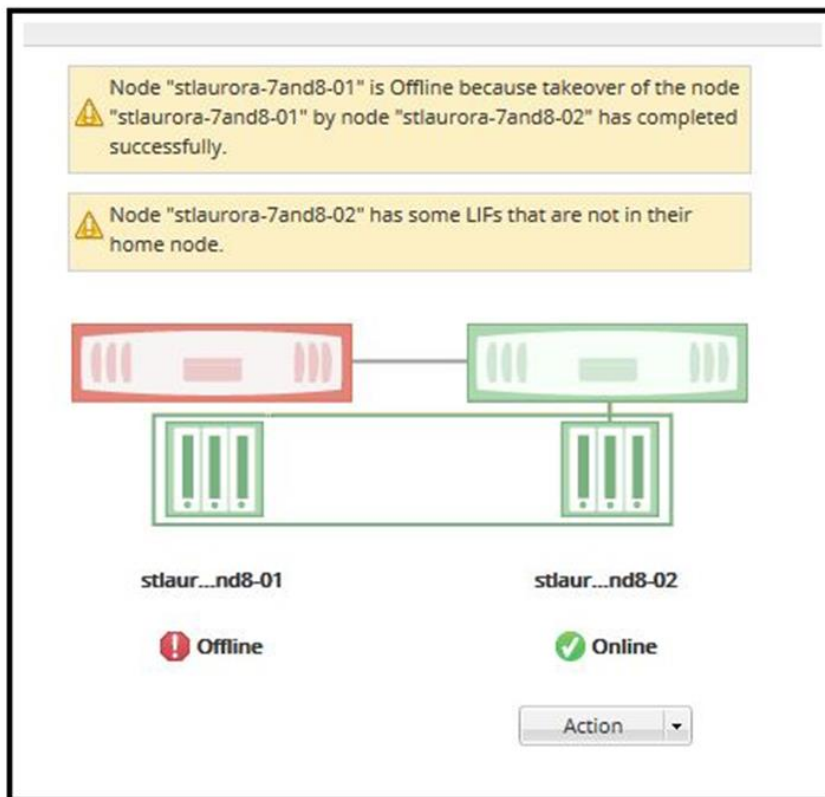
Test Details and Results

The following procedure was used for the storage system resilience test:

1. Apply a Splunk workflow to a healthy system.
2. Induce a controller failure to observe the response from Splunk.
3. Induce a panic on one of the storage controllers while a workflow is running.

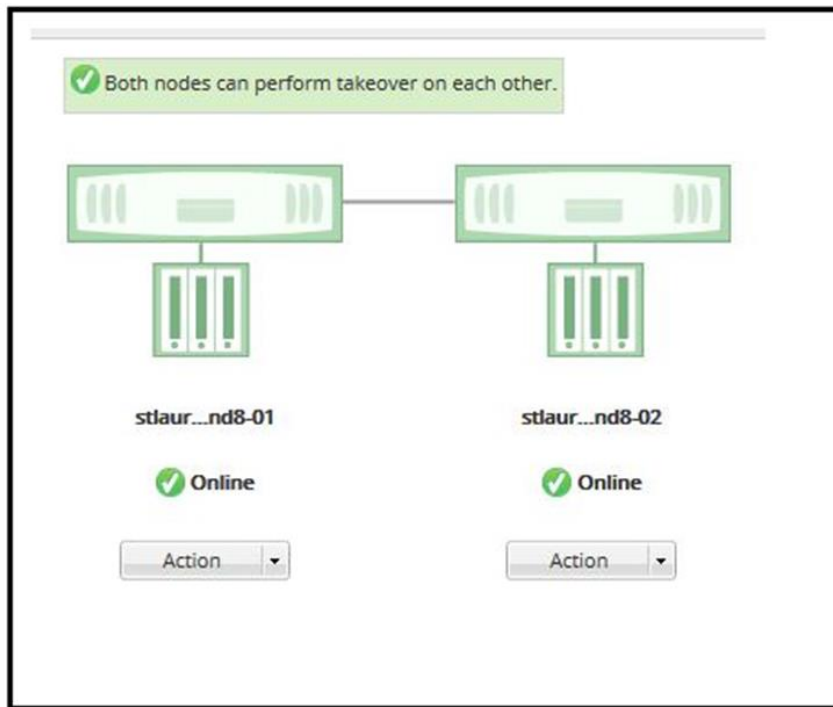
```
10.63.158.93 - PuTTY
login as: admin
Using keyboard-interactive authentication.
Password:
stlaurora-7and8::> run local
Type 'exit' or 'Ctrl-D' to return to the CLI
stlaurora-7and8-01> priv set diag
Warning: These diagnostic commands are for use by NetApp
        personnel only.
stlaurora-7and8-01*> panic
```

4. Observe the effect on the storage system from OnCommand System Manager.

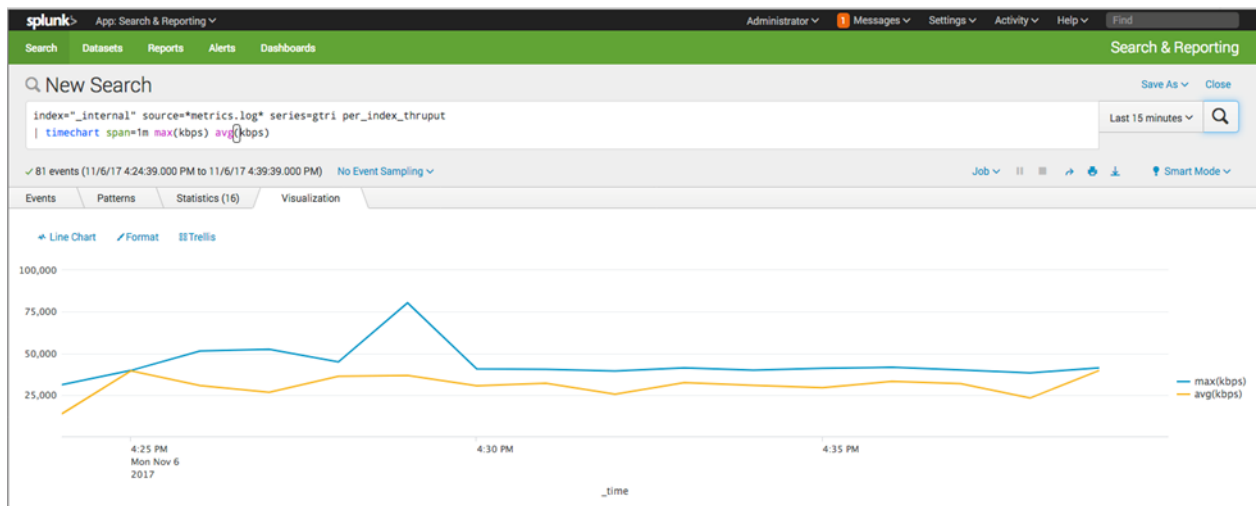


Test Results

At this point, the surviving controller took over for the failed controller and served data from the failed controller's disks. Several minutes later, the affected controller was recovered, a giveback was performed, and we observed that both controllers were up and healthy.



5. Monitor Splunk performance throughout the test.



Test Results

Panic was induced one minute after the workflow was started. Performance was monitored for one minute after giveback was completed. There was no visible effect on Splunk performance during that time, and the workflow completed successfully. Results may vary regarding the performance impact, depending on controller resource utilization.

7.5 Disk Failure and Reconstruct Test – AFF A700

Test Details and Results

The following procedure was used for the disk failure and reconstruct test:

1. Fail a disk during a Splunk index operation.
 - We chose disk 1.0.15 in System Manager for the test.

The screenshot shows the NetApp OnCommand System Manager interface. The top navigation bar includes links for Dashboard, Applications, LUNs, SVMs, Network, Storage Tiers, Hardware and Diagnostics (selected), Protection, and Configurations. The main content area is titled 'Disks' and has two tabs: 'Summary' and 'Inventory' (selected). Below the tabs are buttons for 'Assign', 'Zero Spares', and 'Refresh'. A table lists the disk inventory with columns: Name, Container Type, Home Owner, Current Owner, and Type. Disk 1.0.15 is highlighted in blue. Below the table, the 'Details' section for disk 1.0.15 is shown, listing its aggregate, vendor ID, zeroing status, serial number, and broken details.

