Smart Store Deep Dive

Da Xu
Senior Engineer Manager | Splunk

Bill Ern
Splunk Product Owner | Lockheed Martin

splunk> .conf19

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results may differ materially. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

Agenda

1. High Level Overview
2. Buckets
   - Hot/Warm/Cold
3. Searching
4. SmartStore at Lockheed Martin
   - Architecture
   - Learnings

Overview

splunk> .conf19

High Level Overview

SmartStore brings about the separation of Storage and Compute

Before SmartStore: the filesystem was the storage tier

- Keep buckets on disk for search
- Keep buckets on disk for retention

After SmartStore: the filesystem is for compute!

- Keep buckets on disk for search
- ~~Keep buckets on disk for retention~~

Architecture

Components

Classic Architecture

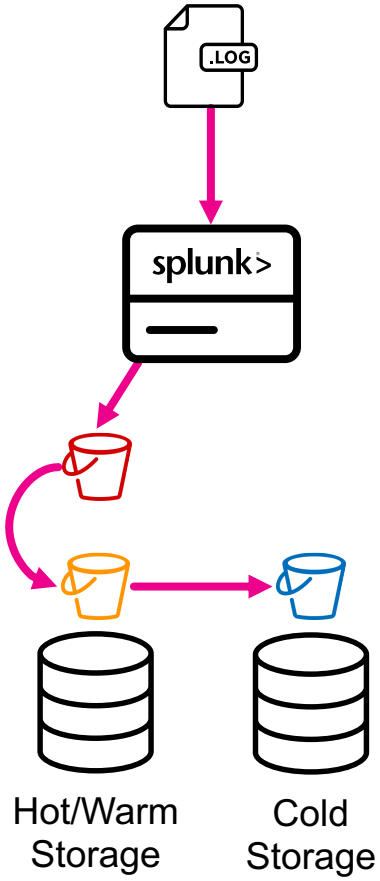

S2 Architecture

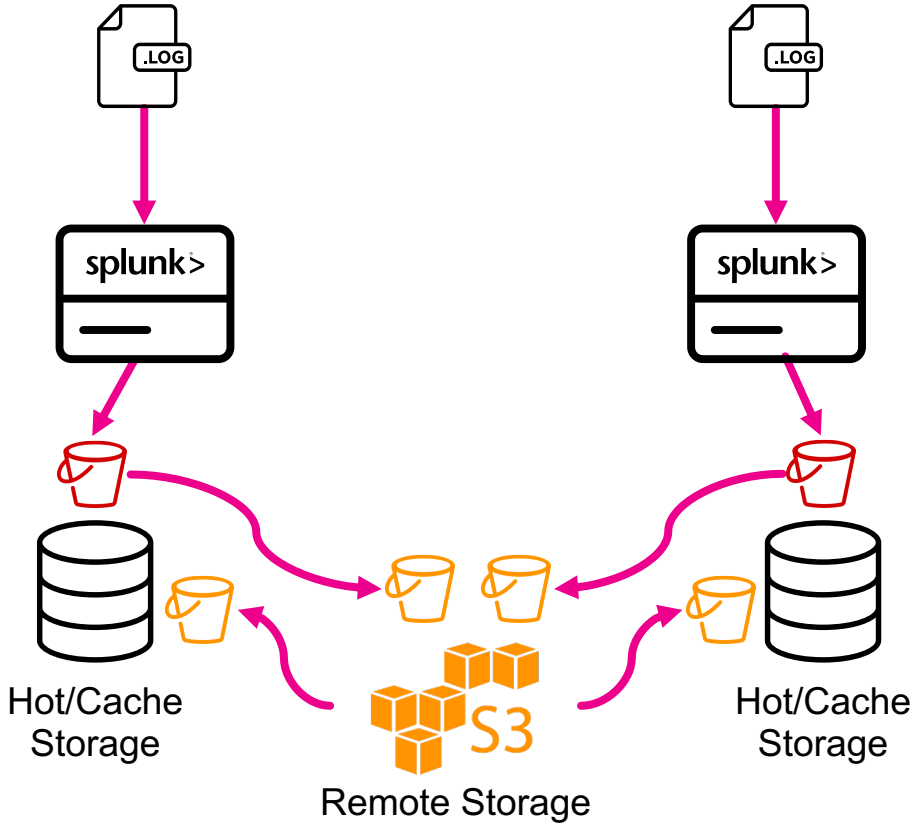

Key Advantages

Our Storage Tier is no longer tied to hardware

- Separation of storage and compute
- Indexer failures is no longer tied to storage failure

Local Storage is now simply a Search-Cache

- No longer need to size local storage to hold long-term retention
  - Local storage used to store - 90 days, 6 months, 3 years, etc
- Just need enough local storage for search
  - Most search is just for 1 day or 7 days

