.conf19

splunk>

Operational Efficiencies with Splunk SmartStore

October 29, 2019

Somu Rajarathinam

Technical Director | Pure Storage

@purelydb | www.somu.us

Operational Efficiencies with SmartStore

On Pure FlashBlade

SmartStore – What & Why?

Operational Efficiency Tests

- Data Rebalance
- Indexer Node Failure

Why SmartStore?

CLASSIC SPLUNK

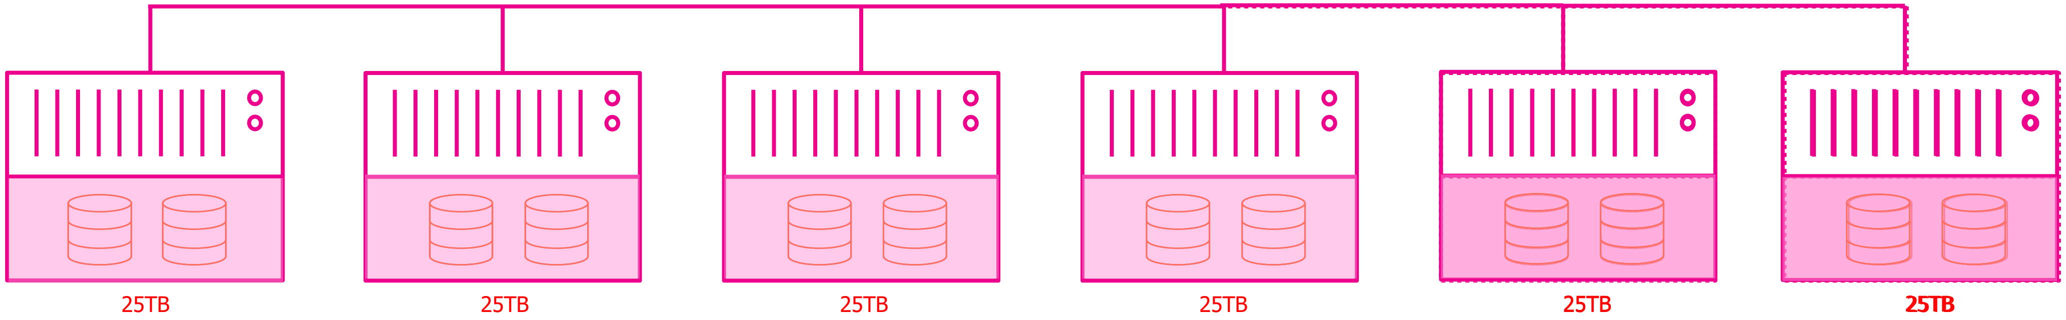

SPLUNK SMARTSTORE

Grow compute independently

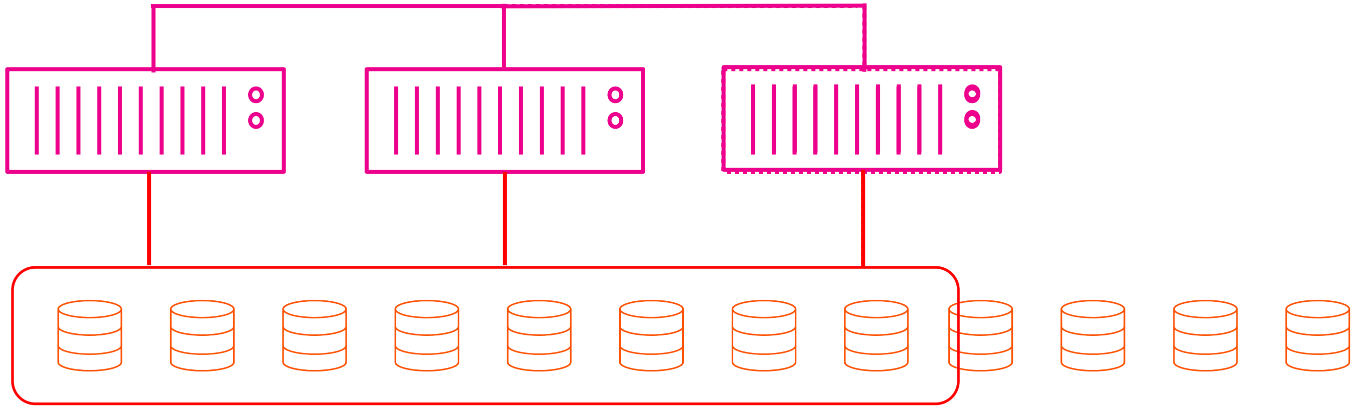

Grow storage independently

What is Splunk SmartStore?