Name	Container Type	Home Owner	Current Owner	Type
1.0.17	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.15	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.11	Shared	stlaurora-7and8-01	stlaurora-7and8-01	SSD
1.0.1	Shared	stlaurora-7and8-01	stlaurora-7and8-01	SSD
1.0.0	Shared	stlaurora-7and8-01	stlaurora-7and8-01	SSD
1.0.14	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.19	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.10	Shared	stlaurora-7and8-01	stlaurora-7and8-01	SSD
1.0.7	Shared	stlaurora-7and8-01	stlaurora-7and8-01	SSD
1.0.20	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.16	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD
1.0.21	Shared	stlaurora-7and8-02	stlaurora-7and8-02	SSD

Details	
Aggregate:	aggr0_stlaurora_7and8_02_0,aggr1,aggr2
Vendor ID:	NETAPP
Zeroing (%):	Zeroed
Serial No:	S396NA0H800957
Broken Details:	-NA-

2. Fail disk 1.0.15 in immediate mode without allowing any time for the contents to be copied to a replacement disk.

```

10.63.158.93 - PuTTY
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~> disk fail -disk 1.0.15 -immediate true

Warning: The system will not prefail the disk and its contents will not be
        copied to a replacement disk before being failed out. Do you want to
        fail out the disk immediately? {y|n}: y

stlaurora-7and8:~> █

```




3. Break status of disk 1.0.15.

NetApp OnCommand System Manager

DashboardApplicationsLUNsSVMsNetworkStorage TiersHardware and Diagnostics

Disks

SummaryInventory

 Assign  Zero Spares  Refresh

Name	Container Type	Home Owner
1.0.17	Shared	stlaurora-7and8-02
1.0.15	Broken	stlaurora-7and8-02
1.0.11	Shared	stlaurora-7and8-01
1.0.1	Shared	stlaurora-7and8-01
1.0.0	Shared	stlaurora-7and8-01

4. Manually unfail the disk after two minutes.

```

10.63.158.93 - PuTTY
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~>
stlaurora-7and8:~> disk unfail -disk 1.0.15

Warning: Failed disk "1.0.15" may have aggregate labels and file system data
present. In that case, this command will attempt to bring this disk
back into the aggregate with which this disk had formerly been
associated and preserve file system data. Are you sure you want to
continue with disk unfail? {y|n}: y

stlaurora-7and8:~> █

```

Test Results

The previously failed disk was recovered as a spare. Reconstruction had already begun on another spare disk.

NetApp OnCommand System Manager

DashboardApplicationsLUNsSVMsNetworkStorage TiersHardware and Diagnostics

Aggregates

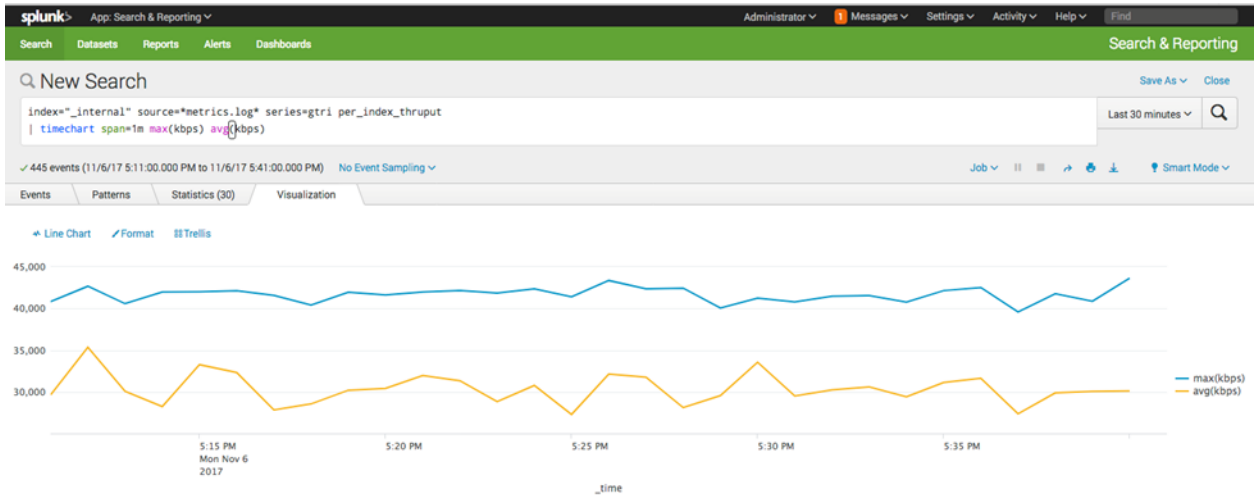
Aggregate: aggr1

OverviewDisk InformationVolumesPerformance

Name	Disk Status	Physical Size
1.0.6	Normal	894.25 GB
1.0.7	Normal	894.25 GB
1.0.8	Normal	894.25 GB
1.0.9	Normal	894.25 GB
1.0.12	Normal	894.25 GB
1.0.13	Normal	894.25 GB
1.0.14	Normal	894.25 GB
1.0.4	Reconstruction 43% completed	894.25 GB

Test Results

Reconstruction completed in approximately 20 minutes as the Splunk workflow continued to run without any noticeable impact. There was no visible impact on performance as the Splunk workflow completed.



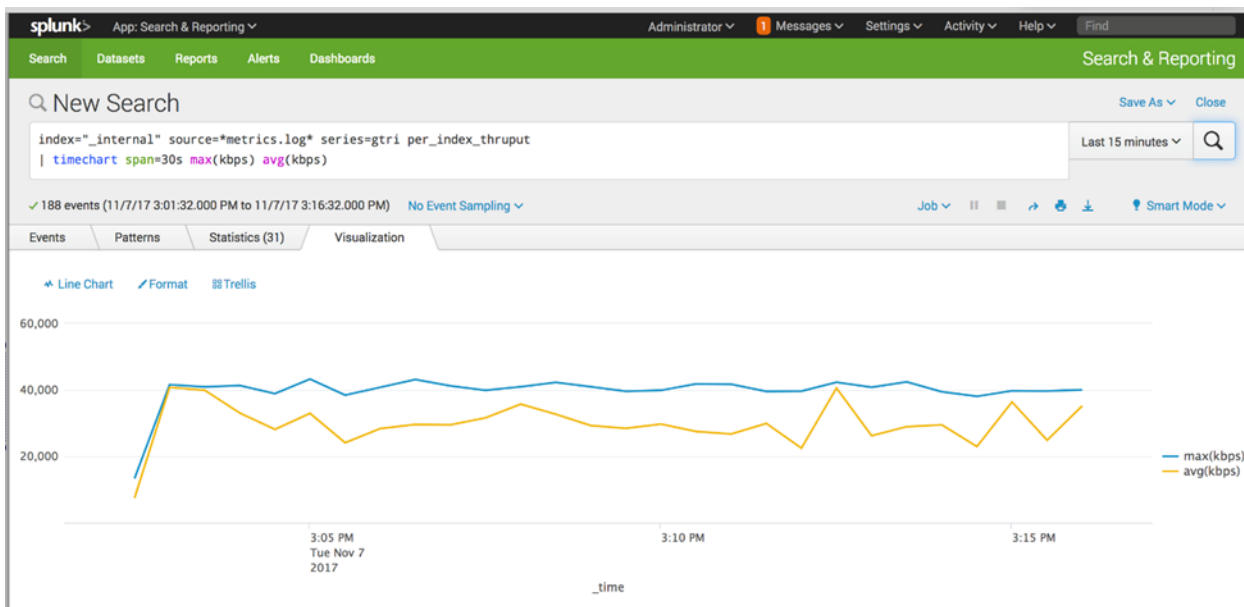
7.6 Splunk Indexer Failure Test – AFF A700

Test Details and Results

The purpose of this test was not to capture the effect of indexer failure on performance, but to demonstrate that NetApp storage fully supports indexer node failover. We determined the impact on performance by CPU and memory utilization on the servers, not by storage I/O bandwidth utilization. Because the storage in this configuration is shared at the controller and aggregate levels, and connectivity between the servers and storage is achieved by using 16Gbps FC, failure of an indexer would have no impact on storage performance.

The following procedure was used for the Splunk indexer failure test:

1. Manually fail four out of eight indexers, one by one, by using the Linux `kill` command to end the Splunk processes, while an indexing workflow is in progress.



Test Results

There was zero impact on data availability and performance, and the workflow was not interrupted as the indexers failed. Storage I/O bandwidth utilization remained the same, as did Splunk query performance. The use of NetApp AFF storage fully supported Splunk's indexer failover capability.

7.7 Creation of Data Copies with FlexClone Technology Test – AFF A700

Test Details and Results

The purpose of this test was to demonstrate the use of NetApp FlexClone technology to create free, instant copies of production data. We completed the following procedure:

1. Create the FlexClone volumes of the data volumes used by four of the indexers, as shown in the following screenshot.

NetApp OnCommand System Manager

Dashboard Applications LUNS SVMs Network Storage Tiers Hardware and Diagnostics Protection Configurations

splunksvm Overview Volumes Application Provisioning Namespace LUNS Qtrees Quotas SVM Settings

Volumes

+ Create Edit Delete Actions Refresh

Status	Name	Style	Total Space
✓	splunksvm_root	FlexVol	1 GB
✓	stlr2540m1_74_splunk_index_vol1	FlexVol	600 GB
✓	stlr2540m1_75_splunk_index_vol2	FlexVol	600 GB
✓	stlr2540m1_76_splunk_index_vol3	FlexVol	600 GB
✓	stlr2540m1_77_splunk_index_vol4	FlexVol	600 GB
✓	stlr2540m1_78_splunk_index_vol5	FlexVol	600 GB
✓	stlr2540m1_79_splunk_index_vol6	FlexVol	600 GB
✓	stlr2540m1_80_splunk_index_vol7	FlexVol	600 GB
✓	stlr2540m1_81_splunk_index_vol8	FlexVol	600 GB
✓	stlr2540m1_82_splunk_frd_vol1	FlexVol	600 GB
✓	stlr2540m1_83_splunk_frd_vol2	FlexVol	600 GB
✓	stlr2540m1_84_splunk_frd_vol3	FlexVol	600 GB
✓	stlr2540m1_85_splunk_cntl_mon_vol1	FlexVol	600 GB

Create FlexClone Volume

Name: stlr2540m1_74_splunk_index_vol1_clone_07112017_152639_99

☒ Thin Provisioning
Allocate space for the volume as it's used. Otherwise, the system reserves space for the entire volume.