Key Advantages

Classic Architecture

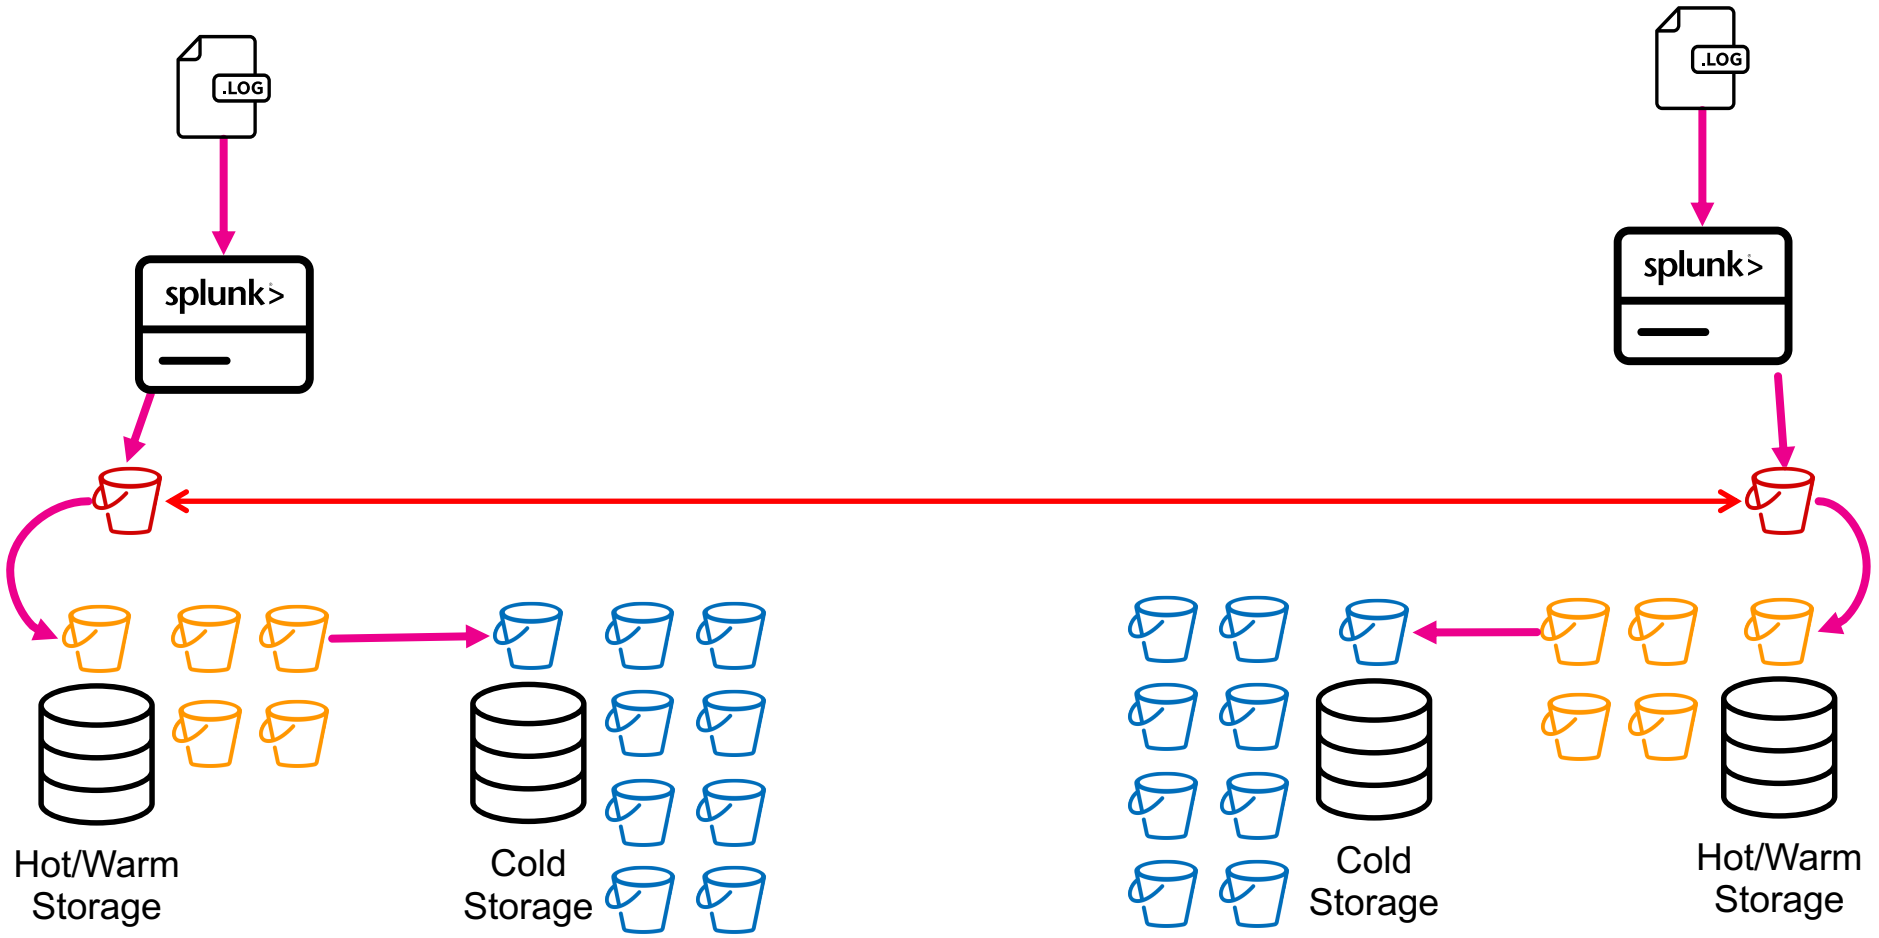

Key Advantages

Classic Architecture

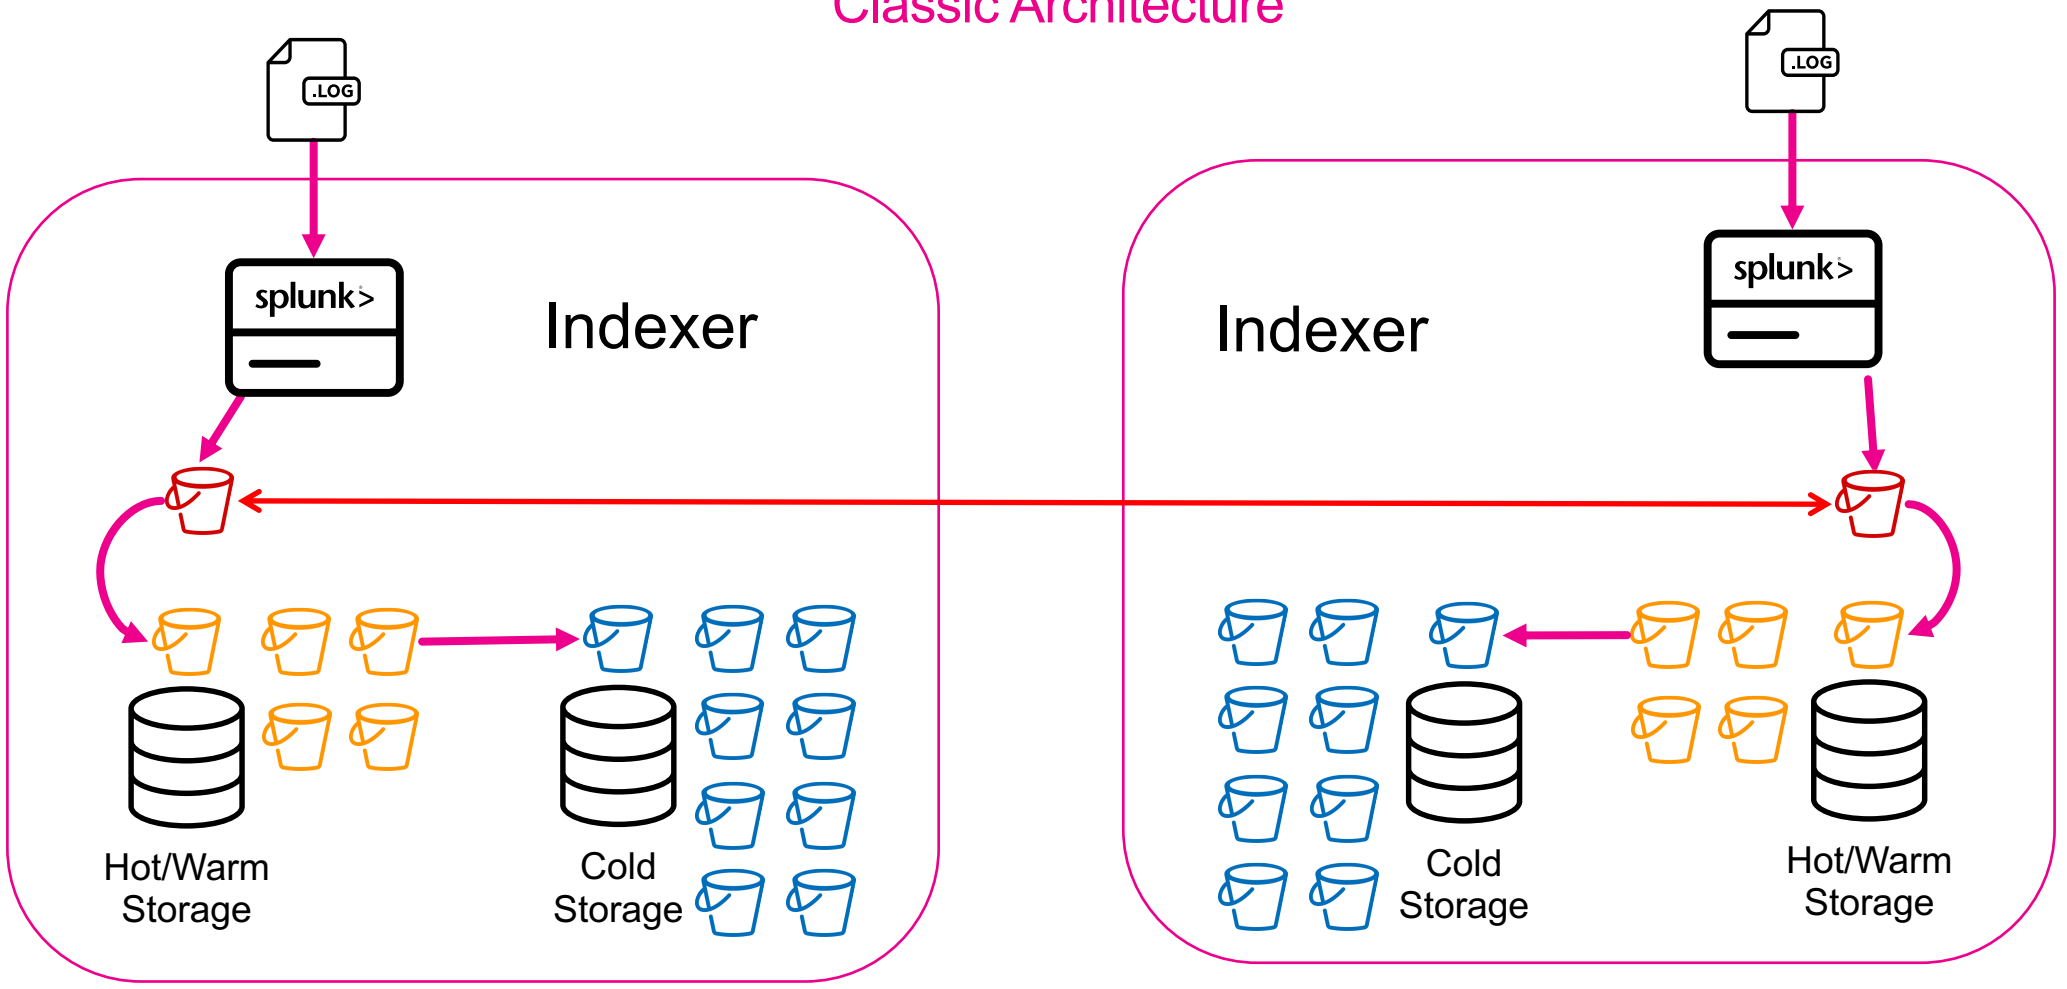

Key Advantages

Classic Architecture

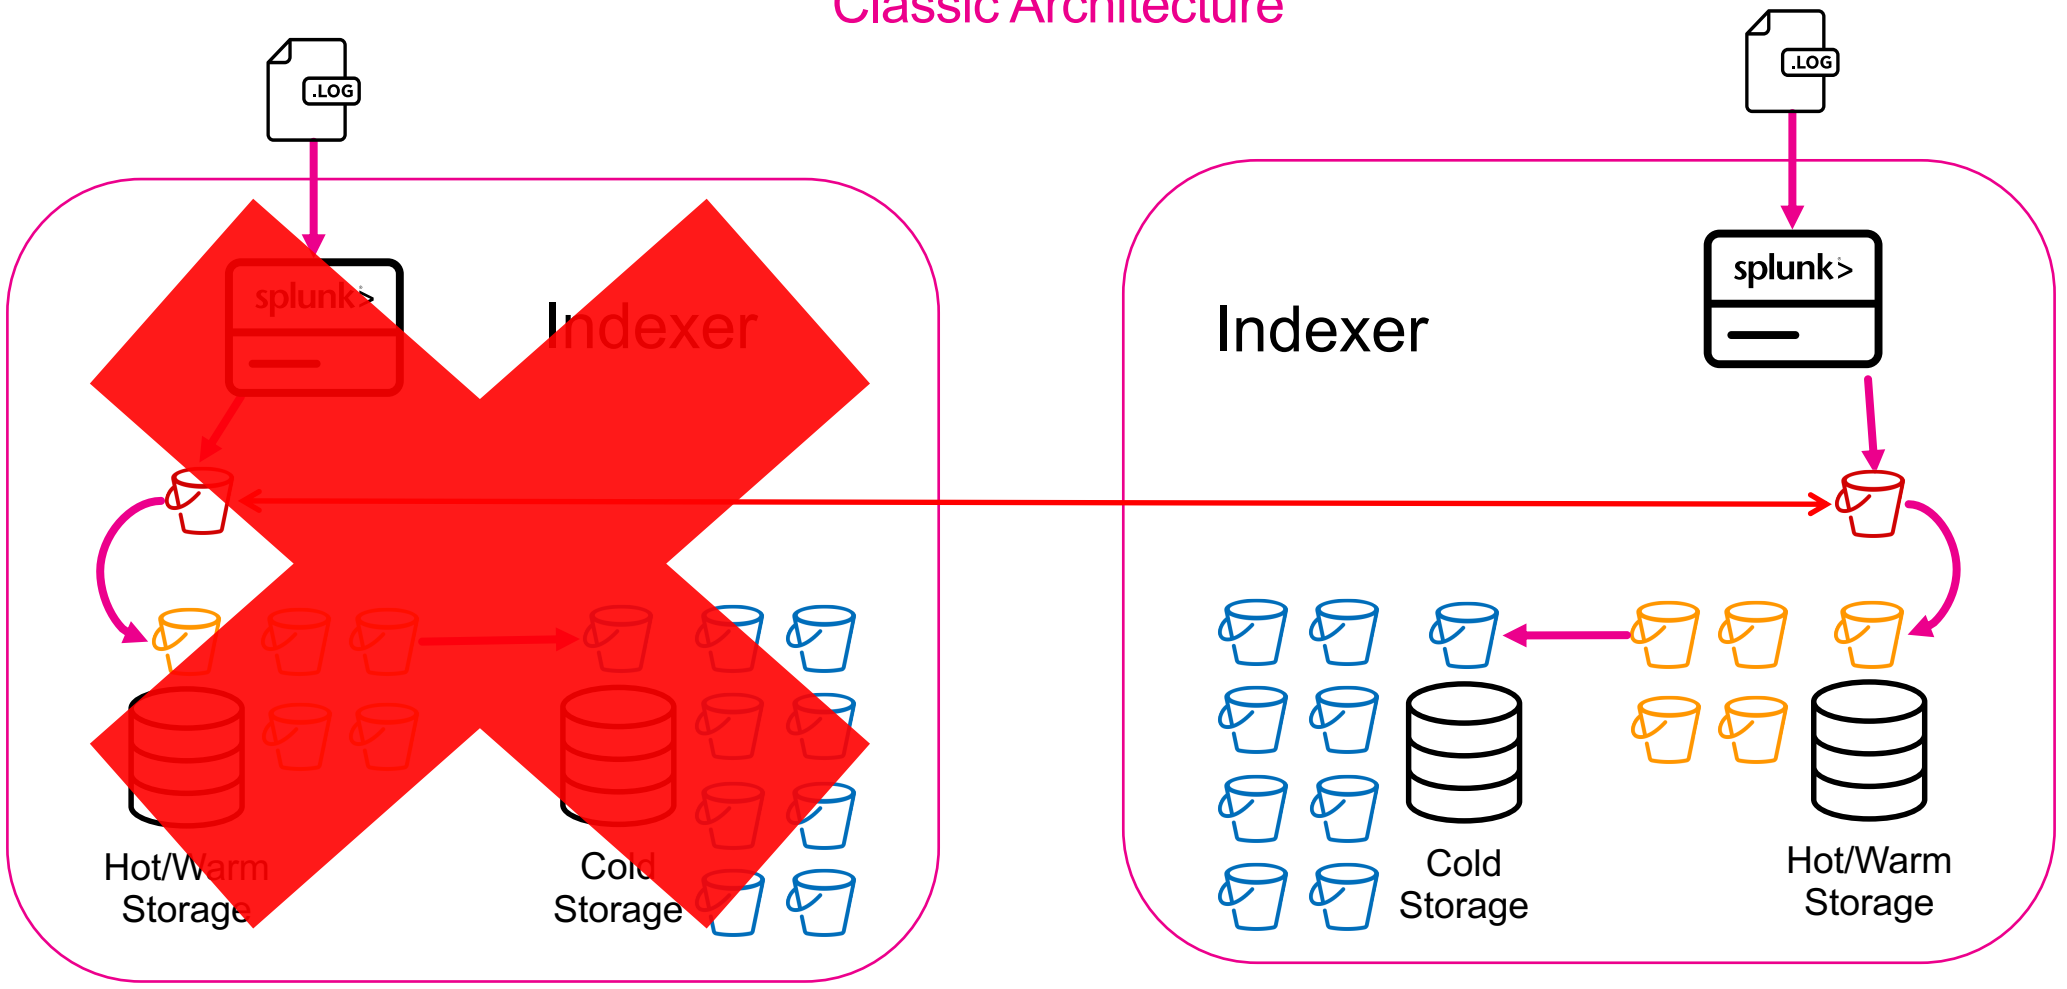

Key Advantages

Smart Store Architecture

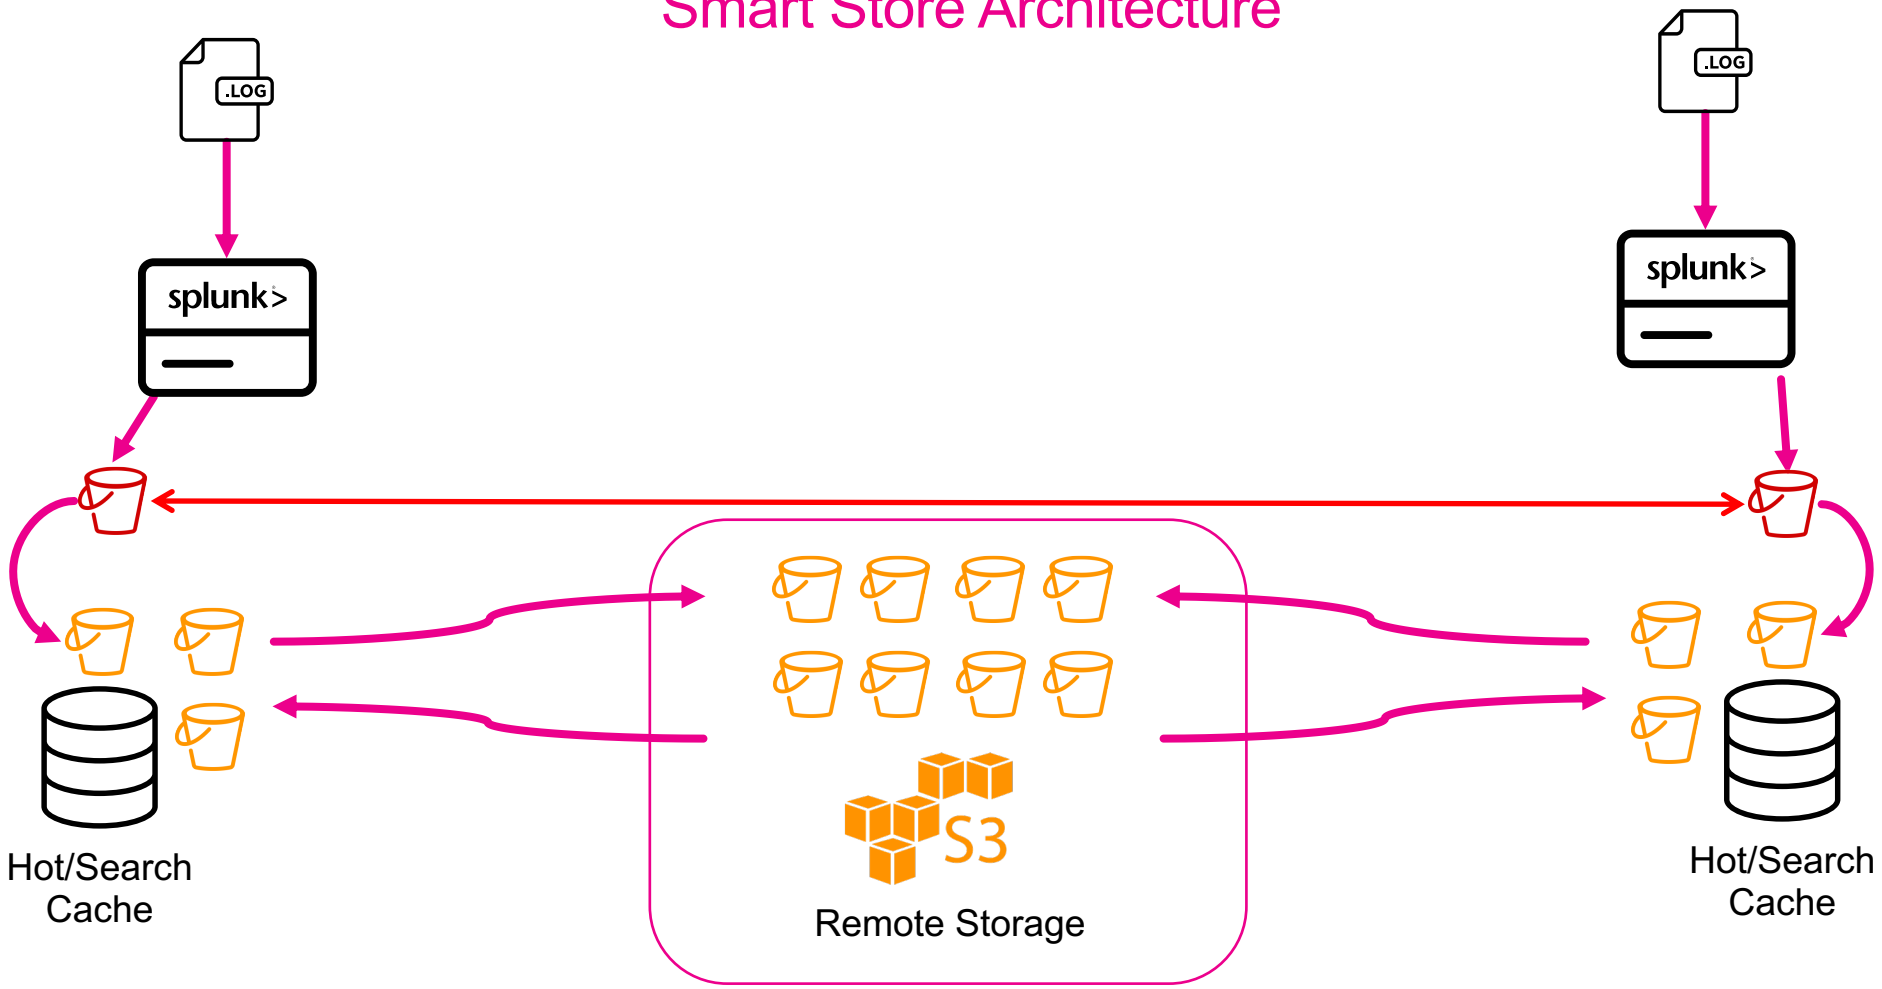

Key Advantages

Smart Store Architecture

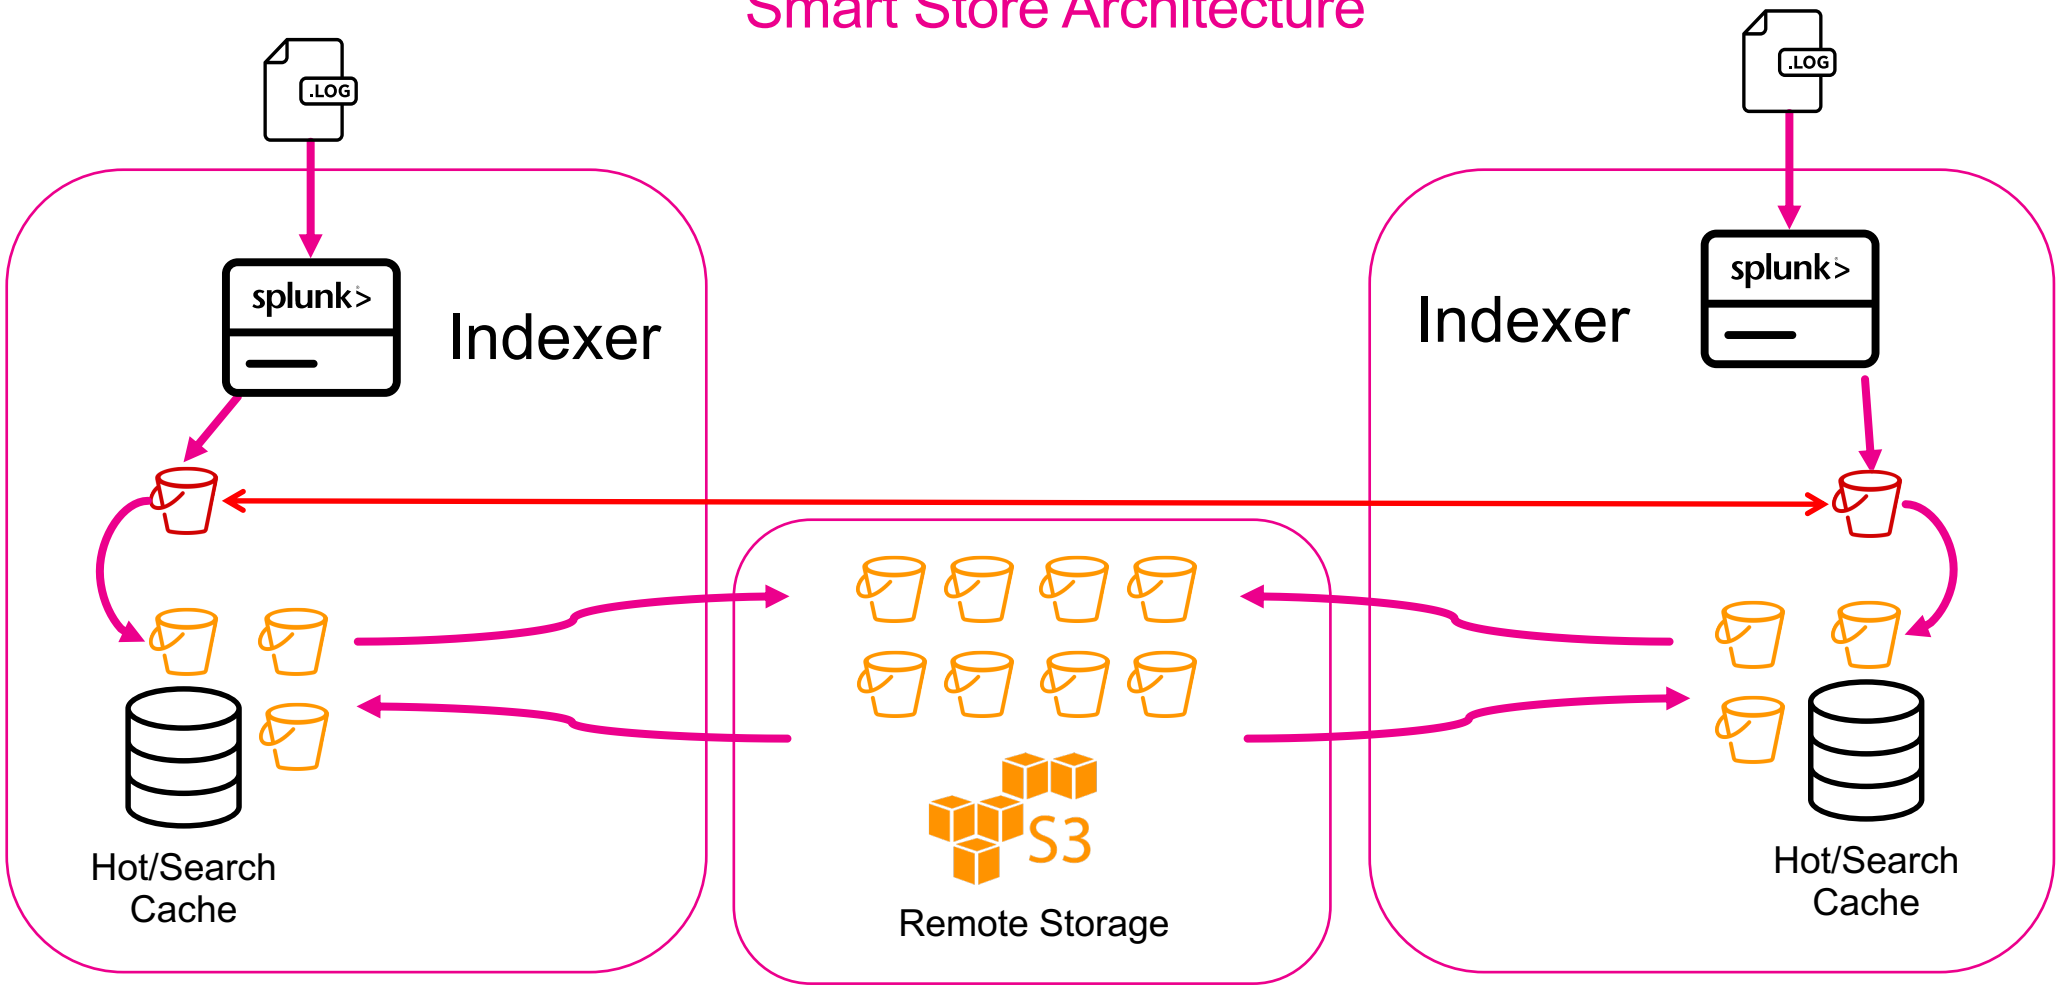

Key Advantages

Smart Store Architecture

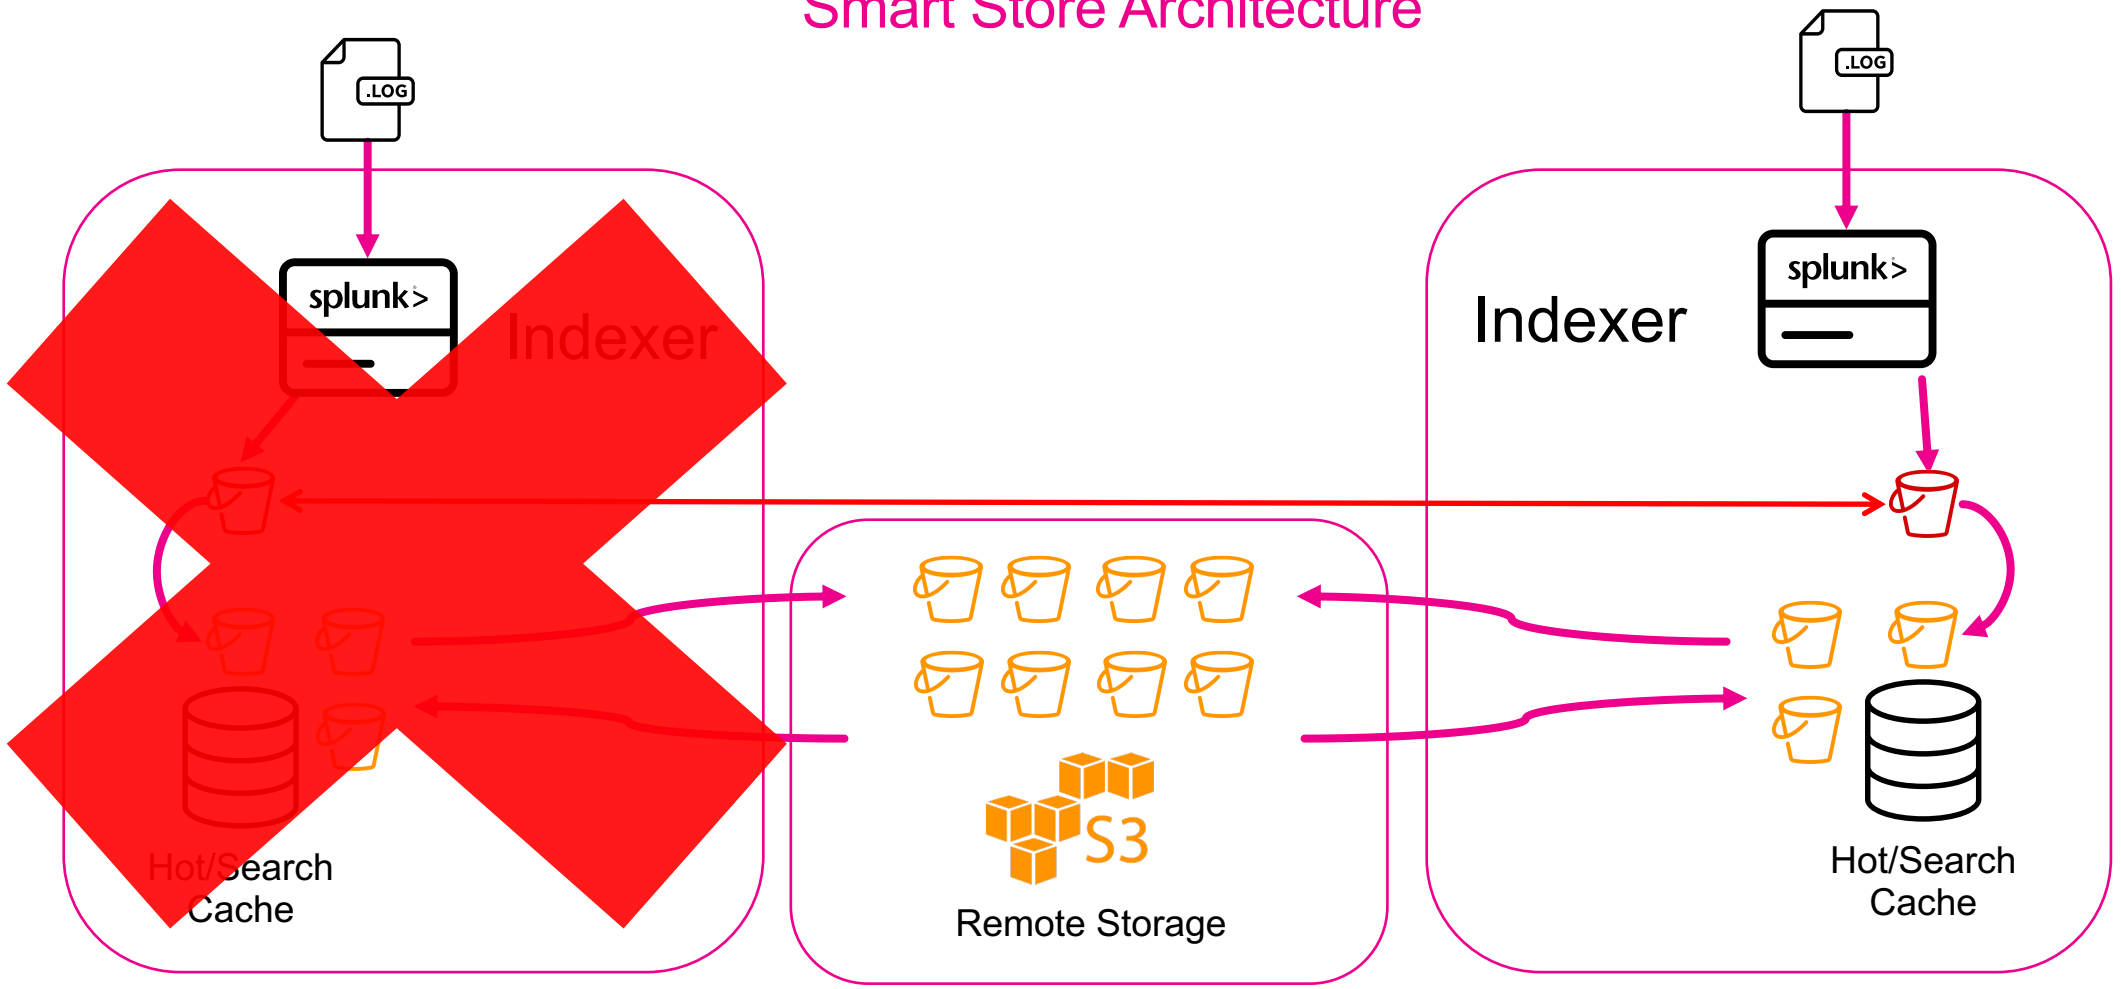

Deep Dive – Buckets

splunk> .conf19

Buckets Overview - Hot

Hot Buckets are the same as non SmartStore

- When using Indexer Clustering, hot buckets are replicated just like normal