SmartStore – A new Splunk indexer functionality

- Disaggregated storage from compute enables efficient resource usage
- Dynamically scale compute & storage on-demand
- Reduce the overall TCO by achieving cost savings with flexible storage options
- Simplify Indexer maintenance without impacting data integrity
- Storage Platform provides all data services (protection, replication, etc)

Splunk SmartStore Operational Efficiencies

Tests & Results

splunk> .conf19

Operational Efficiency Tests

On Pure FlashBlade

- Data Rebalance
- Indexer Node Failure

Data Rebalance

What is data rebalance?

- Balance the storage distribution across the indexer peer nodes
- Redistribute bucket copies so each peer node has similar number of copies
- Operates on Warm and cold buckets only

When do you invoke data rebalance?

- Immediately after adding a new indexer
- Uneven forwarding of log data

Caveats

- Do not use searchable data rebalance with SmartStore indexes

Data Rebalance

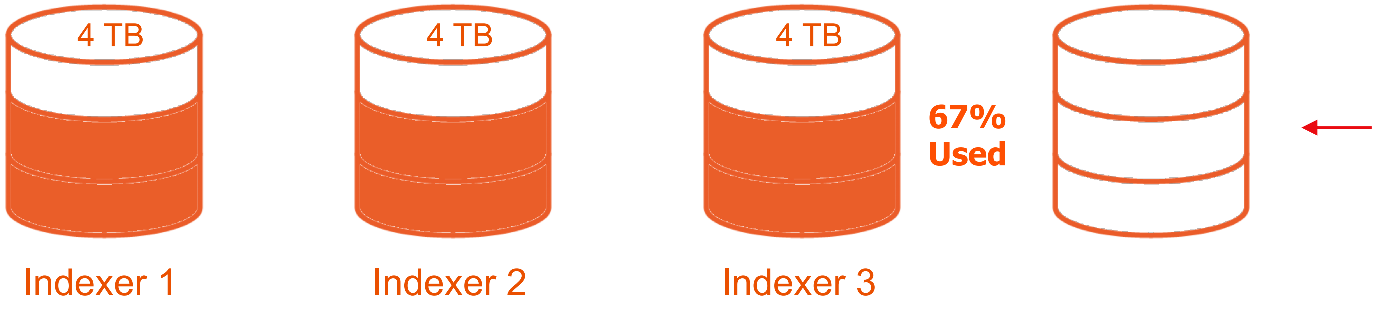

Adding storage capacity requires adding servers

Perform Data Rebalance to distribute the storage across the peer nodes

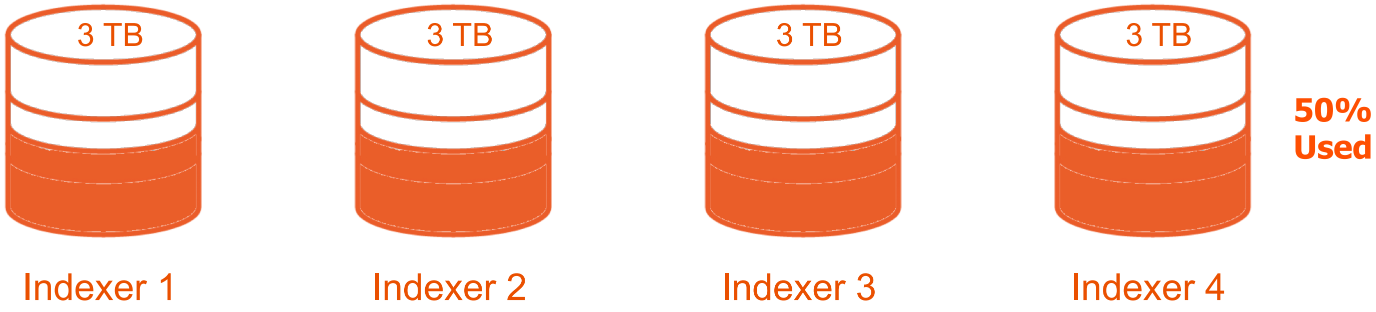

Operational Efficiency Test

INDEXER NODE ADDITION & DATA REBALANCE

Data rebalance redistributes data across the indexer cluster.