FlexClone parent Snapshot copy

☒ Create new Snapshot copy now

☐ Use an existing Snapshot copy

Name	Date
------	------

Clone Cancel

2. Map the LUNs in the FlexClone volumes to four unused servers.
3. Mount the cloned file systems to the unused servers.
 - A list of the newly created FlexClone volumes displays.

NetApp OnCommand System Manager

Dashboard Applications LUNS SVMs Network Storage Tiers Hardware and Diagnostics Protection Configurations

splunksvm Overview Volumes Application Provisioning Namespace LUNS Qtrees Quotas SVM Settings

Volumes

+ Create Edit Delete Actions Refresh

Status	Name	Style	Aggregates	Thin Provisioned	Available Space	Total Space
✓	stlr2540m1_74_splunk_index_vol1_clone_07112017_152639_99	FlexVol	aggr1	Yes	76.64 GB	600 GB
✓	stlr2540m1_75_splunk_index_vol2_clone_07112017_153925_14	FlexVol	aggr1	Yes	70.7 GB	600 GB
✓	stlr2540m1_76_splunk_index_vol3_clone_07112017_153944_56	FlexVol	aggr1	Yes	71.27 GB	600 GB
✓	stlr2540m1_77_splunk_index_vol4_clone_07112017_154003_90	FlexVol	aggr1	Yes	71.89 GB	600 GB

Test Results

The actual physical footprint of the clones ranges from 8.09MB to 48.63MB, while the storage capacity of each clone is 600GB. Relatively speaking, these FlexClone volumes really are free data copies.

```

10.63.158.93 - PuTTY
stlaurora-7and8:>
stlaurora-7and8:>
stlaurora-7and8:>
stlaurora-7and8:> vol show -vserver splunksvm -volume stlrx2540m1_74_splunk_index_vol1_clone_07112017_152639_99 -fields size,physical-used
vserver    volume                                     size    physical-used
-----
splunksvm  stlrx2540m1_74_splunk_index_vol1_clone_07112017_152639_99 600GB 48.63MB

stlaurora-7and8:> vol show -vserver splunksvm -volume stlrx2540m1_75_splunk_index_vol2_clone_07112017_153925_14 -fields size,physical-used
vserver    volume                                     size    physical-used
-----
splunksvm  stlrx2540m1_75_splunk_index_vol2_clone_07112017_153925_14 600GB 30.34MB

stlaurora-7and8:> vol show -vserver splunksvm -volume stlrx2540m1_76_splunk_index_vol3_clone_07112017_153944_56 -fields size,physical-used
vserver    volume                                     size    physical-used
-----
splunksvm  stlrx2540m1_76_splunk_index_vol3_clone_07112017_153944_56 600GB 8.09MB

stlaurora-7and8:> vol show -vserver splunksvm -volume stlrx2540m1_77_splunk_index_vol4_clone_07112017_154003_90 -fields size,physical-used
vserver    volume                                     size    physical-used
-----
splunksvm  stlrx2540m1_77_splunk_index_vol4_clone_07112017_154003_90 600GB 21.48MB

stlaurora-7and8:>

```

- A list of the new LUNs along with the volume clones containing them is displayed.

NetApp OnCommand System Manager									
<div> Dashboard Applications LUNs SVMs Network Storage Tiers Hardware and Diagnostics Protection Configurations </div>									
<div> LUN Management Initiator Groups Portsets </div>									
<div> Create Clone Edit Delete Status Move Storage QoS Refresh </div>									
Name	SVM	Container Path	Space Reservation	Available Size	Total Size	% Used	Type	Status	
stlrx2540m1-74_lun1	splunksvm	/vol/stlrx2540m1_74_splunk_index_vol1	Enabled	339.35 GB	550.01 GB	38.3%	Linux	Offline	
stlrx2540m1-74_lun1	splunksvm	/vol/stlrx2540m1_74_splunk_index_vol1_clone_07112017_152639_99	Enabled	342.63 GB	550.01 GB	37.71%	Linux	Online	
stlrx2540m1-75_lun1	splunksvm	/vol/stlrx2540m1_75_splunk_index_vol2	Enabled	367.41 GB	550.01 GB	33.2%	Linux	Offline	
stlrx2540m1-75_lun1	splunksvm	/vol/stlrx2540m1_75_splunk_index_vol2_clone_07112017_153925_14	Enabled	367.41 GB	550.01 GB	33.2%	Linux	Online	
stlrx2540m1-76_lun1	splunksvm	/vol/stlrx2540m1_76_splunk_index_vol3	Enabled	382.99 GB	550.01 GB	30.37%	Linux	Offline	
stlrx2540m1-76_lun1	splunksvm	/vol/stlrx2540m1_76_splunk_index_vol3_clone_07112017_153944_56	Enabled	382.99 GB	550.01 GB	30.37%	Linux	Online	
stlrx2540m1-77_lun1	splunksvm	/vol/stlrx2540m1_77_splunk_index_vol4	Enabled	379.92 GB	550.01 GB	30.93%	Linux	Offline	
stlrx2540m1-77_lun1	splunksvm	/vol/stlrx2540m1_77_splunk_index_vol4_clone_07112017_154003_90	Enabled	379.92 GB	550.01 GB	30.93%	Linux	Online	
stlrx2540m1-78_lun1	splunksvm	/vol/stlrx2540m1_78_splunk_index_vol5	Enabled	397.28 GB	550.01 GB	27.77%	Linux	Online	
stlrx2540m1-79_lun1	splunksvm	/vol/stlrx2540m1_79_splunk_index_vol6	Enabled	386.83 GB	550.01 GB	29.67%	Linux	Online	

- After the new LUNs are brought online, map them to the initiator groups.

NetApp OnCommand System Manager

Dashboard
Applications
LUNs
SVMs
Network
Storage Tiers
Hardware and Diagnostics
Protection
Configurations

splunksvm
Overview
Volumes
Application Provisioning
Namespace
LUNs
Qtrees
Quotas
SVM Settings

LUN Management
Initiator Groups
Portsets

Create
Clone
Edit
Delete
Status
Move
Storage QoS
Refresh

Name	Container Path
stlrx2540m1-74_lun1	/vol/stlrx2540m1_74_splunk_index_vol1
stlrx2540m1-74_lun1	/vol/stlrx2540m1_74_splunk_index_vol1_clone_07112017_152639_99
stlrx2540m1-75_lun1	/vol/stlrx2540m1_75_splunk_index_vol2
stlrx2540m1-75_lun1	/vol/stlrx2540m1_75_splunk_index_vol2_clone_07112017_153925_14
stlrx2540m1-76_lun1	/vol/stlrx2540m1_76_splunk_index_vol3
stlrx2540m1-76_lun1	/vol/stlrx2540m1_76_splunk_index_vol3_clone_07112017_153944_56
stlrx2540m1-77_lun1	/vol/stlrx2540m1_77_splunk_index_vol4
stlrx2540m1-77_lun1	/vol/stlrx2540m1_77_splunk_index_vol4_clone_07112017_154003_90
stlrx2540m1-78_lun1	/vol/stlrx2540m1_78_splunk_index_vol5

LUN Properties

Name: stlrx2540m1-74_lun1

Policy Group: None

Container Path: /vol/stlrx2540m1_74_sp...

Minimum Throughput: NA

Size: 550.01 GB

Maximum Throughput: NA

Status: Online

Move Job Status: NA

Type: Linux

LUN Clone: false

Move Last Failure: NA

Edit LUN

General

Initiator Groups

Map

Initiator Group Name	Type	LUN ID (Optional)
<input checked="" type="checkbox"/> stlrx2540m1-74.stl.n...	Linux	0
<input type="checkbox"/> stlrx2540m1-86.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-85.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-84.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-83.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-82.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-81.stl.n...	Linux	
<input type="checkbox"/> stlrx2540m1-80.stl.n...	Linux	

☐ Show All Initiator Groups

Add Initiator Group

Save

Save and Close

Cancel

Note: Each volume clone contains a single LUN, and each LUN contains a copy of the logical volume and file system from the parent volume.

We completed the following procedure on each new indexer server to mount the file systems that exist on the cloned LUNs:

1. Run the `rescan-scsi-bus.sh -a` command to make the cloned LUNs visible to the indexer server.
2. Run the `vgscan` command to discover the volume group on the LUN clone.
3. Run the `lvscan` command to discover the logical volume on the LUN clone volume group.
4. Create a mount point for the new file system.
 - For example, execute the `mkdir /splunk` command.
5. Add the new file system to the Linux `/etc/fstab` file on the new indexer server and then mount the cloned file system by running the `mount /splunk` command.

```
#
# /etc/fstab
# Created by anaconda on Tue Aug 22 14:13:01 2017
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/rhel_stlrx2540m1--74-root / xfs defaults 0 0
UUID=b9501a5c-a347-404b-8fcd-28fe95388d40 /boot xfs defaults 0 0
/dev/mapper/rhel_stlrx2540m1--74-home /home xfs defaults 0 0
/dev/mapper/rhel_stlrx2540m1--74-swap swap swap defaults 0 0
/dev/splunk_vg_1/splunkindex /splunk xfs allocsize=128m,noatime,nobarrier,nodiratime 0 0
```

6. Run the Linux `find` command to confirm that the data exists in the cloned file system.

```
[root@stlrx2540m1-74 ~]# find /splunk -maxdepth 3 -exec du -sm {} \;
152909 /splunk
152909 /splunk/splunktest
152909 /splunk/splunktest/gtri
152909 /splunk/splunktest/gtri/db
0 /splunk/splunktest/gtri/colddb
1 /splunk/splunktest/defaultdb
1 /splunk/splunktest/defaultdb/db
0 /splunk/splunktest/defaultdb/colddb
[root@stlrx2540m1-74 ~]#
```

7. Configure the Splunk cluster to use the four new indexer servers.

Test Results

We started the Splunk processes, and the cluster was ready for workflow execution.

