



Technical Report

NetApp Architecture for Splunk

Walter Schroeder, Matt Hurford, Daniel Chan
Field Center of Innovation, NetApp

Brett Matthews, Splunk

May 2015 | TR-4260

Abstract

This technical report describes the integrated architecture of NetApp® and Splunk. Optimized for node storage balance, reliability, performance, and storage capacity and density, this design employs the managed DAS model, with higher scalability and lower TCO. In addition, this document summarizes the performance test results obtained from a universal storage benchmarking tool for various configurations for reference.

TABLE OF CONTENTS

1	Introduction	4
2	Splunk and Use Cases	4
2.1	Splunk Deployment Architectures	6
3	Splunk Use Cases.....	7
3.1	Splunk for IT Operations Management and Business Insights.....	7
4	NetApp E-Series Overview	10
5	Architecture Overview	11
6	Testing	14
7	Conclusion	16
Appendix A: Updated Testing Results		16
	Testing and Results.....	16
	Conclusion.....	18
Appendix B: Bonnie++		19
	Summary	19
	Source Tar File for Bonnie++.....	19
	Using Bonnie++	19
Appendix C: Splunk Apps for NetApp		21
	NetApp Performance App for Splunk.....	21
	Splunk App for NetApp StorageGRID.....	22
	Splunk App for Data ONTAP	23
Appendix D: References		25
	NetApp Documentation	25
	Splunk Documentation	25
	Splunk Apps	25
	Splunk Answers.....	25
	Books	25

LIST OF FIGURES

Figure 1)	Primary functions of Splunk.	5
Figure 2)	Centralized topology overview.	6
Figure 3)	Decentralized topology.....	7

Figure 4) Hybrid topology.	7
Figure 5) Event search.	8
Figure 6) Web store operational visibility.	8
Figure 7) Environment state.	9
Figure 8) Operation visibility.	9
Figure 9) 200GB/day design architecture.	12
Figure 10) 500GB/day design architecture.	13
Figure 11) Archive-based design architecture.	14
Figure 12) Bonnie++ output from CPOC.	15
Figure 13) Bonnie ++ test on a 4 indexer.	17
Figure 14) Bonnie++ on 16 indexers.	18
Figure 15) Example output of Bonnie++.	20
Figure 16) Bonnie++ results with machine memory “-r 24097.”	20
Figure 17) Bonnie++ results with 2 x machine memory “-r 48194.”	21
Figure 18) E-Series Performance Monitor.	22
Figure 19) Audit log dashboard.	23
Figure 20) Data ONTAP storage performance system.	24

1 Introduction

Splunk and NetApp have jointly developed reference architecture to provide guidance for successful Splunk deployments. This was done in conjunction with Splunk power partners to make sure that the deployment of the architecture is as simple as possible and covers regular issues around performance and sizing. All the elements for common deployments, from search engines through indexers to NetApp E-Series storage, have been considered and quantified.

The architecture is based on a common deployment and can be used as the basis for other Splunk architectures, including HA and replication. The architecture is modular to enable scalable deployments; therefore, as you deploy more indexers and/or search engines, storage is applied at an approximate rate.

A complete bill of materials (BOM) is also provided so that distributors can create a single SKU for their Splunk and NetApp partners to use for ordering. The BOMs are based on three options that relate to the standard Splunk licensing model.

This report is for NetApp, Splunk, and relevant partner SEs, as well as potential customers who want to implement a successful Splunk solution.

Currently there is no validated reference architecture for Splunk. Splunk believes that customers, in the absence of a validated architecture, are repurposing equipment for their Splunk deployments and this practice has resulted in suboptimal installations and many support calls and customer satisfaction issues.

The solution is architected to take into account the license levels typically sold by Splunk.

2 Splunk and Use Cases

Splunk is the leading operational intelligence software that enables you to monitor, report, and analyze live streaming and historical machine-generated data, whether it's located on premises or in the cloud. An organization's IT data is a definitive source of intelligence, because it's a categorical record of activity and behavior, including user transactions, customer behavior, machine behavior, security threats, and fraudulent activity. Splunk helps users gain visibility into this machine data that they can use to improve service levels, reduce IT operations costs, mitigate security risks, enable compliance, and create new product and service offerings.

Splunk collects all of your data sources—streaming and historical—by using a technology called *universal indexing*. Splunk is scalable enough to work across all your data centers and powerful enough to deliver real-time dashboard views to any level of the organization. Splunk offers solutions for IT operations, applications management, security and compliance, business analytics, and solutions for industrial data.

Splunk's architecture provides linear scalability for indexing and distributed search. Splunk's own implementation of MapReduce allows large-scale search, reporting, and alerting. Splunk takes a single search and allows you to query many indexers in massive parallel clusters. With the addition of index replication, you can specify how many copies of the data you want to make available to meet your availability requirements. Some of the deployment topologies are described in section 2.1, "Splunk Deployment Architectures."

The Splunk platform is open and has SDKs and APIs, including a REST API and SDKs for Python, Java, JavaScript, PHP, Ruby, and C#. This enables developers to programmatically interface with the Splunk platform. With Splunk you can develop your own applications or templates for deployment on your infrastructure.

Figure 1 shows the primary functions of Splunk.

Figure 1) Primary functions of Splunk.



Forwarders

A forwarder is a streamlined version of Splunk Enterprise that is used to send data to another instance of Splunk Enterprise.

Splunk forwarders come in two packages: the full Splunk distribution and a dedicated Universal Forwarder. The full Splunk distribution can be configured to filter data before transmission, execute scripts locally, or run the Splunk web server. The Universal Forwarder is an ultralightweight agent designed to collect data in the smallest possible footprint. Both types of forwarders provide automatic load balancing, SSL encryption and data compression, and the ability to route data to several Splunk instances or third-party systems.

The Deployment Server

Splunk comes with the ability to manage a distributed deployment without the need for third-party software. The deployment server helps you to synchronize the configuration of your search heads during distributed searching, as well as your forwarders to centrally manage your distributed data collection.

The Indexer

The indexer indexes the raw ingested data. The indexer transforms the raw data in events (individual pieces of data, that is, a single line for a log) and places the results into an index.

The core of the Splunk infrastructure is the indexer. This component requires the most IOPS because it performs constant writes during indexing and bursts of reads during search. An indexer has two functions. It accepts and processes new data, adds it to the index, and compresses it on the disk. The indexer also services search requests by looking through the data using its indexes and returning the results over a compressed communication channel.

Indexers scale out almost limitlessly and with almost no degradation in overall performance, allowing Splunk to scale from single-instance small deployments to truly massive big data deployments.

The Search Head

The search head is what most users initially interact with. It is the web server that provides the web-based user interface for the users to create their searches. As most of the data interpretation happens at search time, the role of the search head is to translate user and app requests into actionable searches for its indexers and display the results. The Splunk web UI is customizable, either through its own view and app system or by embedding Splunk searches in your own web apps using an API or SDKs.

The Cluster Master

The master node manages the cluster when configured. It coordinates the replicating activities of the peer (indexers) nodes and informs the search head about where to find the data. It also helps to manage the configuration of peer nodes and orchestrates remedial activities if a peer goes down.