- Non-searchable data rebalance
- 79K buckets (30 TB) on 8 indexers
- Added the 9th indexer node

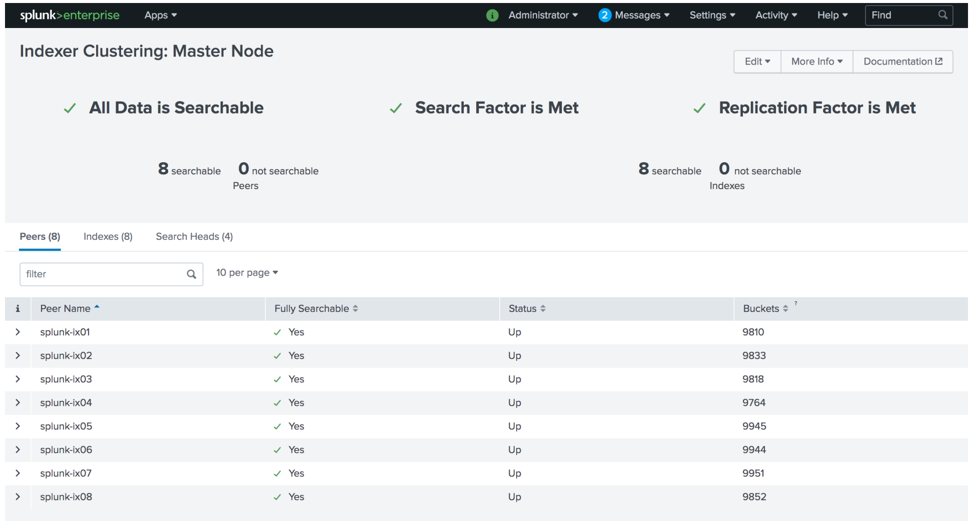

Indexer Clustering: Master Node

✓ All Data is Searchable ✓ Search Factor is Met ✓ Replication Factor is Met

8 searchable 0 not searchable
Peers

8 searchable 0 not searchable
Indexes

Peers (8) Indexes (8) Search Heads (4)

filter 10 per page

i	Peer Name	Fully Searchable	Status	Buckets
›	splunk-ix01	✓ Yes	Up	9810
›	splunk-ix02	✓ Yes	Up	9833
›	splunk-ix03	✓ Yes	Up	9818
›	splunk-ix04	✓ Yes	Up	9764
›	splunk-ix05	✓ Yes	Up	9945
›	splunk-ix06	✓ Yes	Up	9944
›	splunk-ix07	✓ Yes	Up	9951
›	splunk-ix08	✓ Yes	Up	9852

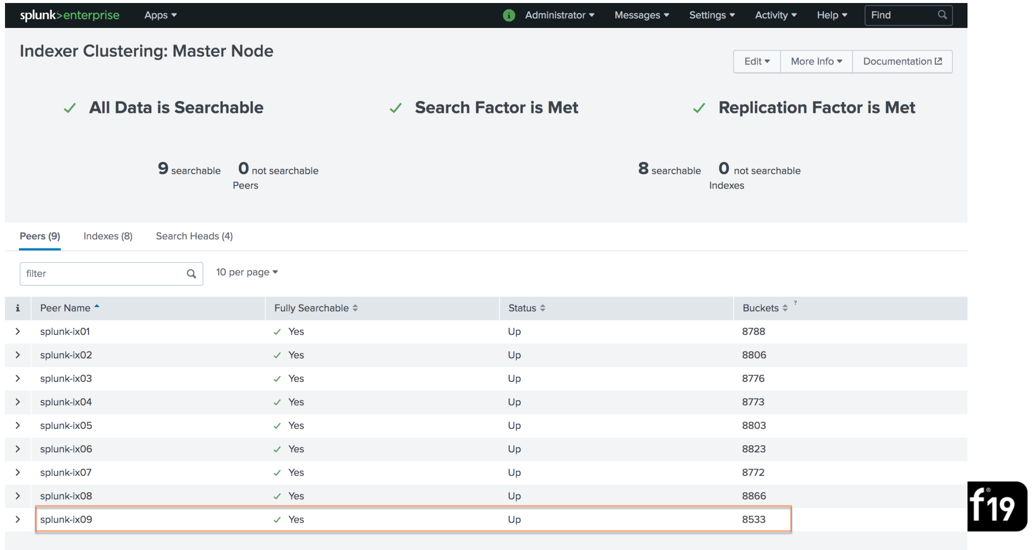

Indexer Clustering: Master Node

✓ All Data is Searchable ✓ Search Factor is Met ✓ Replication Factor is Met

9 searchable 0 not searchable
Peers

8 searchable 0 not searchable
Indexes

Peers (9) Indexes (8) Search Heads (4)

filter 10 per page

i	Peer Name	Fully Searchable	Status	Buckets
›	splunk-ix01	✓ Yes	Up	8788
›	splunk-ix02	✓ Yes	Up	8806
›	splunk-ix03	✓ Yes	Up	8776
›	splunk-ix04	✓ Yes	Up	8773
›	splunk-ix05	✓ Yes	Up	8803
›	splunk-ix06	✓ Yes	Up	8823
›	splunk-ix07	✓ Yes	Up	8772
›	splunk-ix08	✓ Yes	Up	8866
›	splunk-ix09	✓ Yes	Up	8533