There is no information about hot buckets on remote storage

Search works the same as non SmartStore

Buckets Overview - Hot

When the bucket transitions to warm, that's when it will upload to remote storage

- In the clustering case, only the source will initially upload
- The targets will start a timer, and upload later ONLY if the source failed to upload within the timer expiration

After uploading the bucket to remote storage, we also keep the bucket on local disk.

- It's a recent hot bucket – and will likely be searched!
- In clustering, we keep the source, but the targets will be evicted to free up space

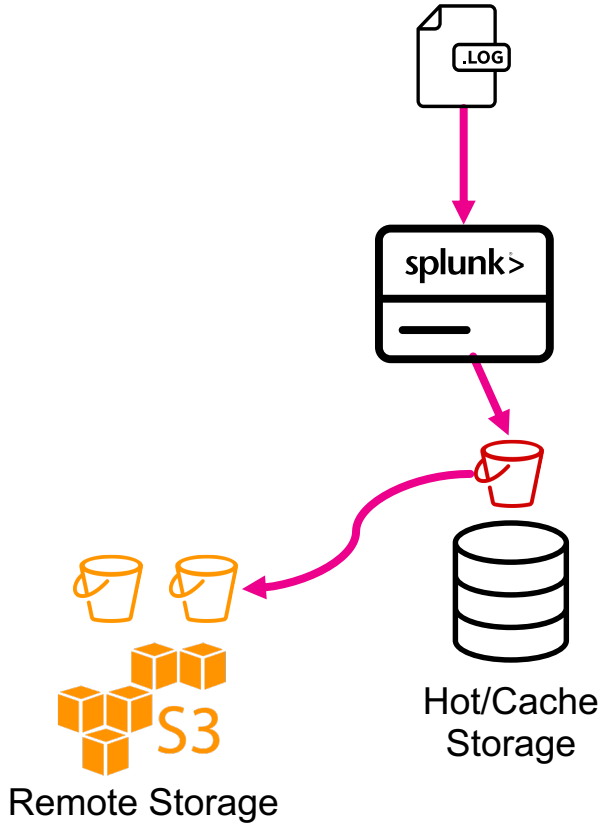