The Reference Server

Splunk's sizing is based on commodity x-86 servers and is generally made up of the following:

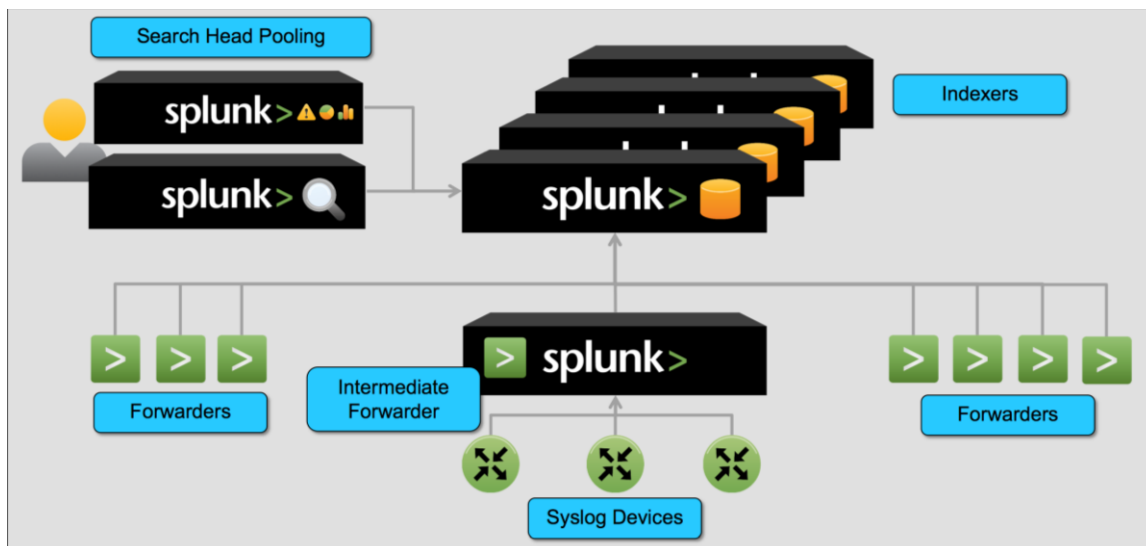
- Dual quad-core CPUs at 3.0 GHz (dual six-core is commonly used)
- 8GB of RAM (16GB is commonly used)
- 64-bit OS
- Disk performance should be at least 800+ IOPS per server

2.1 Splunk Deployment Architectures

Centralized Topology

The centralized topology allows a Splunk indexing cluster to be deployed at a location. The data that feeds the cluster also resides in the same location.

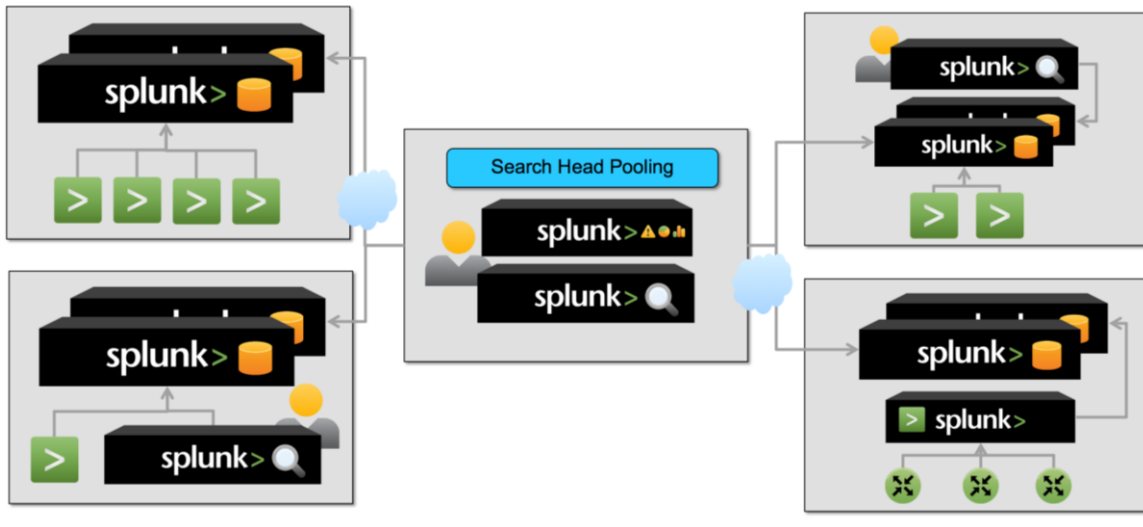
Figure 2) Centralized topology overview.



Decentralized Topology

A decentralized topology allows multiple clusters and single indexers to be federated by search heads. The location of these environments is distributed, and the data sources could also be colocated.


Figure 3) Decentralized topology.



Hybrid Topology

As the name suggests, a hybrid topology is a deployment that has one or many larger indexing clusters with smaller satellite clusters, all federated through search.

Figure 4) Hybrid topology.

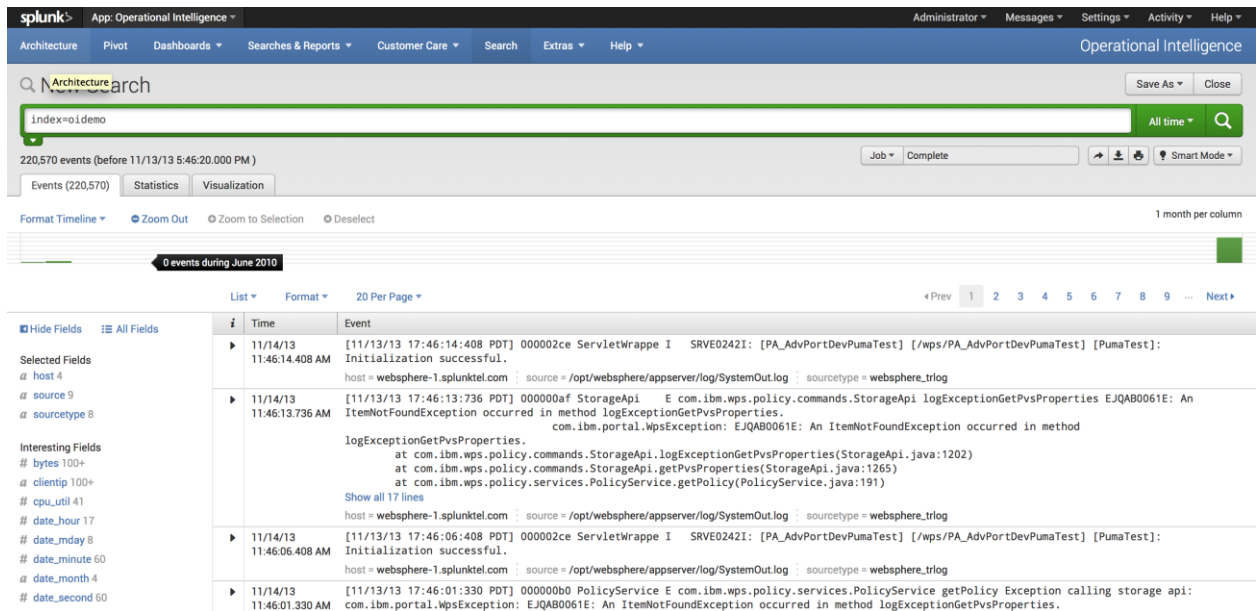


3 Splunk Use Cases

3.1 Splunk for IT Operations Management and Business Insights

Many different technologies and devices are layered and entwined to deliver business services. Virtualization and cloud computing multiply this complexity, especially when there are outages or performance issues. IT operations management teams and administrators waste valuable time moving from one console to another, logging onto one element after another, trying to track down the data they need to make sure of high performance and availability.

Figure 5) Event search.



Splunk collects and indexes data generated by your IT infrastructure. This data includes logs, file configurations, performance metrics, SNMP traps, and custom application logs. After the data is indexed, it provides a centralized federated view of all the machine information. Splunk's search interface is used to create the searches and rapidly prototype results in real time over historical data that is managed in Splunk.

Some organizations use Splunk for web store or e-commerce operational visibility and also to gain business insights from the same data. Because all data is indexed and unmodified and has full fidelity, you can derive security, operational, and business contexts from the same dataset.