Operational Efficiency Test

Data Rebalance comparison between SmartStore and Classical Splunk

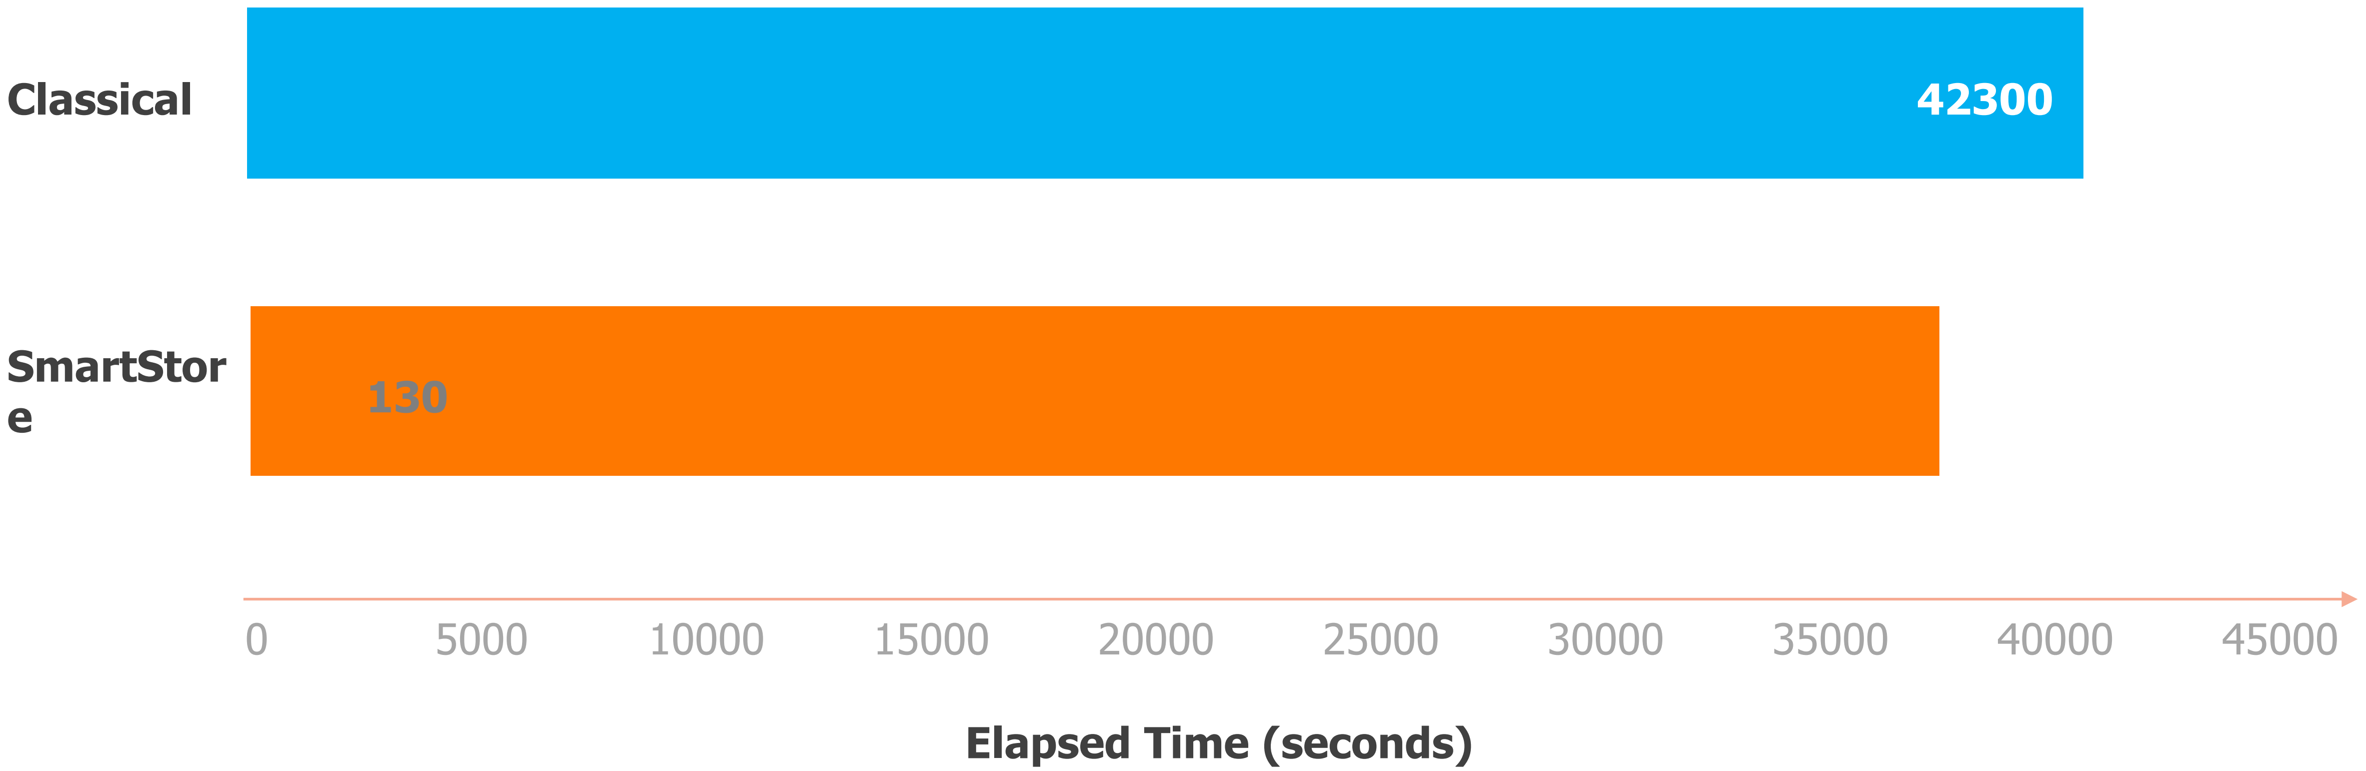

Operational Efficiency Test

INDEXER NODE ADDITION & DATA REBALANCE

Data rebalance comparison of SmartStore against Classical Splunk

- 2 minutes 10 seconds under SmartStore