7.8 Index Performance Testing

Test Details

For this test, we used the custom script mentioned in section 4.1, to generate approximately 1TB of simulated log data. That data was spread across the three universal forwarders and then streamed to the eight indexer peer nodes, where it was indexed and stored on the AFF A800. We captured both the peak

and average indexing rates by using the custom Splunk dashboard. The following Splunk configurations were tested:

1. Internal DAS (SSDs and HDDs) for hot/warm and cold buckets (for comparison).
2. AFF A800 with iSCSI for hot/warm buckets.
3. AFF A800 with 32Gbps FC SAN for hot/warm buckets.
4. AFF A800 with NVMe/FC (NVMe over 32Gbps FC) for hot/warm buckets.

Configurations 2 through 4 used a FAS2552 with spinning disks over 10GbE and NFS for cold buckets. Details of these configurations are provided in section 7.

Test Results

Table 10 shows the test results from the AFF A800 configuration, along with the test results from an identical Splunk configuration using DAS. For both the average and peak indexing rates, the AFF A800 configuration outperformed the commodity servers with DAS for all protocols tested.

Table 10) Splunk indexing rate performance comparison for 1TB data ingest.

Average Indexing Rate for 1TB (AFF A800 versus Internal DAS Configuration)		
Configuration Summary	Avg Rate Observed versus DAS	Peak Rate Observed versus DAS
FC SAN with 8 indexer peer nodes	5% <i>faster</i> than DAS	7% <i>faster</i> than DAS
iSCSI with 8 indexer peer nodes	5% <i>faster</i> than DAS	7% <i>faster</i> than DAS
NVMe/FC (32Gbps FC)	6% <i>faster</i> than DAS	7% <i>faster</i> than DAS

Figure 6 and Figure 7 are graphical views of the data in Table 10, emphasizing the indexing performance difference between the AFF A800 and DAS configurations. For each data ingest test, we captured the following data points at 5-minute intervals in a Splunk dashboard:

- Average index rate across all eight indexer peer nodes.
- Peak indexer peer node index rate

Like the other tests, we used the custom script described in section 4.1, to generate 1TB of log data for the index workflow I/O, spread across all 8 peer nodes in the indexer cluster. The indexer cluster master does a good job of directing log data input across the peer nodes in a balanced manner, but not perfectly. Each universal forwarder in our deployment has been configured to forward a different type of log to the indexer cluster. Depending on which type of log file is being indexed on each peer node, some indexer peers are busier than others. As a result, there can be quite a spread between the peak and average ingest rates for the indexer peer nodes. This is normal for Splunk. To get a complete picture of indexer performance, it's helpful to monitor both the average and peak indexing rates. During our tests, it was observed that the peak indexing rates were often roughly twice the average rates. Figure 9 shows sample dashboard output from one of our tests, demonstrating the difference between average and peak index performance.

Figure 6) Average indexing rate comparison: AFF A800 versus DAS.

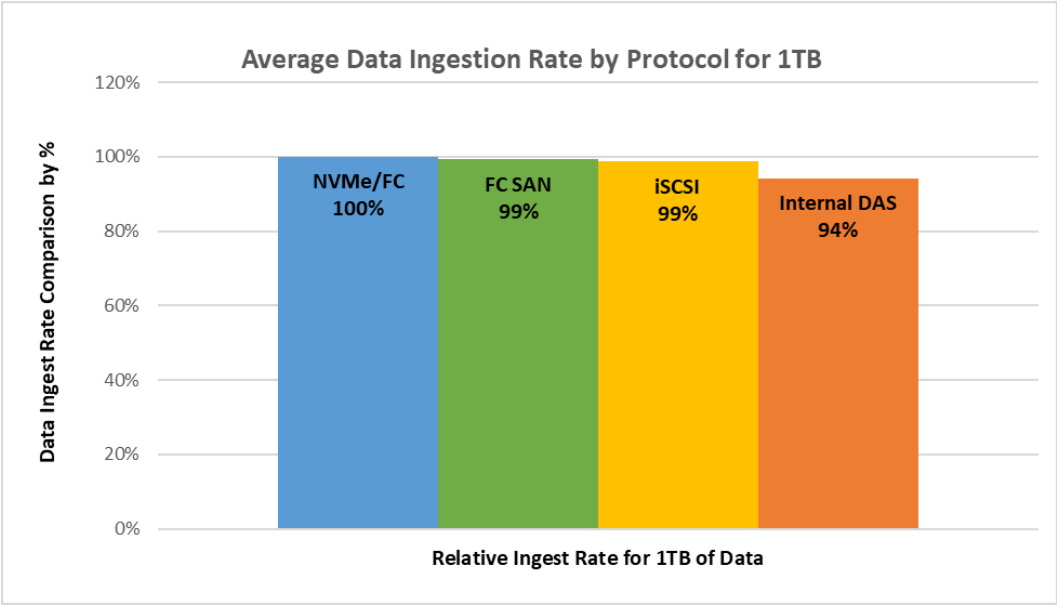


Figure 7) Peak indexing rate comparison: AFF A800 versus DAS.

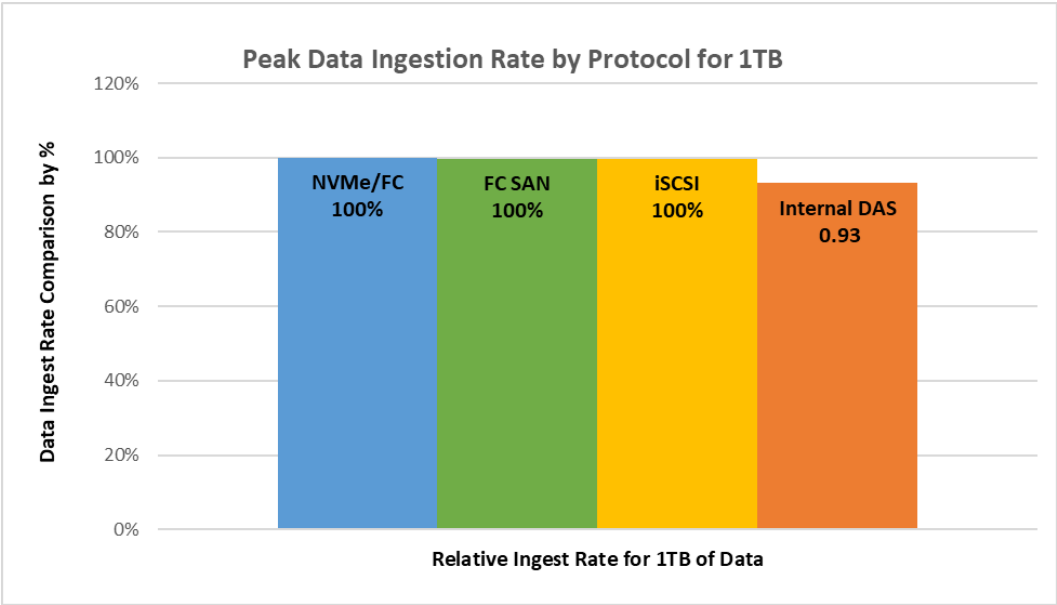
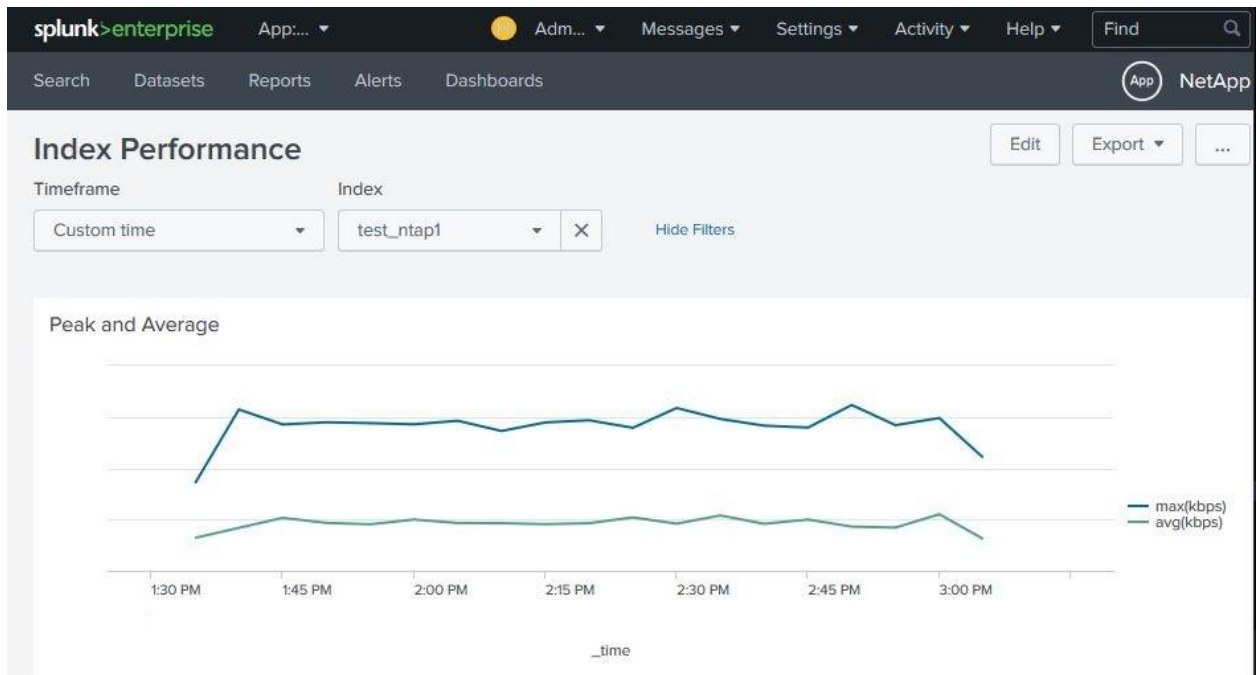


Figure 8) Graphical example of peak and average index performance.