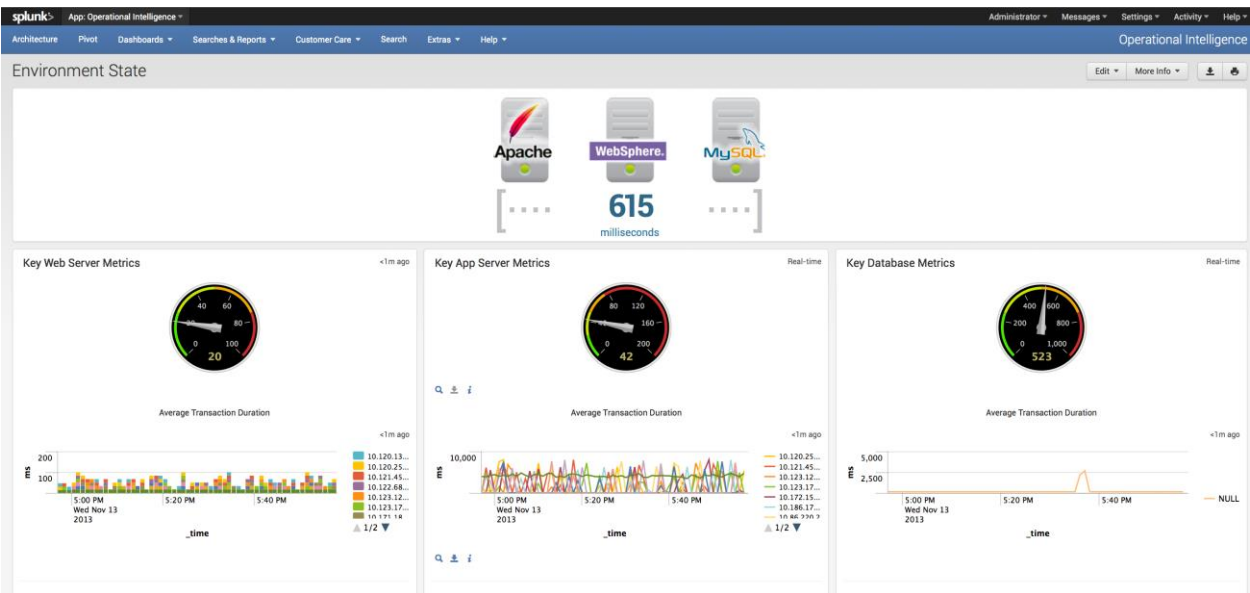
Figure 6 shows a typical landing page for an operational team, focusing on capacity and SLA infractions.

Figure 6) Web store operational visibility.



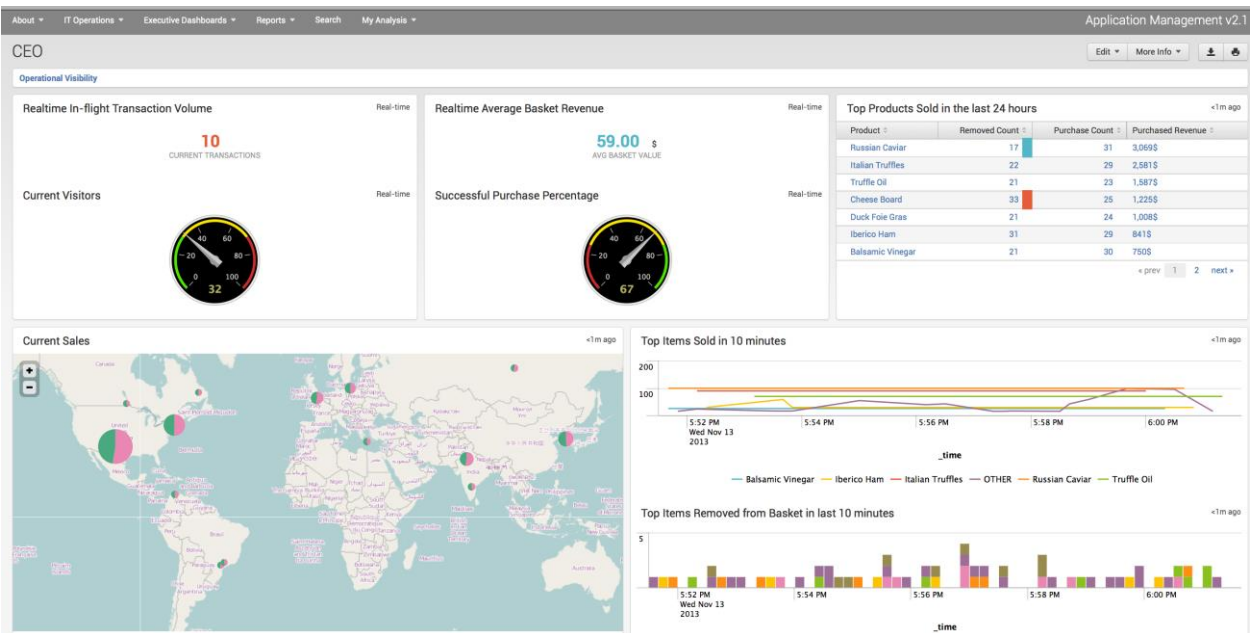
In Figure 7, the same data is represented with a more detailed breakdown of the key components of the service. You can see an end-to-end view, which indicates a problem with the database component.

Figure 7) Environment state.



In Figure 8, the same data is being presented to an executive. The web logs include a status code in the dashboard that provides business information, when combined with external data through a Splunk lookup and view of the top items sold. This view is defined in the search language and presented in a dashboard view in the Splunk UI, offering real-time visibility of the health of the environment and how the business is tracking with regard to online sales.

Figure 8) Operation visibility.



4 NetApp E-Series Overview

NetApp E-Series systems are high-performance, modular, block-based, seventh-generation storage arrays. Along with industry-leading performance and manageability, they offer a cost-effective platform for Splunk deployments. E-Series models include the lower end E2600 and E5400; the higher end E5500; and the all-flash array, EF540. E-Series systems offer large streaming write bandwidth (3.1GB/s for the E5400) and very high IOPS (greater than 350,000 4K IOPS for the EF540).

Splunk normally streams writes and random reads, and E-Series systems provide the best-in-class architecture for this operation. The E5400 storage array includes mechanical engineering and provides dense, scalable, highly reliable bandwidth and capacity. The disk controller firmware supports an optimal mix of high-bandwidth large-block streaming and small-block random I/O.

Splunk manages its data via the concept of buckets. It has hot, warm, cold and frozen. The hot bucket is used for data that has just been ingested and would be most frequently searched. As the data ages then the data is moved down the warm, cold frozen chain. The E-series is capable of having multiple types of drives (SSD, SAS, NL-SAS) in a single array it can accommodate a wide range of performance requirements across the different buckets in a single array.

As E-Series systems use NL-SAS rather than SATA for the larger capacity drives, the search speed across the cold buckets is generally increased by more than 20%.

NetApp E-Series storage arrays are the storage foundation that provides hundreds of petabytes of robust, highly available, back-end scalability and automated path failover, with redundant components and online administration, creating a rock-solid platform on which to build analytic platforms such as Splunk.

E-Series arrays:

- Provide high read and write throughput, so that data is committed to disk quickly and efficiently
- Minimize downtime with redundant power supplies and controllers
- Have extensive diagnostic and protection features, contributing to high levels of data integrity

The NetApp E5400 modular data storage system is architected for performance, density, and modular flexibility for wide-ranging data-intensive workloads. The E5400 meets an organization's demanding performance and capacity requirements without sacrificing simplicity and efficiency. Its balanced performance supports high-performance file systems, bandwidth-intensive streaming applications, and transaction-intensive workloads. The E5400 multiple drive shelf options enable custom configurations for any environment.