- 11 hours and 45 minutes under Classical
- Faster under SmartStore as actual data is not redistributed but just the buckets' pointers

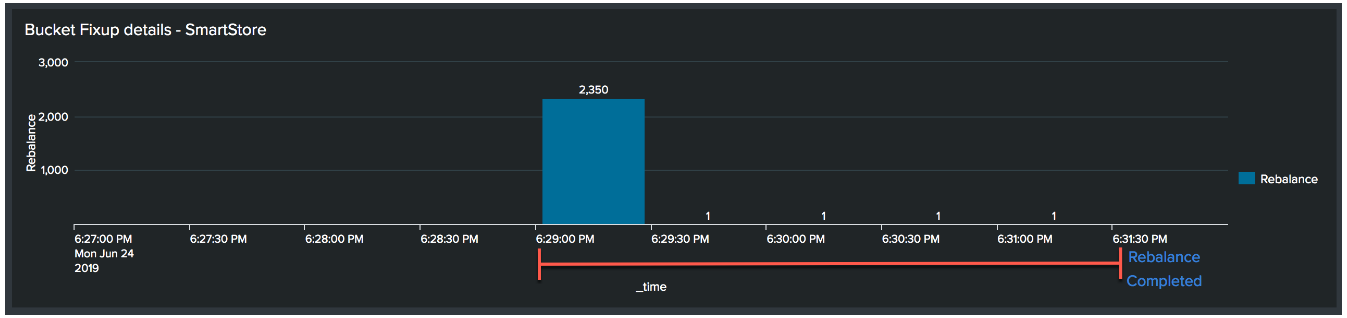

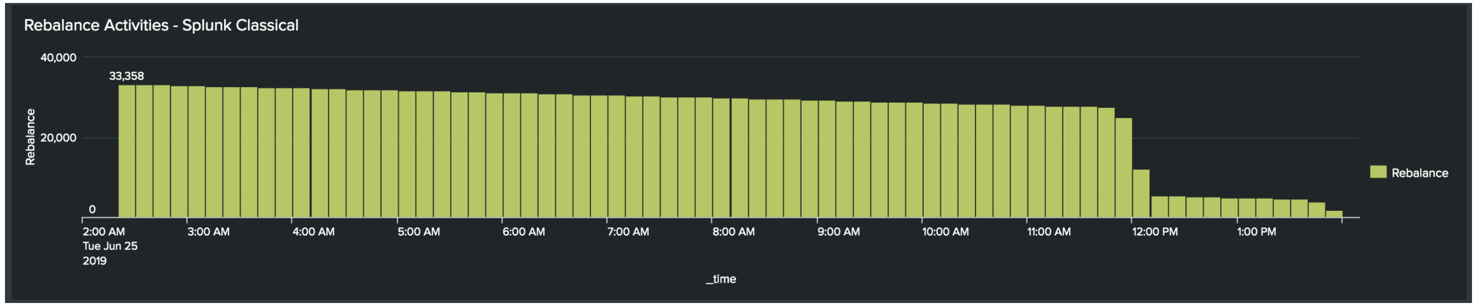

Indexer Node Failure

What happens when an indexer node goes down?