With internal DAS Splunk deployments, the user must add servers and storage whenever additional Splunk resources are needed. If additional storage is needed, both servers and storage must be added. If additional compute resources are needed, both servers and storage must be added. Consequently, there may be a perpetual imbalance in resources, either unused storage or unused compute resources. With the NetApp AFF solution for Splunk, we decouple storage from compute. As the need for Splunk resources grows, the administrator only needs to address the resource deficit. In other words, if additional compute is required, only servers need be added. If additional storage is needed, storage can be added without adding servers. Compute and storage can scale independently. With that in mind, we monitored indexer host CPU utilization during all ingest tests. Average CPU idle time across the eight indexer peer nodes for each configuration tested is shown in Table 11. Iostat was used to capture that data with a sampling rate of one per minute.

Table 11) Average indexer CPU utilization during data ingest.

Average Idle Time During Data Ingest Tests per Tested Configuration		
Configuration Summary	Average % CPU Idle Time	Minimum % CPU idle Time
FC SAN with 8 indexer peer nodes	93%	77%
iSCSI with 8 indexer peer nodes	93%	79%
NVMe/FC (32Gbps FC)	92%	78%
Internal DAS	93%	77%

As listed in Table 12, the average CPU idle time for data ingest, regardless of storage configuration was around 93%, with minimum CPU idle time between 77% and 79%. In this case, with the internal DAS configuration, adding indexer storage would require the addition of unneeded compute resources. With the other three configurations, where storage has been decoupled from compute, storage can be added without adding servers. In fact, with the FAS AFF configurations that were tested, it's possible to reduce the server count and still maintain good performance.

7.9 Performance Baseline Test: Splunk Search Performance

Test Details

Next, search performance tests were carried out using same indexes that were created during our index performance tests. The following types of searches were tested:

- **Dense search.** A query that scans data and returns matches for at least 1% of events scanned. The dense search used in these tests returned 1 match per 100 lines of log data scanned.

Note: Examples of a dense search include searches that return the number of errors on a web server or all failed login events on a database server.

- **Sparse search.** A query that matches between 0.01% and 1% of events scanned. The sparse search used in these tests returned 1 match per 1,000 lines.

Note: Sparse searches are often referred to as “needle in a haystack” searches.

- **Super-sparse search.** A query that matches between .0001% and 0.01% of events scanned. The super-sparse search used in these tests returned 1 match out of 1,000,000 lines scanned.
- **Rare search.** A query that matches between .00001% and .0001% of events scanned. The rare search used in these tests returned 1 match per 10,000,000 lines scanned.

Each type of search was performed with indexers actively indexing data and with indexers in a static state (with no data ingest).

Search Performance Test Results

Table 12 compares the test results for these queries and for queries performed with a similarly configured DAS Splunk environment. The key metric of these tests is completion time in seconds.

Table 12) Splunk search performance: AFF A700 versus DAS.

Search Time Comparison – AFF A800 versus Internal DAS								
	With Ingest				With No Ingest			
	Dense	Sparse	Super Sparse	Rare	Dense	Sparse	Super Sparse	Rare
FC SAN	18% faster	17% faster	Same as DAS	Same as DAS	17% faster	17% faster	Same as DAS	Same as DAS
NVMe/FC	14% faster	17% faster	Same as DAS	Same as DAS	12% faster	16% faster	Same as DAS	Same as DAS
iSCSI	12% faster	17% faster	Same as DAS	Same as DAS	10% faster	17% faster	Same as DAS	Same as DAS

The results clearly show that the NetApp AFF A800 outperformed DAS for dense and sparse queries, regardless of whether data was being ingested or not. Those queries returned the greatest amount of data. The super sparse and rare searches returned the least amount of data and completed in about the same amount of time as DAS.

For the FC SAN, NVMe/FC, and iSCSI AFF A800 configurations, search performance significantly exceeded that of internal DAS for both dense and sparse searches conducted both with and without data ingest running. The difference ranged from 10% to 18%. Performance for super sparse and rare searches were the same for all configurations tested.

Figures 9 through 12 are graphical representations of these results. The AFF A800 results are represented as a percentage, as compared to DAS, where DAS was set to 100%. For example, in Figure 9, the dense search during data ingest, using NVMe/FC storage, finished in 86% the amount of time required for DAS.

Figure 9) Dense search during data ingest completion time comparison.

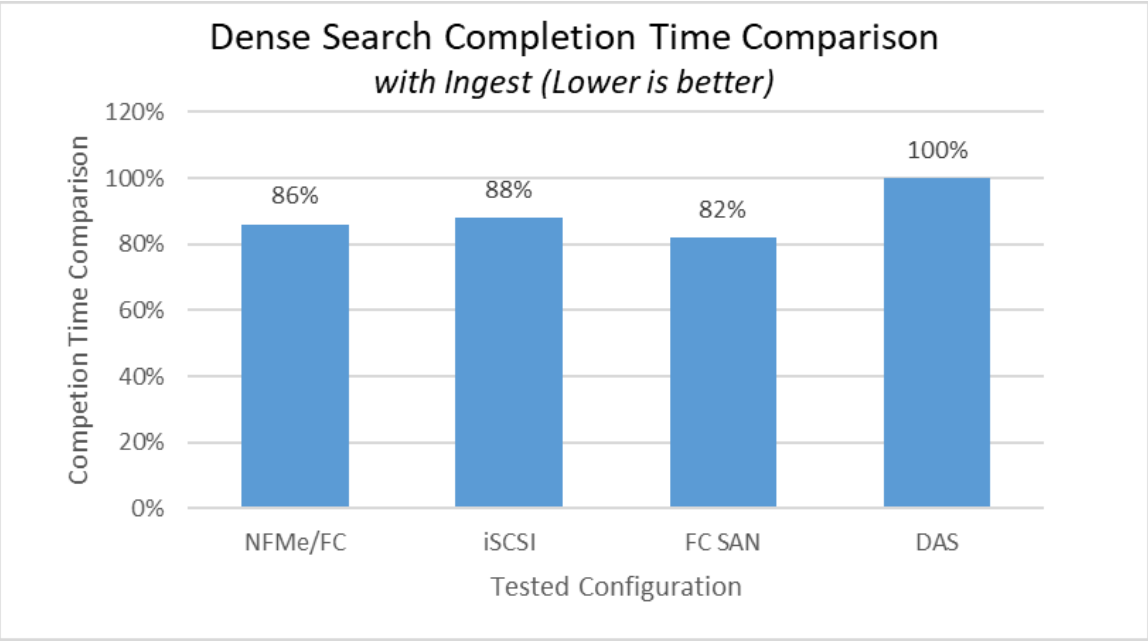


Figure 10) Dense search without data ingest completion time comparison.

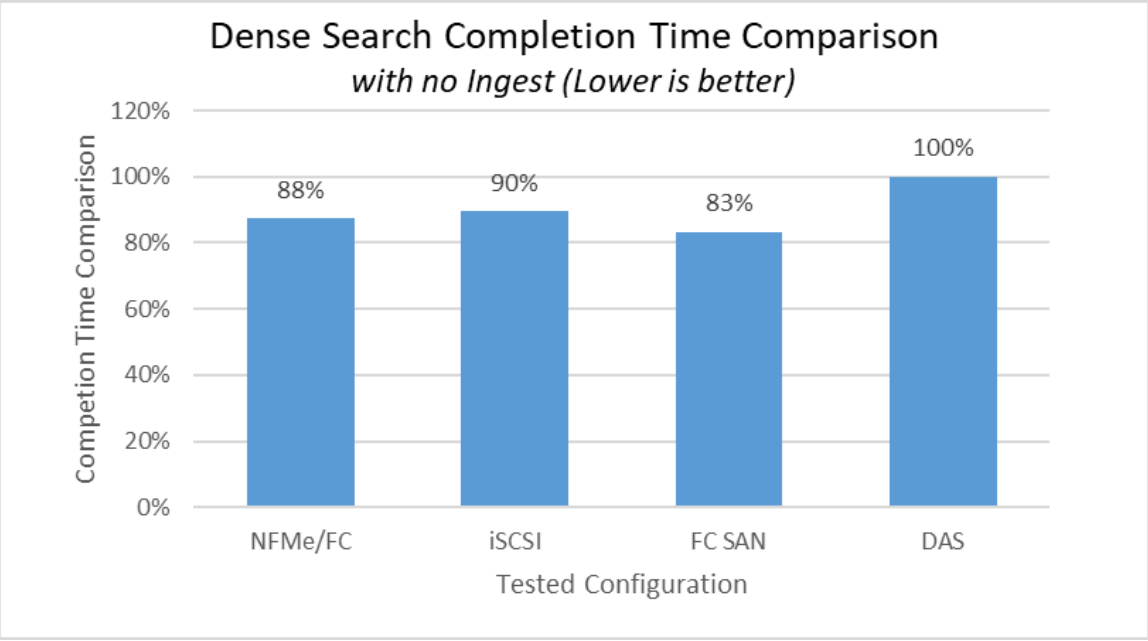


Figure 11) Sparse search during data ingest completion time comparison.