Based on 20 years of storage-development experience and a field-proven architecture, the E5400 delivers the highest reliability and 99.999% availability. Redundant components, automated path failover, and online administration mean that organizations stay productive 24/7/365. The E5400 advanced protection features and extensive diagnostic capabilities achieve consistently high levels of data integrity.

For more information, see <http://www.netapp.com/us/products/storage-systems/e5400/>.

NetApp SANtricity® is the management platform for the E-Series array, offering a powerful, intuitive administrative interface.

SANtricity storage management software enables storage administrators to achieve maximum performance and utilization of their storage through custom performance tuning and extensive configuration flexibility. This distinctive combination is especially important in high-performance environments with varying and often drastically different workloads and performance demands.

SANtricity software is designed to support the highest levels of data availability, integrity, and protection to maximize application uptime. Its automated I/O path failover and extensive online configuration, reconfiguration, and maintenance capabilities support 99.999% availability. With advanced technologies such as proactive monitoring, background repair, and extensive diagnostic features, SANtricity makes sure that your data is fully protected when it reaches the storage system.

For more information, go to <http://www.netapp.com/in/products/storage-systems/e2600/e2600-software.aspx>.

5 Architecture Overview

NetApp decided on a modular pod style, shared-nothing architecture to complement the distributed nature of typical Splunk implementations.

The array can be configured in several ways for different Splunk deployments, but we have narrowed the focus down to three of the most common configurations.

All testing was done using Bonnie ++ (linux.die.net/man/8/bonnie++), which Splunk recommends for testing and sizing configurations.

Splunk had been implemented on NetApp storage for many years. Plug-ins are available for the entire product range including FAS, Hadoop and the NetApp StorageGRID® object-based storage solution. We believe that deployments of Splunk can benefit from reference architecture, and NetApp recommends keeping Splunk deployments on a dedicated infrastructure. If Splunk is deployed on the infrastructure that it is trying to monitor, it can cause undesirable side effects during critical events by exponentially increasing the load on what it is designed to monitor.

There are successful installations of Splunk on FAS, but Splunk also provides its own set of data management features, such as compression and indexing, and so derives little benefit from the advanced data management features of FAS such as deduplication, compression, and replication.

We used a “pod-style” architecture for the NetApp reference architecture for Splunk. The base design is to have an E-Series array for each of two search engines and four indexers. The indexers and search engines are sized for common implementations on a Linux platform. We used a SAS interconnect between the indexers and the arrays. The SAS interface was used for its low-latency interconnect and low cost of deployment. No special drivers are needed for the LSI SAS cards used, because they are in the base Linux OS. For example, if the Splunk sizing exercise required 16 indexers, then four E-Series arrays would be required.


The POD design can be used in the three deployment architectures of Centralized, Decentralized and Hybrid discussed in section 2. Splunk is flexible enough that all three POD types described below could be deployed in any combination as close to the data to be ingested for speed or as far away as needed for protection.

This section details the sample BOMs. These BOMs are expected to cover more than 80% of implementations. The most common deviation is caused by the number of concurrent searches being performed, which affects the size and number of search engines.

For the up to 200GB/day design:

- 2 search engines
Cisco C220 servers
Dual socket, 16 cores + 32GB RAM
- 2 index servers
Cisco C220 servers
Dual socket, 8 cores + 24GB RAM
LSI 6Gb SAS HBA
- 1 E5424 system
24 900GB 10K drives: RAID 5 or RAID 6
2 6Gb SAS interfaces, one from each server: direct connect to servers

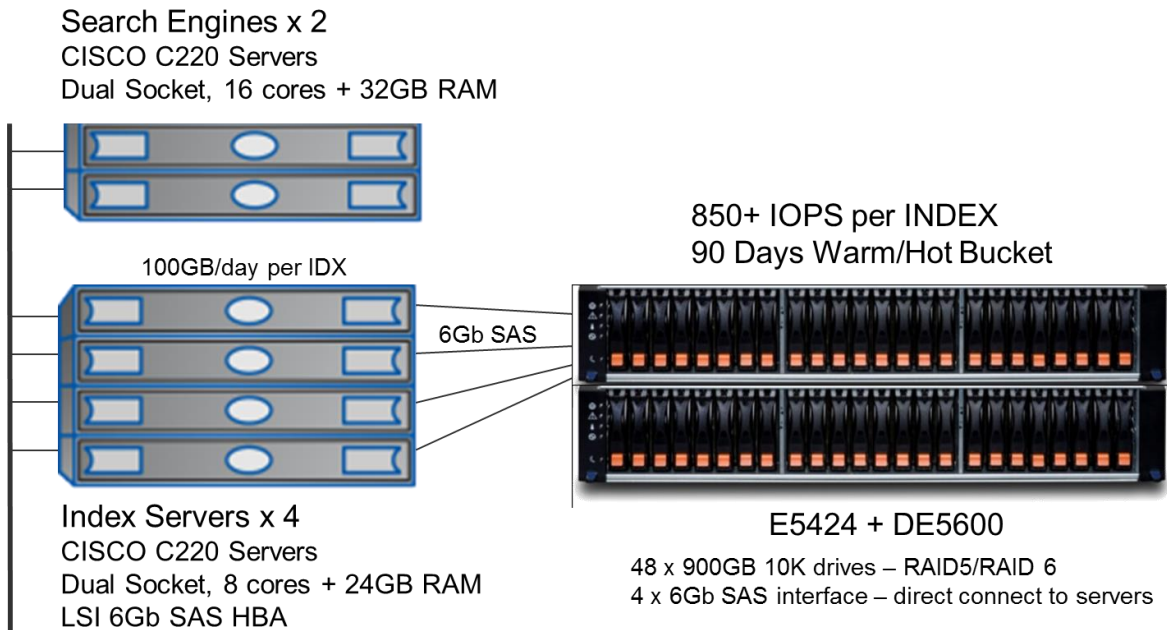
Figure 9) 200GB/day design architecture.



For the up to 500GB/day design, we added two more indexers and the extra capacity required:

- 2 search engines
Cisco C220 servers
Dual socket, 16 cores + 32GB RAM
- 4 index servers
Cisco C220 servers
Dual socket, 8 cores + 24GB RAM
LSI 6Gb SAS HBA
- 1 E5424 system
48 900GB 10K drives: RAID 5 or RAID 6
4 6Gb SAS interfaces: direct connect to servers


Figure 10) 500GB/day design architecture.



For the archive-based design, extra capacity was added with the 60-drive shelf:

- 2 search engines
Cisco C220 servers
Dual socket, 16 cores + 32GB RAM
- 4 index servers
Cisco C220 servers
Dual socket, 8 cores + 24GB RAM
LSI 6Gb SAS HBA
- 1 E5424
48 900GB 10K drives: RAID 5 or RAID 6
4 6Gb SAS interfaces: direct connect to servers
- 1 E5460 system
60 3TB NL-SAS 7.2K drives
150TB of usable “cold bucket” capacity

Figure 11) Archive-based design architecture.



The workload profile of Splunk (sequential write and random reads) is well suited to the E-Series architecture. The I/O size of Splunk varies frequently, even within a deployment, depending on the data stream and the number of data streams being ingested and their compressibility. This data stream, combined with the expected search workload, determines the number of indexers and search nodes required. It is up to a Splunk power partner to design the specific architecture.

The architecture design employs a shared-nothing approach, allowing each indexer to think that it has dedicated storage. The E-Series storage solution provides storage services to the indexers. Each indexer accesses its own private volume of storage on the E-Series array. The Splunk solution provides four volumes of private storage with the enterprise-grade capabilities often found in a traditional SAN environment. E-Series systems are frequently configured (and were originally designed) for SAN environments, but they are configured specifically to function as four volumes of dedicated storage. Combining four volumes into a single package gives each indexer 20 drives, as well as array intelligence to dramatically improve data availability and manageability.