- Cluster Master coordinates bucket fixing to return the cluster to a complete state where the cluster has
  - One primary copy of each bucket
  - A full set of searchable copies for each bucket matching SF
  - A full set of copies for each bucket matching RF

What activities does Bucket-fixing involves when a node goes down?

- Reinstate for any primary copies on the surviving peer nodes (Quick)
- Convert non-searchable bucket copies to searchable on other peer nodes (Time consuming activity)
- Replace all bucket copies of the downed node by streaming a copy of each bucket between the surviving nodes (Time consuming)

Indexer Node Failure

SMARTSTORE BEHAVIOR

What happens when an indexer node goes down in SmartStore?

- Cluster Master initiates the fix-up operation
  - Peers that has a copy of buckets are instructed to copy them to other peers until Replication Factor/Search Factor is met
  - Hot buckets are copied in full
  - For Cached bucket, only the metadata is pushed
    - Peer nodes do not need the full contents of cached buckets, as it can be fetched from the remote object store

Operational Efficiency Test

INDEXER NODE FAILURE/OFFLINE

Offline the indexer node with enforce-counts ensures all data is replicated before the node is shutdown.

- 79K buckets (30 TB) on 10 indexers
- Removed an indexer node
- In Splunk Classical, this can take hours or days
- Completes under 9 minutes under SmartStore

If number of peer nodes down >= RF Indexer cluster does not lose any SmartStore warm buckets as all Warm buckets reside on the remote store

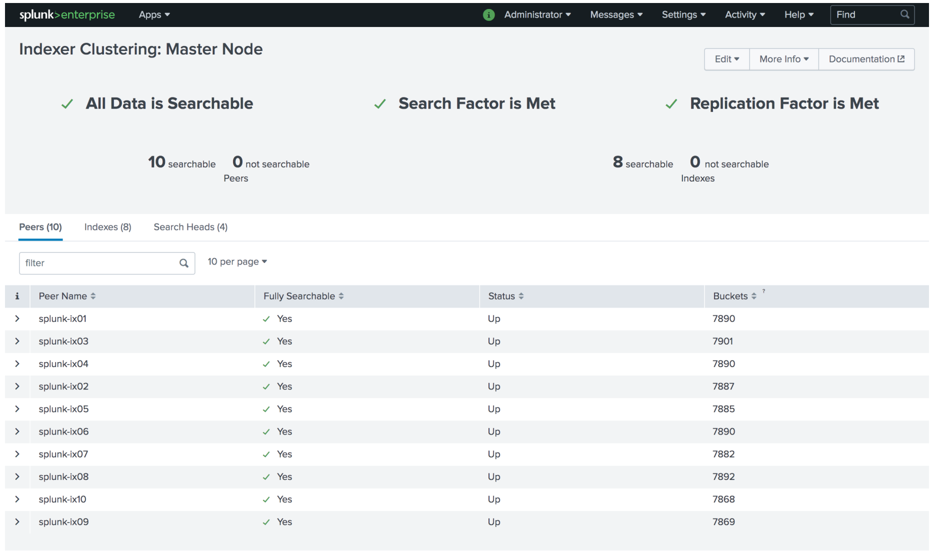

Indexer Clustering: Master Node

✓ All Data is Searchable ✓ Search Factor is Met ✓ Replication Factor is Met

Peers: 10 searchable 0 not searchable

Indexes: 8 searchable 0 not searchable

Peers (10) Indexes (8) Search Heads (4)

filter 10 per page ▾

i	Peer Name	Fully Searchable	Status	Buckets
›	splunk-ix01	✓ Yes	Up	7890
›	splunk-ix03	✓ Yes	Up	7901
›	splunk-ix04	✓ Yes	Up	7890
›	splunk-ix02	✓ Yes	Up	7887
›	splunk-ix05	✓ Yes	Up	7885
›	splunk-ix06	✓ Yes	Up	7890
›	splunk-ix07	✓ Yes	Up	7882
›	splunk-ix08	✓ Yes	Up	7892
›	splunk-ix10	✓ Yes	Up	7868
›	splunk-ix09	✓ Yes	Up	7869