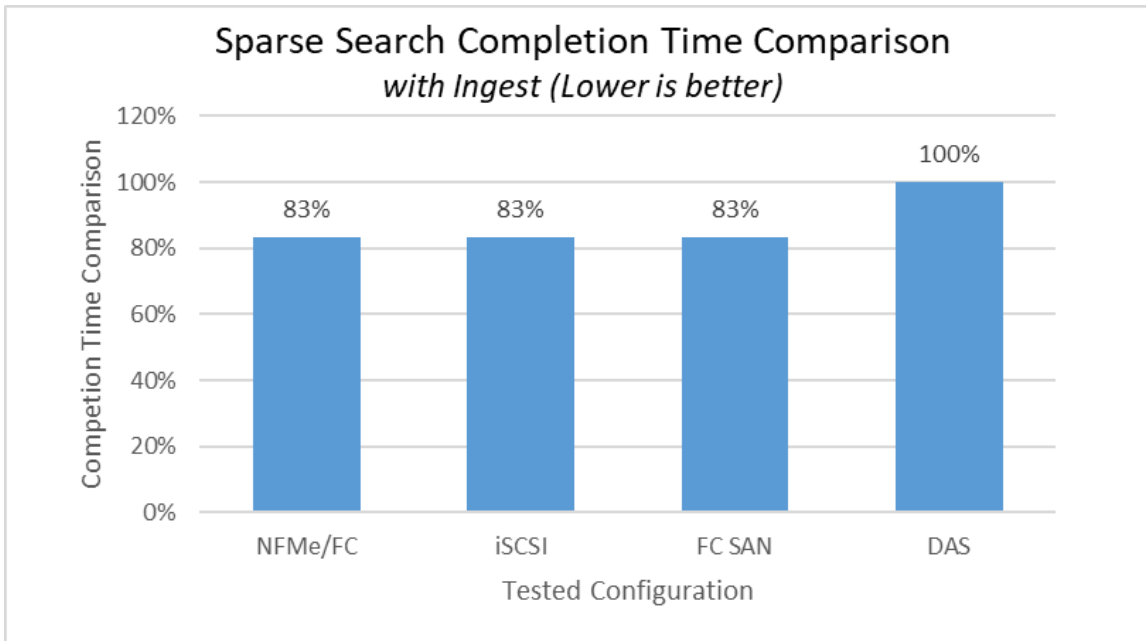
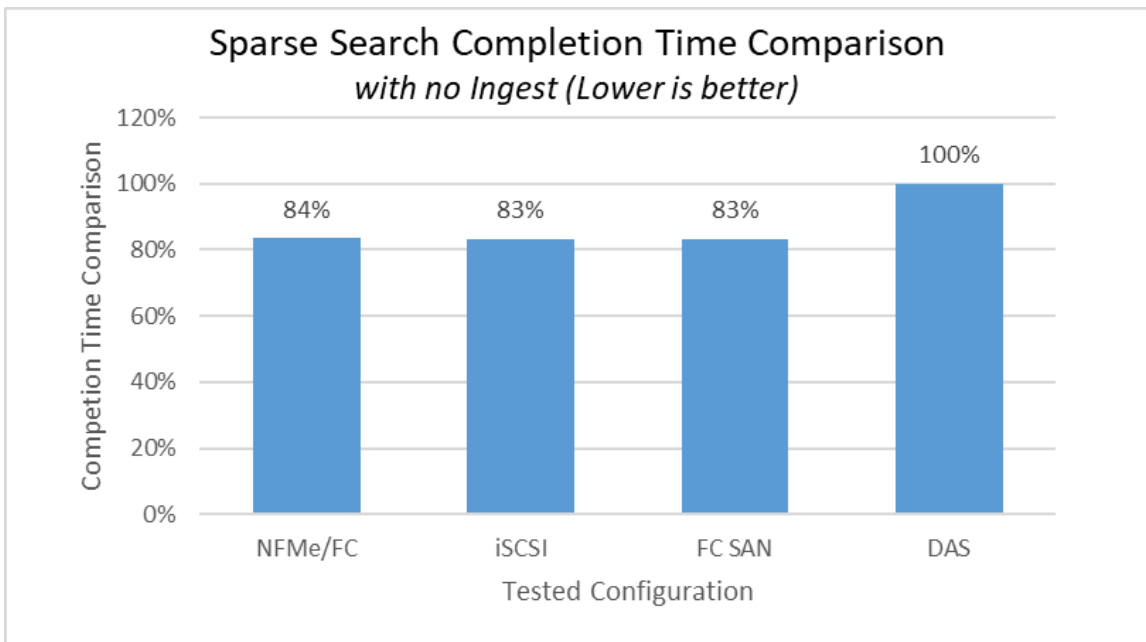


Figure 12) Sparse search without data ingest completion time comparison.



7.10 ONTAP Storage Efficiency Test

Test Details

After completing each index performance benchmark test, we queried storage efficiency by using OnCommand System Manager and command line to determine the following:

1. Overall storage space savings due to ONTAP inline storage efficiency.

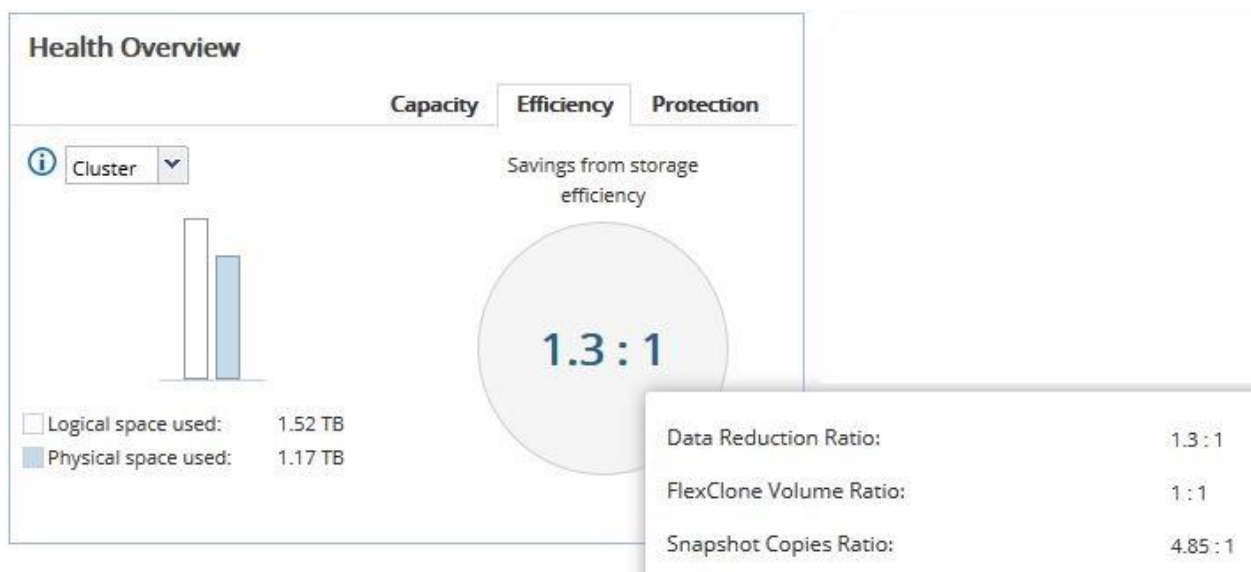
2. Storage efficiency details, to determine the percent contribution of each storage efficiency feature to the overall savings.

This data was captured for data only, before any Snapshot copies or FlexClone volumes were created. We then created a single snapshot for each storage volume and recorded overall storage efficiency. After that, we added a single FlexClone for each volume and recorded overall storage efficiency.

Test Results

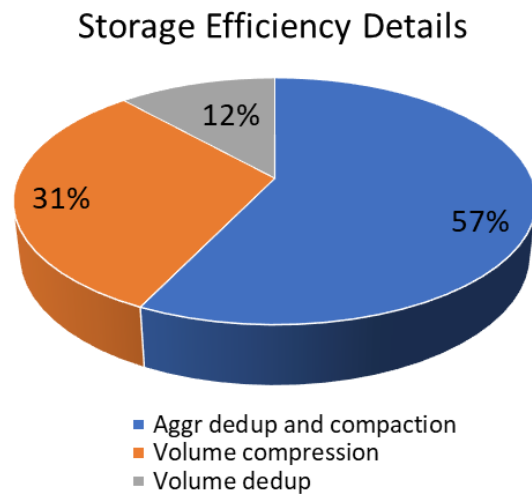
As shown in Figure 13, the overall storage efficiency for the aggregates used in our Splunk configuration was reported to be 1.3:1 for 1.52TB of Splunk indexer data (without Snapshot copies or FlexClone volumes). The storage savings is the result of inline deduplication, compaction, and compression. Figure 13 shows the ONTAP storage efficiency ratio as reported graphically by OnCommand System Manager. The ratio represents the comparison of the logical and the physical space used.