Buckets Overview – Warm/Cold

Warm (and cold) buckets **may or may not** be fully existing!

- The folder is there, and they are in memory, but the actual contents may be missing!
- If the content is missing, these are considered evicted buckets
- When a search comes in against an evicted bucket, Splunk will download the bucket to serve the search – more on this process later

Buckets Overview – Warm/Cold

```
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/ | head -n 400 | tail
drwx------    2 splunk splunk   4096 Sep 19 15:27 db_1567409839_1567400823_5827_508A66A1-4E47-497D-A038-E8D376644D72
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566884199_1566868870_5851_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566522183_1566522177_5858_83C9AB49-D062-4370-89A9-52FF1539D79D
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566912744_1566902296_5860_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:28 db_1567096738_1567058261_5876_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567122755_1567120695_5881_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567124087_1567121687_5884_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567130519_1567126875_5892_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147947_1567131875_5897_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
total 464
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 .
drwx------ 3511 splunk splunk 466944 Oct 21 06:43 ..
splunk@idx-i-0f8facba5ef9ac497:~$
```

Getting Data In

Clustered Deployments

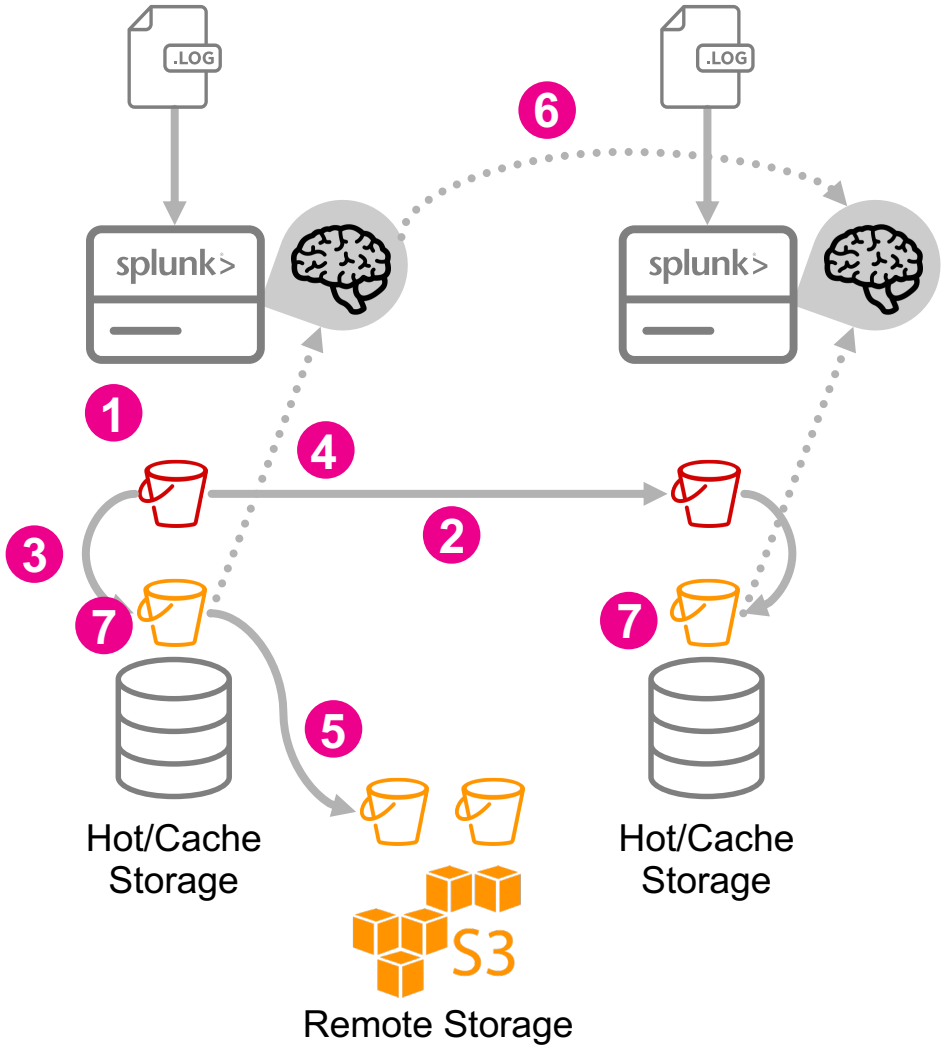

1. Data arrives and is written to a Hot bucket
2. Hot bucket streams to cluster peer(s) according to RF
3. Replication completes and the buckets roll to warm
4. Buckets are registered with their cache managers
5. Cache manager on source peer uploads the bucket to the remote store
6. Source peer notifies replication peers that the bucket was uploaded successfully
7. Cached copies remains on the peers until evicted by the local cache manager