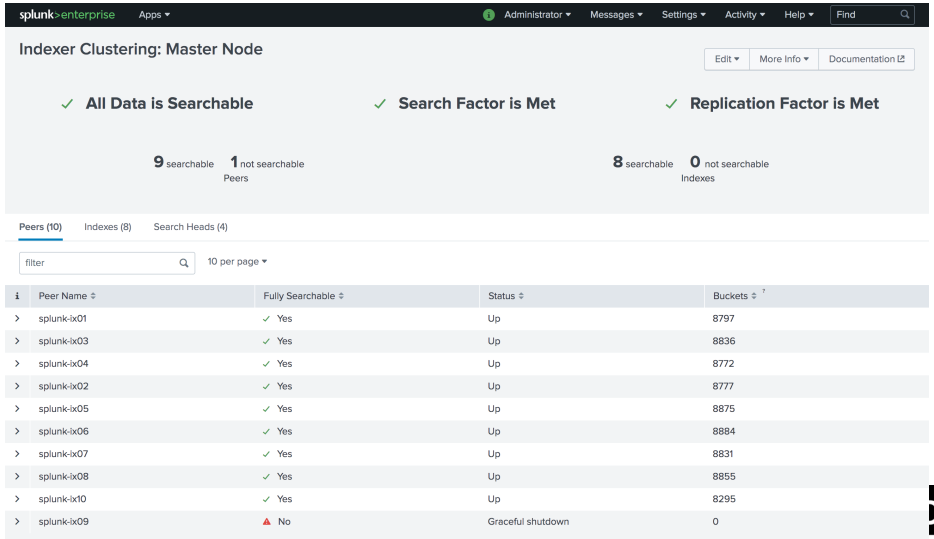

Indexer Clustering: Master Node

✓ All Data is Searchable ✓ Search Factor is Met ✓ Replication Factor is Met

Peers: 9 searchable 1 not searchable

Indexes: 8 searchable 0 not searchable

Peers (10) Indexes (8) Search Heads (4)

filter 10 per page ▾

i	Peer Name	Fully Searchable	Status	Buckets
›	splunk-ix01	✓ Yes	Up	8797
›	splunk-ix03	✓ Yes	Up	8836
›	splunk-ix04	✓ Yes	Up	8772
›	splunk-ix02	✓ Yes	Up	8777
›	splunk-ix05	✓ Yes	Up	8875
›	splunk-ix06	✓ Yes	Up	8884
›	splunk-ix07	✓ Yes	Up	8831
›	splunk-ix08	✓ Yes	Up	8855
›	splunk-ix10	✓ Yes	Up	8295
›	splunk-ix09	⚠ No	Graceful shutdown	0

Operational Efficiency Test

INDEXER NODE FAILURE/OFFLINE

Indexer node offline comparison of SmartStore against Classical Splunk

- 79K buckets (30 TB) on 10 indexers
- 9 minutes under SmartStore
- 2 hours and 35 minutes under Classical