Figure 13) Savings from storage efficiency as shown in OnCommand System Manager (only data).



From the ONTAP command line, we obtained additional details about the total savings. These details are presented in Figure 14.

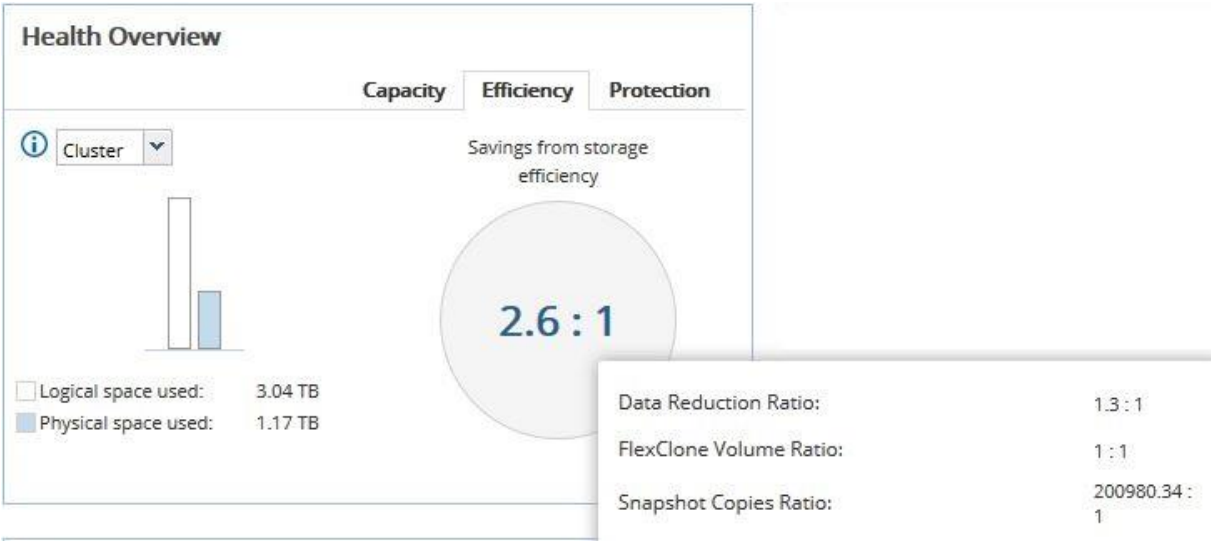
Figure 14) ONTAP savings from storage efficiency details.



The greatest percentage (57%) of our storage savings was obtained from aggregate deduplication and compaction. 32% of savings was obtained from volume compression and 12% was obtained from volume deduplication.

After creating a Snapshot copy on all data volumes, the storage efficiency ratio doubled, going from 1.3 to 2.6, as shown in Figure 15.

Figure 15) Savings as shown in OnCommand System Manager (data + 1 Snapshot copy).



Finally, a FlexClone volume was created from each volume, causing the efficiency ratio to increase by nearly 50%, nearly triple the storage efficiency ratio observed for data only. Figure 16 shows storage efficiency as graphically reported by OnCommand System Manager.

Figure 16) Savings as shown in OnCommand System Manager (data + 1 Snapshot copy + 1 FlexClone).

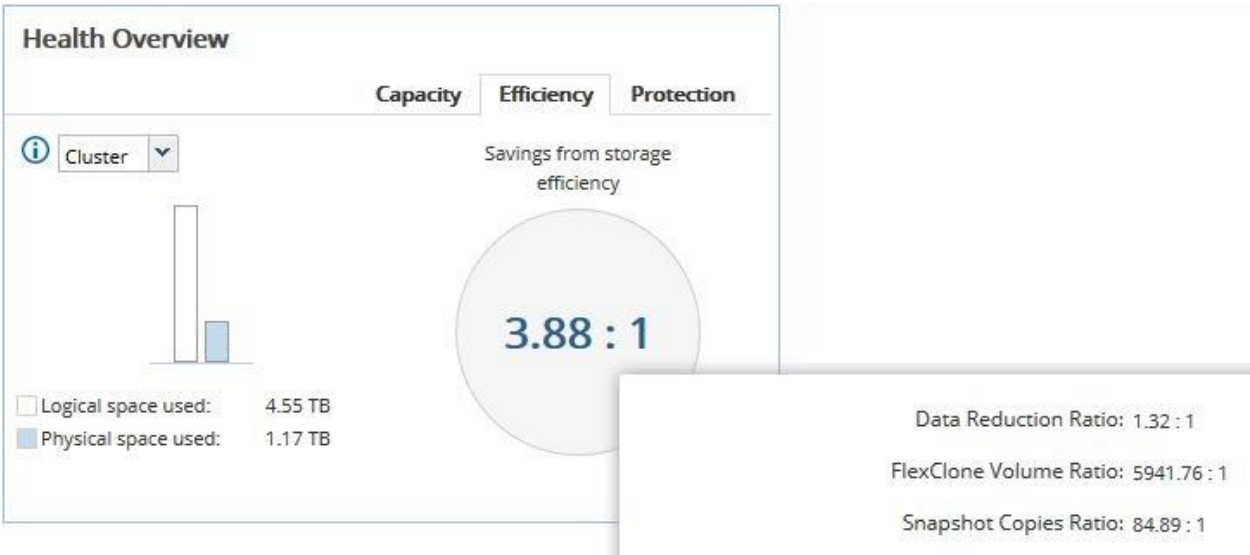


Table 13 summarizes the storage efficiency observations described previously.

Table 13) Storage efficiency for Splunk Indexer data summary.

Savings from Storage Efficiency Summary		
Description	Storage Efficiency Ratio	Savings Details
Data only	1.3:1	
Volume deduplication		12%
Volume compression		31%
Aggregate-level deduplication and compaction		57%
Data + 1 Snapshot copy per volume	2.6:1	
Data + 1 Snapshot copy + 1 FlexClone volume	3.88:1	

For data only, savings were obtained from a combination of volume deduplication, volume compression, aggregate-level deduplication, and aggregate-level compaction.

- 6% of storage savings came from ONTAP volume deduplication
- 36% of storage savings came from ONTAP volume compression
- 57% of storage saved came from ONTAP aggregate level deduplication and compaction

Inline storage efficiency is enabled automatically for NetApp AFF storage arrays, and has negligible, if any, effect on performance.

The multiplying effect on storage efficiency by the use ONTAP Snapshot copies and FlexClone volume should be noted. Creating one Snapshot copy backup of each storage volume doubles the storage efficiency achieved for data alone. Creating one Snapshot copy and one FlexClone volume of each storage volume nearly triples the storage efficiency ratio.

The logical size of the Splunk data was about 1.52TB, but the actual physical space used was only about 1.17TB. The space savings is the result of ONTAP inline storage efficiency. Although results can vary based on actual data, the results obtained were typical across all the AFF A800 configurations tested.

Best practices for Splunk using DAS require storage media to be configured as RAID 10, which involves mirroring. In this configuration, Splunk has access to only half of a server's internal storage. With NetApp AFF storage, our data required 1.17TB of physical storage. The same amount of data would require 3.04TB of server based internal RAID 10 storage.

8 Recommendations

8.1 Splunk Settings

Considerable effort went into tuning Splunk at the application level for best performance with NetApp AFF storage. As it turns out, Splunk is well written set of applications, with little need for manual optimization. Other than Splunk's own best practices, there is only one tuning recommendation for Splunk running on NetApp ONTAP storage. This recommendation is to set the `server.conf` parameter `parallelIngestionPipelines` to 2. The default setting is 1. By increasing the value to 2, you can moderately improve the ingesting performance. The `parallelIngestionPipelines` parameter is covered in Splunk documentation, and Splunk best practices should be followed while making the change. This parameter provides parallelization for the indexer and accelerates indexer data parsing and disk writes. A setting of 2 enables a small, but noticeable improvement in the indexing speed. The maximum recommended value is 2.

8.2 Network Configuration

For optimum performance, the Splunk server network should be at least 10GbE with end-to-end jumbo frames enabled.

For IP-based storage (iSCSI and NFS), the indexer servers must have a unique 10GbE private storage network with end-to-end jumbo frames enabled. That network should be used only for storage. Not doing so, will result in instability of the Splunk deployment, as network congestion results in delayed heartbeats and slow replication.

8.3 NFS for Splunk Storage

NFS is a good choice for cold bucket storage but should not be used for hot and warm buckets. NFS was tested in our environment and found to perform at about 50% of internal DAS for data ingest. The use of NFS for hot/warm buckets is also strongly discouraged by Splunk. Given our experience with NFS performance and Splunk's messaging, NetApp recommends against using NFS for hot/warm bucket storage.

8.4 Server-Side Configuration

For details about the server and OS configuration optimization for Splunk, see appendix B.