The solution scales as indexers are added. For every four indexers, a storage node must be included. This allows a simple and very effective scaling model.

Splunk deployments have been observed to include data ingestion rates from 500MB/day to more than 20TB/day. This data can be retained for varying amount of time, to more than 7 years. 2TB/day, retained for 2,555 days (7 years), is potentially more than 50PB of data to be stored, compressed, and indexed by Splunk. Some deployments have even included server-side cache to hold the hot buckets.

Several Splunk platinum partners have used the reference architecture as a basis for deployment; all report excellent results and high satisfaction.

6 Testing

The tests were performed in the NetApp Customer Proof of Concept Lab (CPOC) using Bonnie++ (linux.die.net/man/8/bonnie++), the universal storage benchmarking tool that is used to establish a baseline for storage performance at customer sites. The Splunk white paper in Appendix A is an extract of the full description available from Splunk.

Figure 12 shows the Bonnie++ output from the CPOC testing. The value of interest is in the Random Seeks column. The “random seeks per second” value is the number used to compare different configurations. For the purpose of sizing and validation, none of the other fields are important and can be disregarded. The output of the tool is shown for reference and to indicate where you can find the output required.

The E5400 storage array was used because it is easily available from QuoteEdge and offers the greatest flexibility to the customer in terms of capacity and protocols that can be used.

In a single or dual indexer configuration, testing with Bonnie++ showed more than 1,250 random seeks per indexer. When extra indexers were added, the average random seeks grew to more than 850 per indexer, as shown in Figure 12.

Figure 12) Bonnie++ output from CPOC.

cpu MHz : 2533.415		Sequential Output						Sequential Input				Random Seeks		
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU	
x3550-m3-16.cpoc.local	94G			798021	52	338347	31			710404	37	817.8	1	16

cpu MHz : 2400.110		Sequential Output						Sequential Input				Random Seeks		
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU	
x3550-m3-52.cpoc.local	94G			795211	62	336687	34			677586	40	855.7	2	16

cpu MHz : 2399.831		Sequential Output						Sequential Input				Random Seeks		
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU	
x3550-m3-53.cpoc.local	94G			795777	61	328577	36			693671	44	855.3	2	16

cpu MHz : 2533.449		Sequential Output						Sequential Input				Random Seeks		
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU	
x3550-m3-23.cpoc.local	94G			798520	53	343641	32			706141	29	839.9	1	16

7 Conclusion

The NetApp reference architecture for Splunk is optimized for node storage balance, reliability, performance, storage capacity, and density.

Organizations that use Splunk often use traditional server-based storage with inefficient, hard-to-scale, internal direct-access storage (DAS). The NetApp reference design employs the managed DAS model, with higher scalability and lower TCO. The lower cost comes from choosing the number of clusters and storage you want instead of spending money for features that you don't need. It also comes from nonproprietary lower cost options and by decoupling compute nodes from storage, preventing unnecessary purchases of compute nodes for storage-intensive workloads.

Appendix A: Updated Testing Results

Since the writing of this TR, NetApp has continued to innovate on the E-Series storage arrays. NetApp has introduced several new models of E-Series array: the E2700 and the E5600, as well as the EF560 all-flash storage line. The evolution of the E-Series continues to improve performance, density, and features such as allowing for the hot, warm, and cold storage to be serviced from the same controller set.

Splunk workloads are performance sensitive, and the innovations made with NetApp storage improve on the existing best-in-class architecture. The massive increase in performance and density allows for more indexing capability in a smaller footprint. In addition to the traditional protection mechanisms NetApp provides such as RAID 5 and RAID 6, T10-PI (write failure detection on the disk), we have also introduced Dynamic Disk Pools (DDP). DDP is a next-generation parity-based protection mechanism that improves the recovery for large drive failures to minutes while providing greater protection and performance than RAID 6.

The E2700 replaces the E2600, improving bandwidth and IOPS over its predecessor by a wide margin. The E2700 is capable of supporting 192 disks, and, as with previous generations, one can mix drive types (SSD, SAS, and NL-SAS) at 8GB/s or read at over 80,000 IOPS.

For more information, see <http://www.netapp.com/us/products/storage-systems/e2700/index.aspx>.

The newest member of the high-performance family is the E5600, doubling the performance of the older E5400 in terms of bandwidth and IOPS. The E5600 is now capable of supporting 384 drives (SSD, SAS, and NL-SAS) at 12GB/s and 650,000 IOPS. NetApp has also improved the industry-leading EF540 to the EF560 all-flash array. Here too the performance improvement is an order of magnitude greater than the previous generation. The data represented here is based on the EF5600.

For more information, see <http://www.netapp.com/us/products/storage-systems/e5600/index.aspx>.

NetApp recently had the opportunity to test these new arrays with Bonnie++, with a larger number of indexers than in our previous testing. The results of this testing is included, while additional Splunk use cases will be included in a new technical report.

The opportunity was taken to also incorporate the newer DDP technology in the testing this time. The previous round of testing was done with traditional RAID groups. DDP gives the user a number of advantages, including consistency of setup, performance, and easier management. Additional benefits of DDP are faster rebuilds and lower impact to Splunk while rebuilds of the disk protect are occurring. It is worth pointing out that the E-Series has demonstrated (as validated by IDC) a five 9s uptime (99.999%). As a result, unplanned interruptions are low in any case.

Testing and Results

The testing was done with the cooperation of our CPOC labs. Two main types of tests were run.

- Testing with Bonnie++ was done for a few customers using the current architecture with an average to high ingest rate (100–300GB/day), but had very high query rates. For those customers we were

able to show Bonnie++ IOPS of over 14,000 IOPS per indexer in a 4 indexer to a single E-Series array configuration, as seen in Figure 13.

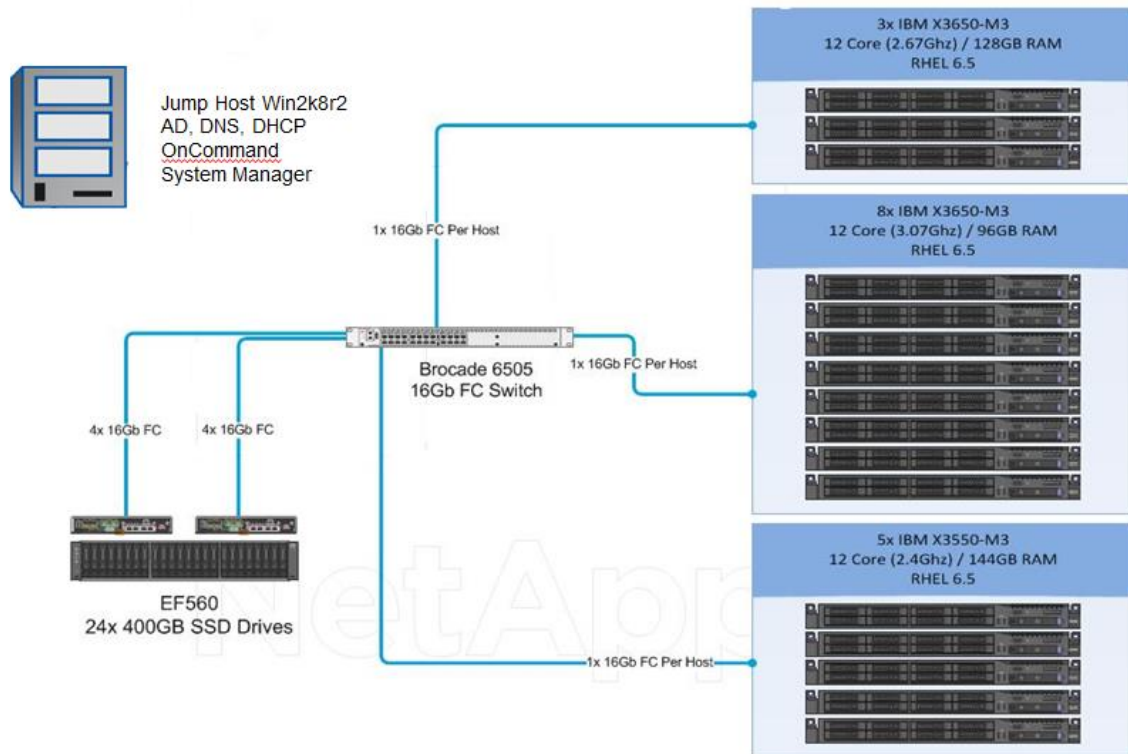
Figure 13) Bonnie ++ test on a 4 indexer.