Deep Dive – Search

splunk> .conf19

Searching

Hot Buckets are the same as non SmartStore

Warm buckets behavior can be different depending on the bucket state:

Bucket State	Description
hot	Same as non SmartStore
warm – all files local	Same as non SmartStore
warm – some files local	Might need to download more: If local files are enough for search, then we don't! Ex: Bloomfilter filters out this bucket ➔ nothing more is needed Tstats search and tsidx is local ➔ nothing more is needed
warm – no files local	Start downloading files for search

Searching with S2

Hot Buckets

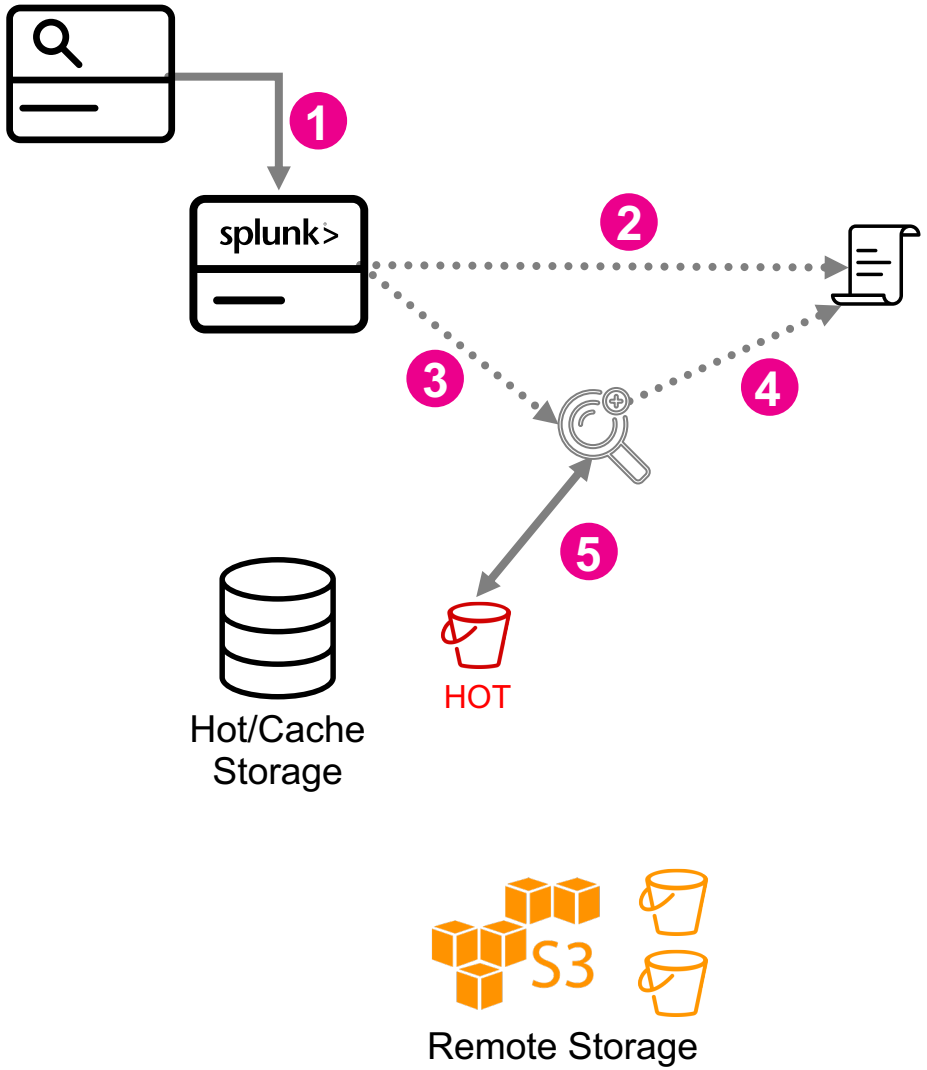

1. Search request is received
2. Indexer generates a list of relevant buckets to be searched
3. Search process is spawned
4. Spawned process reads the bucket list
5. Hot buckets are searched in the same manner as “classic” search

Searching with S2

Cached Buckets

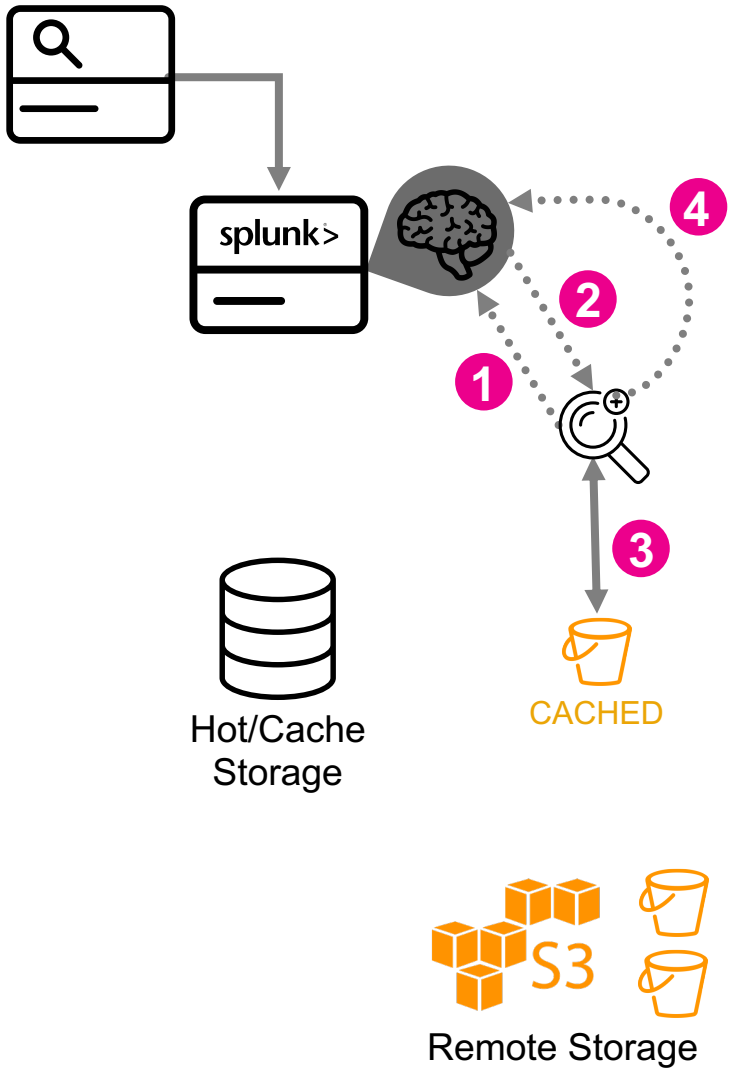

1. Search process “opens” the bucket with the Cache manager
2. Cache manager tells the search process that the bucket is local and available for search
3. Search process searches the bucket
4. Search process ”closes” the bucket with the cache manager