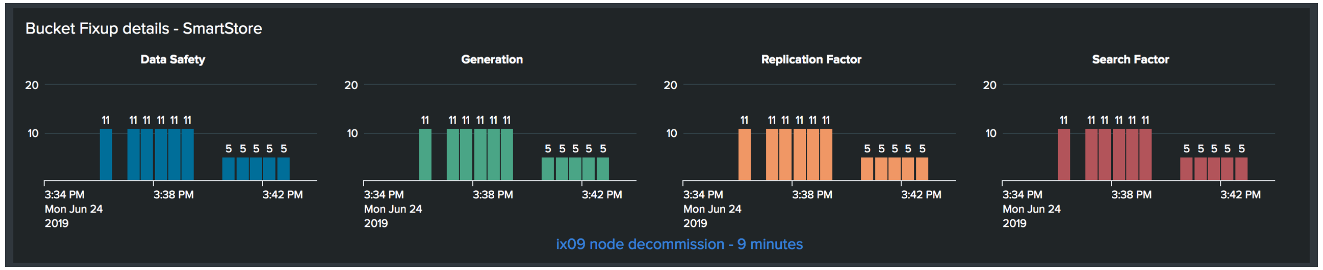
Bucket Fixup details - SmartStore

ix09 node decommission - 9 minutes

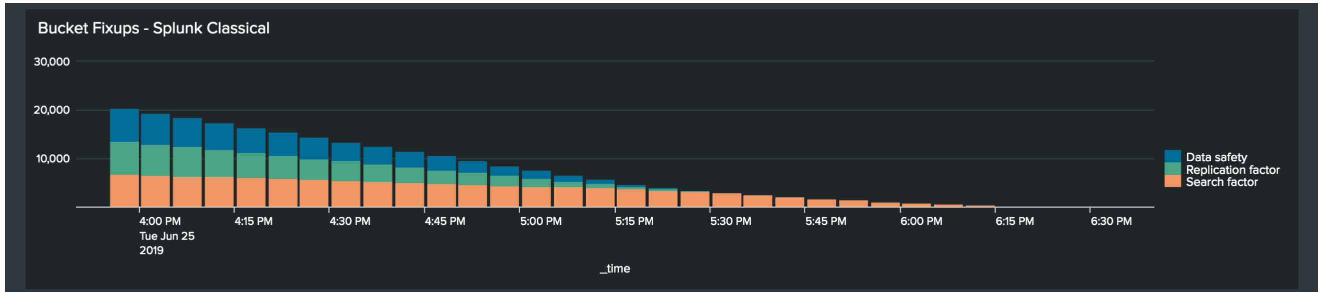
Bucket Fixups - Splunk Classical

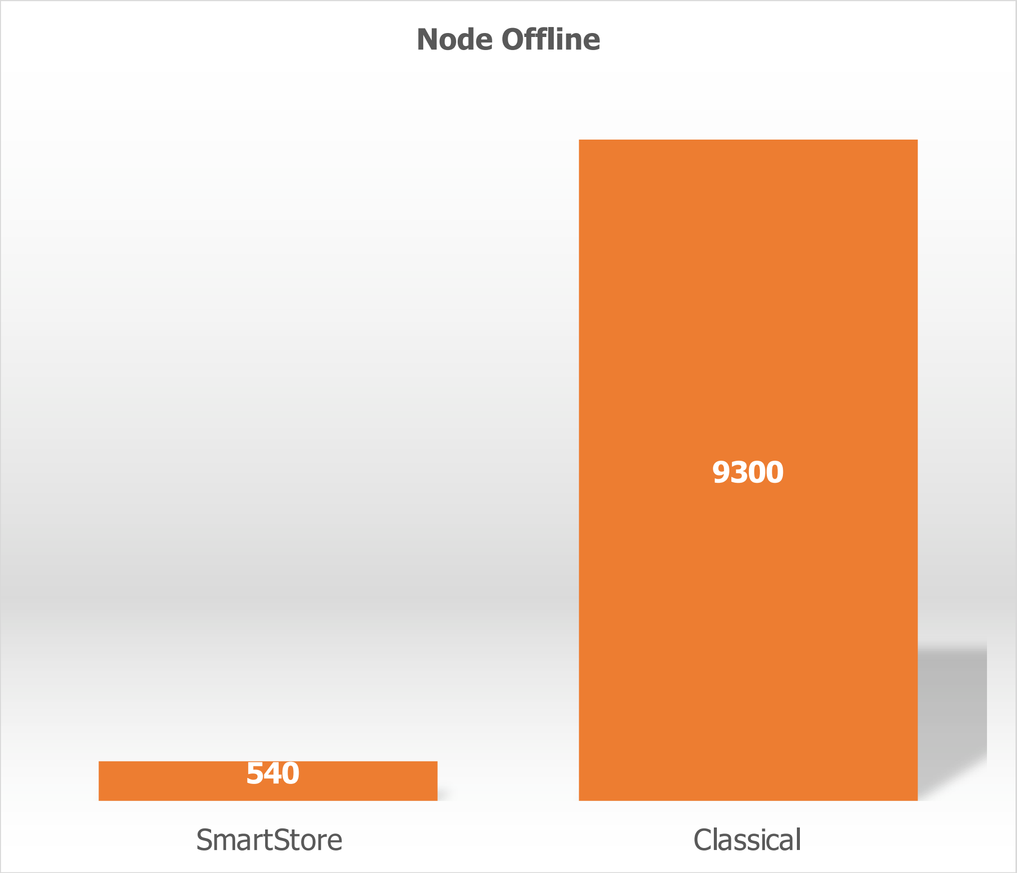
Node Offline

SPLUNK SMARTSTORE on FLASHBLADE™

BENEFITS

Operational Efficiency

- Simplified Indexer maintenance without impacting data integrity
- Increased HA recovery
- Dynamic Cluster scaling by adding compute and storage independently
- Upgrade/replace indexer with simple bootstrap from the object store
- FB provides compression and encryption

Reduction in TCO

- Deploy indexers based on compute requirements
- Lower TCO with reduced indexer infrastructure
- Reduction in Indexer servers and storage

Scalability & High Availability

- High data availability with remote storage tier on FlashBlade
- Performance at scale with cached active dataset
- Architectured for massive scale

For more information

Visit us at **Booth #114** in
source=*Pavilion

splunk> .conf19

.conf19

splunk>

Thank You!

Go to the .conf19 mobile app to

RATE THIS SESSION