		Sequential Output						Sequential Input				Random Seeks		Sequential Create						Random Create					
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files	Create		Read		Delete		Create		Read		Delete
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU		/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	
s1	98G			150974	20	117260	14			502867	24	13014.5	24	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++	
		Sequential Output						Sequential Input				Random Seeks		Sequential Create						Random Create					
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files	Create		Read		Delete		Create		Read		Delete
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU		/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	
s2	98G			138153	19	113688	13			616826	31	14117.5	25	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++	
		Sequential Output						Sequential Input				Random Seeks		Sequential Create						Random Create					
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files	Create		Read		Delete		Create		Read		Delete
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU		/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	
s3	98G			128900	18	130766	21			556057	25	13052.2	21	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++	
		Sequential Output						Sequential Input				Random Seeks		Sequential Create						Random Create					
	Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block				Num Files	Create		Read		Delete		Create		Read		Delete
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/sec	% CPU		/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	/sec	% CPU	
s4	98G			138328	19	115991	13			550087	24	14876.1	24	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++	

- With the successful results seen in Figure 13, it was concluded that an architecture was possible in which a customer could increase the indexers dramatically, increasing the ingest rate capability and the number concurrent searches.

The tests were conducted on the EF560 and the E2700 array models. We increased the attached indexer count to 16 to a single array with 24x 400GB SSD drives inside a DDP pool. 16 separate LUNs were created and presented to each of the 16 indexers, as seen in Figure 14.

Figure 14) Bonnie++ on 16 indexers.



For the EF560 running with Bonnie++ we were able to show an average of 8,400 Bonnie++ IOPS per indexer, giving us an aggregate performance of over 135,000 Bonnie++ IOPS.

For the E2700 with the same DDP configuration, we were able to show an average per indexer of 2,280 Bonnie++ IOPS, giving us an aggregate performance of over 36,500 Bonnie++ IOPS on the low-end platform. This is a considerably larger amount over the 1,200 IOPS suggestion by Splunk. This can all be provided in a 34RU (16x 2RU for the servers and 1x 2RU for the disks) footprint.

Conclusion

NetApp has been able to demonstrate a considerable increase in the capabilities of its latest arrays. The performance increase allows an increase in the flexibility of the Splunk deployments that are now possible. From our originally suggested architecture of 4 indexers per array ingesting around 400–500GB/day, we now deploy in 34 RU, 16 indexers per E-Series array, ingest around 4.8TB/day and query rates that put the architecture at the top of even the most demanding of customer requirements.

The NetApp EF5600 provides a significant performance improvement over traditional deployments. It essentially shrinks multiple racks of servers with internal drives to a single rack with greater performance and reliability. With the NetApp approach anyone can have a hot, warm, and cold storage system capable of outperforming servers with internal SSD as well. The scalability and density provide a most efficient Splunk deployment.

Appendix B: Bonnie++

Summary

To establish a baseline for storage performance, the universal storage benchmarking tool Bonnie++ must be installed on an indexing server.

Bonnie++ is distributed in source form; therefore, you need a machine with a C++ compiler on it. On an initial Splunk installation or preparation, Bonnie++ should be built and installed, so the setup and build sections should need to be done only once per target OS.

Source Tar File for Bonnie++

You can download Bonnie++ at <http://www.coker.com.au/bonnie++/>. If you want to try the latest release, go to <http://www.coker.com.au/bonnie++/experimental/>.

The following examples use Bonnie++ version 1.03e.

Using Bonnie++

This section discusses:

- Usage
- Formatting results
- Interpreting results

Usage

`./bonnie++`

The options for Bonnie++ are listed as follows. You must use the `-u` switch when running as root user:

```
bonnie++ [-d scratch-dir] [-s size(MiB) [:chunk-size(b)]]  
          [-n number-to-stat[:max-size[:min-size]][:num-directories]]]  
          [-m machine-name]  
          [-r ram-size-in-MiB]  
          [-x number-of-tests] [-u uid-to-use:gid-to-use] [-g gid-to-use]  
          [-q] [-f] [-b] [-D] [-p processes] [-y]
```

For more information, refer to the Bonnie++ man page at linux.die.net/man/8/bonnie++.

Formatting Results

Bonnie++ example (HTML):

Run the following command:

```
./bonnie++ -u root -d/home/<username>/scratchdirectory -x 1 -q -f | ./bon_csv2html.in > disk-io-  
test.html
```

In a browser, display `disk-io-test.html`. Figure 15 is an example of the output.

This example was run on a very small Linux virtual machine without the `-r` option for the purpose of illustration. The `-r` option is RAM size in megabytes.

Figure 15) Example output of Bonnie++.

		Sequential Output					Sequential Input		Random Seeks				Sequential Create				Random Create									
	Size:Chunk Size	Per Char	Block		Rewrite		Per Char	Block					Num Files	Create		Read		Delete		Create		Read		Delete		
		K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec					% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	
fedora10-01	2G			22589	6	15909	5			81267	24	275.1	0	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++

Running on Indexers

To establish a baseline for storage performance, the universal storage benchmarking tool Bonnie++ (v 1.0.3e) must be installed and run on the indexing server. The following options and arguments were used:

```
./bonnie++ -u root -d /home/<username>/scratchdirectory -r <memsizeinMB> -x 1 -q -f |
./bon_csv2html.in > disk-io-test_wM.html
```

The options used in this command are:

- -d is the directory to use for the test.
- -r is RAM size in megabytes. If you specify the -r option, other parameters are checked to make sure that they are suitable for a machine with the RAM specified.
- -x is the number of test runs.
- -q is run in quiet mode.
- -f is run in fast mode; skips per-char I/O tests.

Determining Memory on Linux

To determine the amount of memory, run the following command:

```
cat /proc/meminfo
```

Look for the MemTotal: value, which gives the amount of RAM in kB. You need to convert it to megabytes.

As a second benchmark, it is a best practice to double the amount of physical memory as the argument to the --r option in order to limit the effect of memory and controller caching on the results. To accomplish this, perform an additional run using the following options and arguments:

```
./bonnie++ -u root -d /home/<username>/scratchdirectory -r <memsizeinMB*2> -x 1 -q -f |
./bon_csv2html.in > disk-io-test_wM.html
```

Figure 16 and Figure 17 show the results of the sample benchmark on target Splunk indexer and indexes.

Figure 16) Bonnie++ results with machine memory “-r 24097.”

	Sequential Output						Sequential Input				Random Seeks		Sequential Create						Random Create					
Size:Chunk Size	Per Char	Block		Rewrite		Per Char	Block		Num Files	Create			Read		Delete		Create		Read		Delete			
	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU		/ sec			% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	
48194M			291384	39	125372	21		347206	20	653.7	1	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++		

Figure 17) Bonnie++ results with 2 x machine memory “-r 48194.”