Appendix A: Splunk Use Cases

Common use cases for Splunk include:

- Analyze system performance: log monitoring and reporting with trend analysis
- Troubleshoot failure condition: root cause analysis using system logs
- Monitor business metrics: analysis of output from business applications and real-time data
- Search and investigate an outcome

- Create dashboards to visualize and analyze results
- Provide security monitoring and assurance

At a time when large-scale data breaches and cyberthreats often make headlines, security is a primary concern for everyone, from the largest online retailers to anyone who has a bank account. Splunk offers a powerful framework for fraud prevention and security. Use cases include:

- Malware detection and investigation:
 - Detect infected hosts
 - Determine the spread of malware
- Data exfiltration—unauthorized transfer of data over a network:
 - Monitor transactions
 - Isolate suspicious events
- Privileged user monitoring:
 - Prevent or contain advanced attacks
 - Prevent or contain insider threat-based attacks
- Identification of patient zero malware:
 - Identify command-and-control network communications
 - Identify malware-infected hosts
 - Identify first infected host (patient zero) of a malware outbreak
- Detect zero-day attacks—exploitation of unknown software security vulnerability:
 - Splunk Enterprise Security Risk Analysis Framework
 - Security Domain dashboards
- Fraud detection—account takeovers:
 - Detect and investigate
 - Remediation actions initiated by Splunk
- Ensure compliance—detect when critical systems stop sending logs to Splunk:
 - Avoid regulatory compliance issues
 - Investigation and remediation actions
- User Behavior Analytics (UBA) for insider threats:
 - Detect and mitigate insider threats
 - Fully automated and continuous monitoring
- UBA for external threats:
 - Detect cyberattacks of malware and hidden threats
 - Stop external threats before it's too late

Of all the Splunk use cases, security is the most critical for any organization or individual. It's so critical that Splunk offers a product that is dedicated to security. For Splunk to provide the high level of security required today, the product must be configured for continuous uptime and optimal performance. NetApp AFF storage arrays powered by ONTAP provide the features and tools to meet those requirements.

Appendix B: Test Configuration Details

The following three Splunk configurations are referenced in this document:

- Splunk with NetApp AFF A700 FC SAN for indexer storage
- Splunk with NetApp AFF A800 FC SAN for indexer storage

- Splunk with NetApp AFF A800 iSCSI SAN for indexer storage
- Splunk with NetApp AFF A800 NVMe/FC for indexer storage
- Splunk with DAS internal to the servers

For the server configuration details, see the appendix.

Appendix C: Server Hardware

Servers Used with AFF A700 FC SAN Configuration

The following servers were used with the AFF A700 FC SAN configuration:

- Fujitsu PRIMERGY RX2540 M1 servers, each equipped with:
 - 2 CPUs, 16 physical cores total
 - Intel Xeon CPU E5-2670 v3 @ 2.30GHz processors
 - 256GB physical memory
 - Three 300GB SAS OS drives (mirrored)
 - One QLogic QLE2672 QLogic 2-port 16Gb FC Adapter

Figure 2 illustrated this configuration.

Servers Used with AFF A800 and Internal DAS Configuration

The following servers were used with the internal DAS configuration:

Fujitsu PRIMERGY RX2540 M4 servers running SLES 12.3, each equipped with:

- 2 CPUs, 32 physical cores total
- Intel Xeon Gold 6142 CPU @ 2.60GHz processor
- 256GB physical memory
- Three 300GB SAS OS boot drives (mirrored)
- 32Gbps FC, dual controller
- 10GbE dual port

Eight of those servers were used and indexer cluster peer nodes, each equipped with:

- Four Fujitsu 960GB 12Gbps SAS SSDs for hot/warm buckets
- Two 300GB 12Gbps SAS 15K RPM HDDs for cold buckets

The SSDs and HDDs were configured as per Figure 5

Appendix D: Server OS Details – AFF A800 Configuration

The following server OS versions were used:

- SLES 12.3 for performance tests
- RHEL 7.2 for functional and resilience tests

OS Kernel Settings

The system requirements recommendations for the use of Splunk Enterprise on-premises listed in the [Splunk Enterprise installation manual](#), were followed.

Logical Volumes and File System Configuration, Mount Options

Logical Volumes Configuration

The logical volumes were configured as follows:

- For the AFF A800 environments that use FC SAN and iSCSI, the default option with four stripes (-i 4), corresponding to the number of LUNs per volume group, was used.

For example:

```
lvcreate -W y -i 4 -L 620G -n splunk_lv splunk_vg
```

- For the AFF A700 environments that use FC SAN, the default option with one stripe (-i 1), corresponding to the number of LUNs per volume group, was used.

For example:

```
lvcreate -W y -i 1 -L 550G -n splunk_lv splunk_vg
```

File System Configuration

The XFS file system was configured with the default options.

For example:

```
mkfs.xfs -f -L SPLUNK_FCP /dev/splunk_vg/splunk_lv
```

Mount Options

The following mount options were used:

- Mount options for XFS on NVMe/FC (for hot/warm buckets):
`defaults,nobarrier`
- Mount options for XFS on FC SAN (for hot/warm buckets):
`defaults,nobarrier`
- Mount options for XFS on iSCSI (for hot/warm buckets):
`defaults,nobarrier`
- Mount options for XFS on internal DAS (for hot/warm buckets):
`Defaults`
- NFS mount options for FAS2552 volumes (for cold buckets)
`rw,bg,hard,nointr,nolock,rsiz=32768,wsiz=32768,tcp`
- Mount options for XFS on internal DAS (for cold buckets):
`Defaults`

9 Conclusion

Machine data is one of the fastest growing types of big data. That data, often referred to as digital exhaust, is generated in real time with sources including IT infrastructures, web servers, online retail and financial systems, machine sensors, and all the elements that constitute the IoT. Machine data is unstructured, has high velocity, and is large in volume, making it impossible to capture and analyze using traditional analytics tools, such as traditional relational databases.

Enterprise organizations depend on the continuous availability and analysis of data. In the area of healthcare, where critical medical devices are monitored in real time, seconds can mean the difference

between life and death. In the area of security, real-time analytics can help prevent financial fraud and attacks on sensitive computer systems.

Splunk provides the tools and capabilities that allow an enterprise organization to collect that data and extract high value from it, such as patterns of customer behavior, trending data that might predict equipment failure, and indications of financial fraud. Those tools and capabilities include visualization, extremely fast data ingest, real-time analytics, a rich set of APIs, notification capabilities, and extreme scalability.

However, traditional Splunk deployments with DAS are subject to server sprawl. The Splunk best practice recommendation to configure server storage with mirroring means that only half of a server's storage capacity is available for data. Also, adding storage capacity requires servers to be added, even if additional compute capacity is not needed. Splunk also relies on traditional methods for data protection and DR, which are slow and consume valuable storage and compute resources. Backups can't be created often enough to meet low RPOs, and recovery takes too long to meet RTO requirements.

The NetApp ONTAP and Splunk Enterprise solution is the ideal data protection and data management platform for enterprise-class organizations to capture and analyze machine data. As demonstrated by the tests described in this technical report, in most cases NetApp AFF performance exceeds that of DAS and it provides the following benefits:

- Backups can be created almost instantaneously by using NetApp Snapshot technology.
 - Those backups require little storage space
 - They can be used to quickly recover a Splunk system from catastrophic data loss
- Creation of Snapshot copies has zero impact on performance and is nondisruptive.
- NetApp storage is resilient: A panic condition was created on one of the storage controllers, with no effect on availability and performance.

Note: Based on workload, there could be a small impact of short duration on performance, but that was not observed during the tests.

- No effect on performance or on Splunk availability as shown by the disk failure test.
- ONTAP storage efficiency offers a 1.3:1 savings in storage utilization because of ONTAP inline storage efficiency, and an additional 2:1 savings without the traditional internal disk mirroring recommended by Splunk.
- Capability to create instant, free data copies by using FlexClone technology, a key feature for support of DevOps.

The proof points and test results described in this report clearly demonstrate that by decoupling compute from storage and implementing NetApp AFF storage in Splunk deployments, the following results are achieved:

- Higher performance
- A reduction in server requirements
- A reduction in storage requirements
- A more robust Splunk configuration
- Enterprise-class data protection
- Faster recovery from data loss
- Consistent performance, even in the event of storage hardware failure

These results mean reduced costs for power and data center floor space, and a significant increase in system reliability and availability. There is no room for compromise on Splunk system reliability and availability, especially in the areas of healthcare and security. NetApp AFF storage powered by ONTAP 9 is the logical choice for Splunk.

Acknowledgments

- Paul Burland, sales representative, NetApp
- Nilesch Bagad, senior product manager, NetApp
- Krister Eriksson, senior product manager, NetApp
- Mike McNamara, senior manager, product marketing, NetApp
- John Elliott, former senior technical marketing engineer, NetApp.

Most the content for this technical report was provided by John Elliott, who is no longer at NetApp. We are indebted to John's work and thank him for the tremendous effort.

Where to Find Additional Information

To learn more about the information presented in this document, refer to the following documents and websites:

- NetApp All Flash Arrays product page
<https://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- NetApp ONTAP data management software product page
<http://www.netapp.com/us/products/data-management-software/ontap.aspx>
- Splunk
<https://www.splunk.com/>
- Splunk App for NetApp ONTAP
<https://splunkbase.splunk.com/app/1293/>
- TR-4623: NetApp E-Series E5700 and Splunk Enterprise
<http://www.netapp.com/us/media/tr-4623.pdf>
- NetApp Product Documentation
<https://docs.netapp.com>

Version History

Version	Date	Document Version History
Version 1.0	December 2017	Initial release.
Version 2.0	June 2018	Revised to reflect current information.
Version 3.0	January 2019	Updated to include AFF A800 test results and new internal DAS test results.

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 © 2019 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.