	Sequential Output						Sequential Input				Random Seeks		Num Files		Sequential Create						Random Create					
Size:Chunk Size	Per Char		Block		Rewrite		Per Char		Block						Create		Read		Delete		Create		Read		Delete	
	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	K/sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU	/ sec	% CPU				
96388M			272512	42	119310	20			345904	22	382.6	1	16	+++++	+++	+++++	+++	+++++	+++	+++++	+++	+++++	+++			

Because Splunk is I/O bound, a higher `seek` metric results in faster searches. A lower `seek_cpu` metric indicates the potential for better concurrency and more headroom before the CPU becomes the predominant bottleneck.

Ideally, the `seek` metric would be around 800 I/Os, with `seek_cpu` around 1%. A real-world rule of thumb is that around 500 I/Os yields an above-average to excellent search performance.

Therefore, the 653.7 I/Os `seek` performance with 1% `seek_cpu` on the first trial is easily in the top 5% of results that Splunk solution architects have observed. Even the second run, when the amount of physical memory reported to Bonnie++ was doubled, it resulted in a highly respectable 400 I/Os.

The `seq_create` and `seq_create_cpu` metrics measure the ability of the storage subsystem to create data in a sequential manner, which is for the most part the type of operation that Splunk performs during indexing and optimization. In this case, the optimal results for these metrics are “+++++” and “+++,” respectively, indicating that this portion of the test completed so quickly (<500ms) that it was impossible to measure accurately.

Appendix C: Splunk Apps for NetApp

The relationship between NetApp and Splunk extends to developing apps for the NetApp product range. The current products supported with Splunk apps include NetApp clustered Data ONTAP®, Data ONTAP operating in 7-Mode, SANtricity app for E-Series arrays, and StorageGRID.

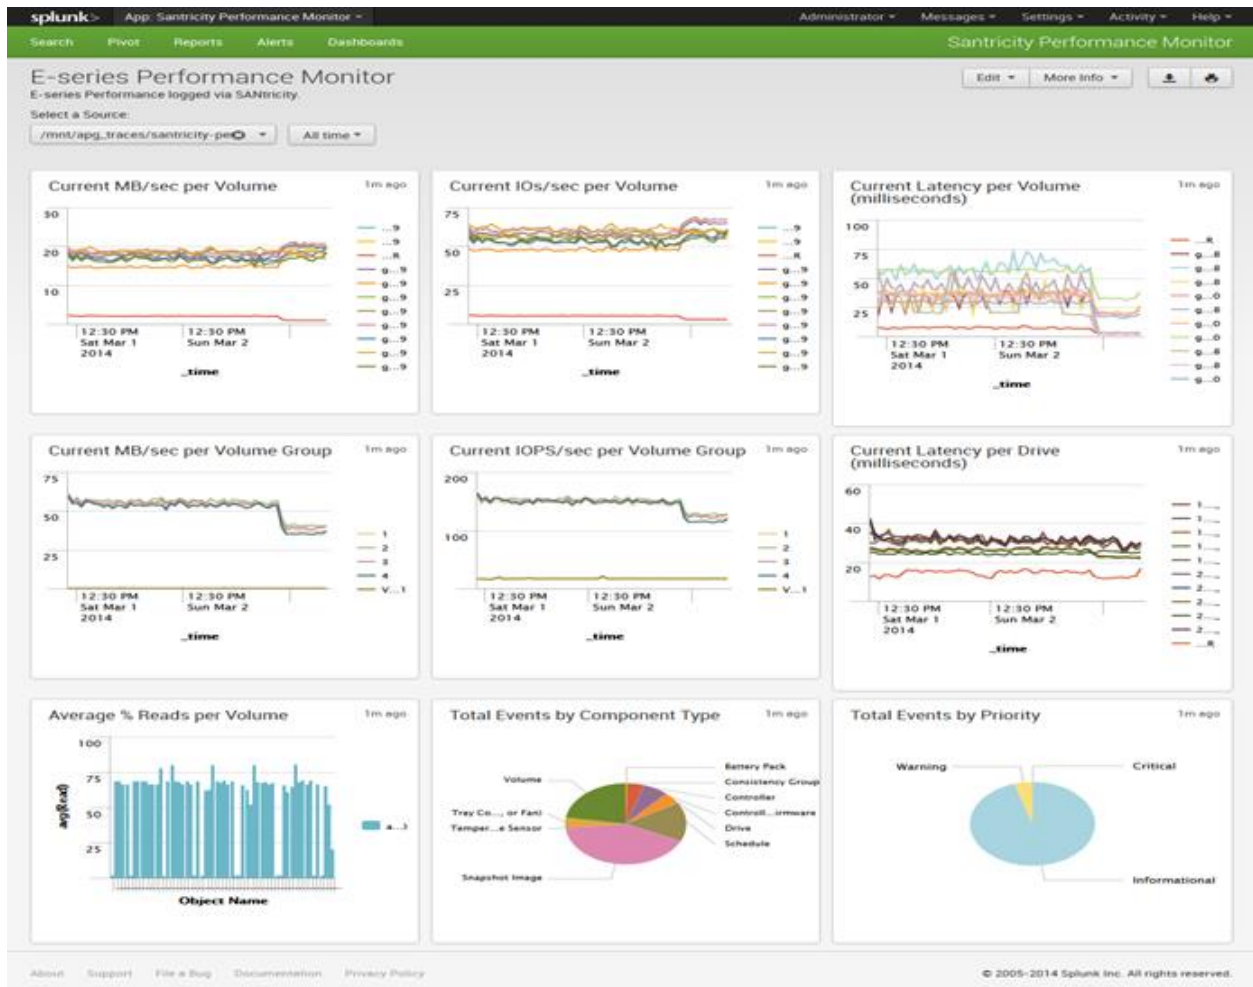
These apps are all available at the Splunk base portal at <http://apps.splunk.com/>.

NetApp Performance App for Splunk

The NetApp SANtricity performance app for Splunk Enterprise provides visibility into the health and performance of NetApp E-Series and EF-Series storage systems. Figure 18 displays the configuration information by using a dashboard view for multiple E-Series storage systems. This makes it easy to drill down to specific configuration information for each array such as IOPS, MB/sec, latency, information on controller, DDP, pools, volume groups, volume, and drives. It also displays MEL information.

The app can be downloaded from <https://apps.splunk.com/app/1932/> and also requires <https://apps.splunk.com/app/1933> to run.

Figure 18) E-Series Performance Monitor.

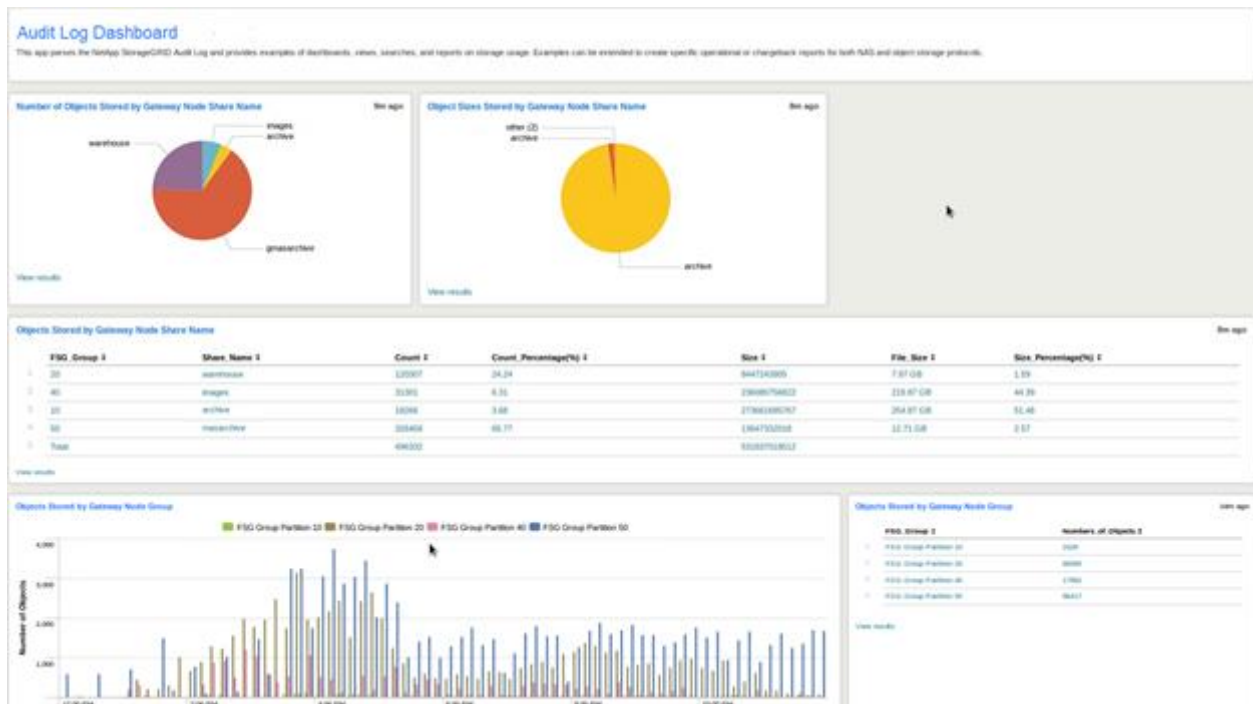


Splunk App for NetApp StorageGRID

The Splunk app for StorageGRID provides real-time visualization and reporting of audit log parsing information for billing and usage monitoring. This allows chargeback and billing, search integration, custom reporting, security diagnostics, and alerts for compliance events. It also provides a view of CIFS/NFS activity and CDMI/SGAPI usage.

The app can be downloaded from <https://apps.splunk.com/app/1436/>.

Figure 19) Audit log dashboard.



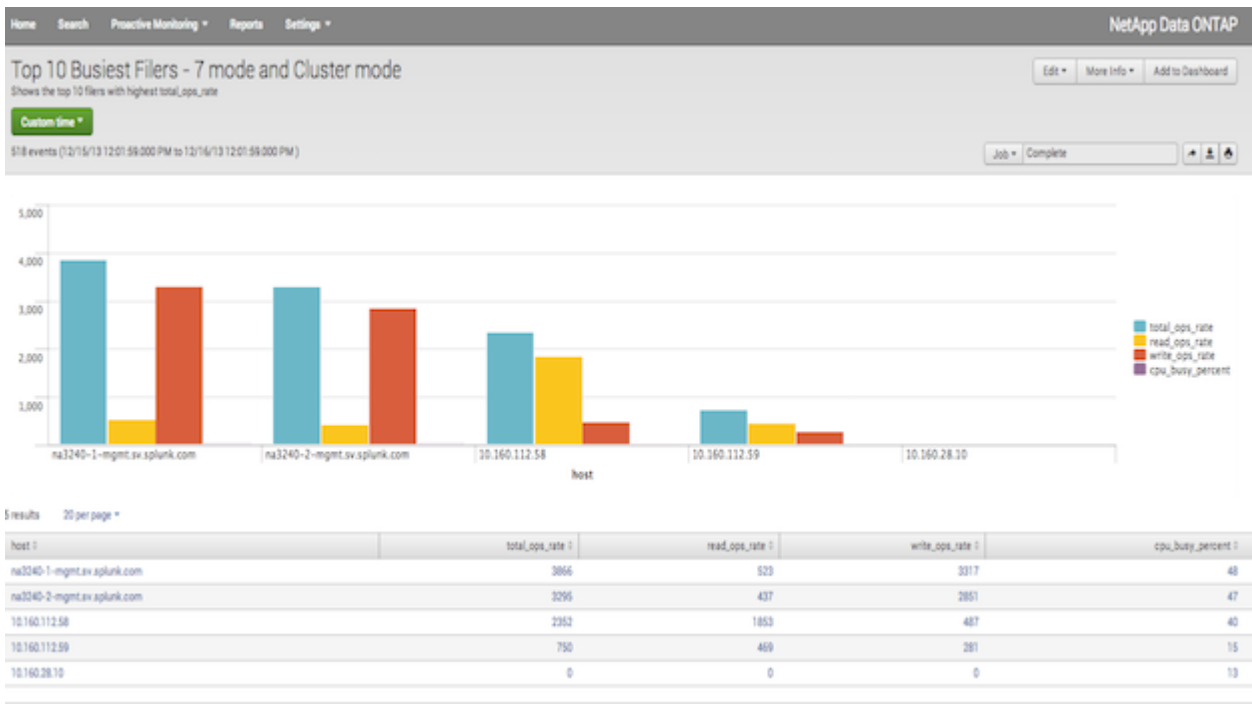
Splunk App for Data ONTAP

This app is a single solution that provides instant operational visibility into the health of your Data ONTAP storage systems. It allows you to quickly visualize configuration, logs, and performance in both clustered Data ONTAP and Data ONTAP operating in 7-Mode, with quick drilldowns to specific subsystems including storage system, aggregate, volume, disk, and other events.

The Splunk app for Data ONTAP is compatible with Data ONTAP 8.x and above.

The app can be downloaded from <https://apps.splunk.com/app/1293/>.

Figure 20) Data ONTAP storage performance system.



Appendix D: References

NetApp Documentation

- For information on NetApp E2700 storage system:
<http://www.netapp.com/us/products/storage-systems/e2700/index.aspx>
- For information on NetApp E5600 storage system:
<http://www.netapp.com/us/products/storage-systems/e5600/index.aspx>

Splunk Documentation

- <http://docs.splunk.com/Documentation>
- For information about installing Splunk:
<http://docs.splunk.com/Documentation/Splunk/latest/Installation/Whatsinthismanual>
- For information about sizing your Splunk deployment:
<http://docs.splunk.com/Documentation/Splunk/latest/Installation/CapacityplanningforalargerSplunkdeployment>
- For information about scaling and clusters:
<http://docs.splunk.com/Documentation/Splunk/latest/Installation/CapacityplanningforalargerSplunkdeployment>
- Splunk's documentation covers index size. In addition, the following blog is a reference for managing index size in Splunk. This blog is quite detailed and covers retention policy:
<http://blogs.splunk.com/2011/01/03/managing-index-sizes-in-splunk/>

Splunk Apps

<http://apps.splunk.com/>

Splunk Answers

<http://answers.splunk.com/>

Books

- *Exploring Splunk*, by David Carasso
<http://www.splunk.com/goto/book>
- *Implementing Splunk: Big Data Reporting and Development for Operational Intelligence*, by Vincent Bumgarner
http://www.amazon.com/Implementing-Splunk-Development-Operational-Intelligence/dp/1849693285/ref=sr_1_2?ie=UTF8&qid=1384323844&sr=8-2
- *Big Data Analytics Using Splunk: Deriving Operational Intelligence from Social Media, Machine Data, Existing Data Warehouses, and Other Real-Time Streaming Sources*, by Peter Zadrozny and Raghu Kodali
http://www.amazon.com/Big-Data-Analytics-Using-Splunk/dp/143025761X/ref=sr_1_3?ie=UTF8&qid=1384323844&sr=8-3&

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

Copyright Information

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

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

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

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

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

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

Trademark Information

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

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4260-0515