Searching with S2

Remote Buckets

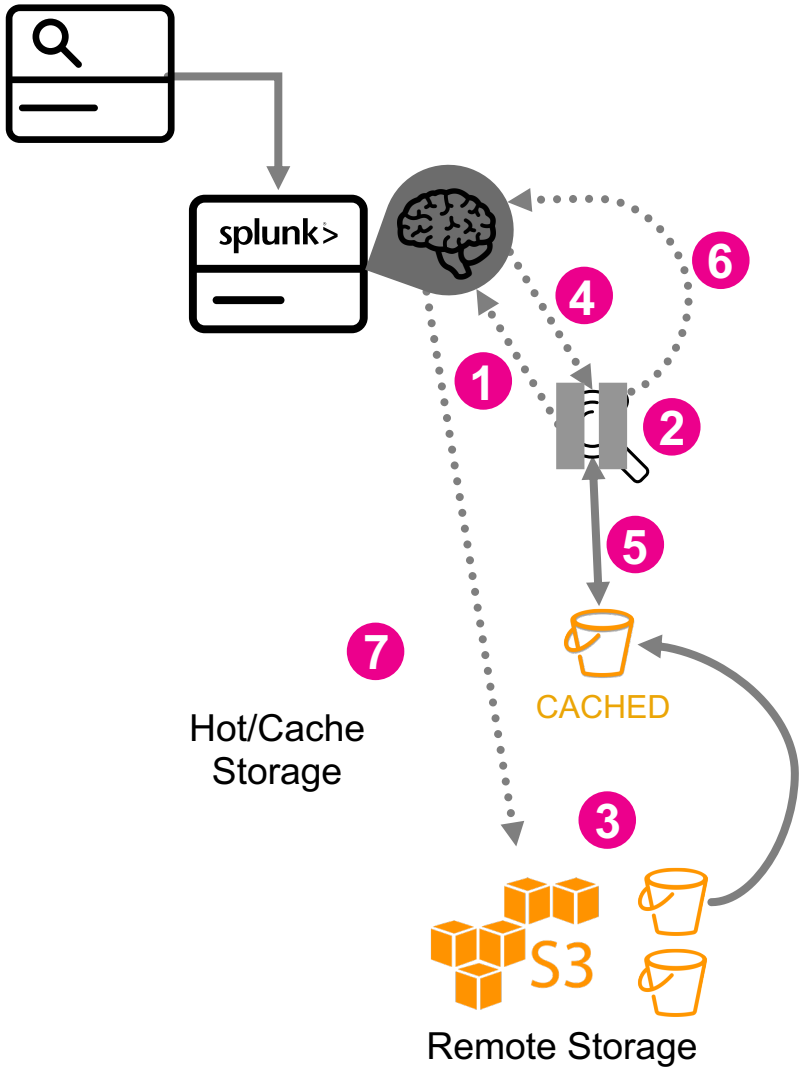

1. Search process “opens” the bucket with the Cache manager, but it isn’t in cache
2. Search process waits
3. Cache manager fetches the bucket from the remote store
4. Cache manager tells the search process that the bucket is local and available for search
5. Search process searches the bucket
6. Search process ”closes” the bucket with the cache manager
7. Bucket remains in cache until evicted by the cache manager

Cache Manager

Localizing Data

Cache manager offers the ability to fetch specific bucket files

Ex: bloomfilter, TSIDX, metadata, journal

Some search commands only need specific files from the bucket

- Don’t need the raw data
  - Ex: metadata, tstats
- Don’t need any bucket content
  - Ex: eventcount, dbinspect

Lookahead

- Cache manager will attempt to pre-fetch buckets needed for a search
  - Heuristic will adjust itself based upon the thruput from the remote store

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/ | head -n 400 | tail
drwx------    2 splunk splunk   4096 Sep 19 15:27 db_1567409839_1567400823_5827_508A66A1-4E47-497D-A038-E8D376644D72
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566884199_1566868870_5851_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566522183_1566522177_5858_83C9AB49-D062-4370-89A9-52FF1539D79D
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566912744_1566902296_5860_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:28 db_1567096738_1567058261_5876_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567122755_1567120695_5881_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567124087_1567121687_5884_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567130519_1567126875_5892_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147947_1567131875_5897_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
total 464
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 .
drwx------ 3511 splunk splunk 466944 Oct 21 06:43 ..
splunk@idx-i-0f8facba5ef9ac497:~$
```

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/ | head -n 400 | tail
drwx------    2 splunk splunk   4096 Sep 19 15:27 db_1567409839_1567400823_5827_508A66A1-4E47-497D-A038-E8D376644D72
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566884199_1566868870_5851_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566522183_1566522177_5858_83C9AB49-D062-4370-89A9-52FF1539D79D
drwx------    2 splunk splunk   4096 Sep 19 15:28 db_1566912744_1566902296_5860_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:28 db_1567096738_1567058261_5876_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567122755_1567120695_5881_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567124087_1567121687_5884_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567130519_1567126875_5892_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147947_1567131875_5897_00D246A1-32FE-4C7F-8A96-548D40C50073
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
total 464
drwx--x---    2 splunk splunk   4096 Sep 19 15:29 .
drwx------ 3511 splunk splunk 466944 Oct 21 06:43 ..
splunk@idx-i-0f8facba5ef9ac497:~$
```

```
splunk@idx-i-0f8facba5ef9ac497:~$ bin/splunk search "| tstats count where index=_internal earliest=1567146441 latest=1567147265"
```

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ bin/splunk search "| tstats count where index=_internal earliest=1567146441 latest=1567147265"

INFO: Your timerange was substituted based on your 'where' clause
count
------
715445
```

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ bin/splunk search "| tstats count where index=_internal earliest=1567146441 latest=1567147265"
```

```
INFO: Your timerange was substituted based on your 'where' clause
count
------
715445
```

```
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073
total 1072
-rw-------    1 splunk splunk     88 Oct 21 06:44 splunk-autogen-params.dat
-rw-------    1 splunk splunk      8 Oct 21 06:44 .rawSize
-rw-------    1 splunk splunk     75 Oct 21 06:44 bucket_info.csv
-rw-------    1 splunk splunk      6 Oct 21 06:44 .sizeManifest4.1
-rw-------    1 splunk splunk 600072 Oct 21 06:44 1567147265-1567146441-16742391476385632200.tsidx
-rw-------    1 splunk splunk     49 Oct 21 06:44 cachemanager_local.json
drwx--x---    2 splunk splunk   4096 Oct 21 06:44 .
drwx------ 3511 splunk splunk 466944 Oct 21 06:45 ..
splunk@idx-i-0f8facba5ef9ac497:~$
```

Example - searching

- cachemanager_local.json
  - **contains the local set of files for a particular bucket**

```
splunk@idx-i-0f8facba5ef9ac497:~$ ls -lh var/lib/splunk/_internaldb/db/db_1564422298_1564421583_4970_1A5C7428-EAC6-4B32-BDDC-0CEABD7BA4A3
total 8.1M
-rw------- 1 splunk splunk 8.0M Oct 23 07:36 1564422298-1564421583-9123203473889183478.tsidx
-rw------- 1 splunk splunk   75 Oct 23 07:36 bucket_info.csv
-rw------- 1 splunk splunk   49 Oct 23 07:36 cachemanager_local.json
-rw------- 1 splunk splunk   89 Oct 23 07:36 splunk-autogen-params.dat
splunk@idx-i-0f8facba5ef9ac497:~$ awk 1 var/lib/splunk/_internaldb/db/db_1564422298_1564421583_4970_1A5C7428-EAC6-4B32-BDDC-0CEABD7BA4A3/cachemanager_local.json
{"file_types":["dma_metadata","tsidx","deletes"]}
```

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ bin/splunk search "index=_internal earliest=1567146441 latest=1567147265 | stats count"
```

Example - searching

```
splunk@idx-i-0f8facba5ef9ac497:~$ bin/splunk search "index=_internal earliest=1567146441 latest=1567147265 | stats count"

INFO: Your timerange was substituted based on your search string
count
------
715445
splunk@idx-i-0f8facba5ef9ac497:~$ ls -latr var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/*
-rw------- 1 splunk splunk 600072 Oct 21 06:44 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/1567147265-1567146
x
-rw------- 1 splunk splunk     88 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/splunk-autogen-par
-rw------- 1 splunk splunk     75 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/bucket_info.csv
-rw------- 1 splunk splunk    955 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/SourceTypes.data
-rw------- 1 splunk splunk   1617 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/Sources.data
-rw------- 1 splunk splunk    302 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/Strings.data
-rw------- 1 splunk splunk    374 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/Hosts.data
-rw------- 1 splunk splunk  19603 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/bloomfilter
-rw------- 1 splunk splunk    138 Oct 21 06:45 var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/cachemanager_local

var/lib/splunk/_internaldb/db/db_1567147265_1567146441_5898_00D246A1-32FE-4C7F-8A96-548D40C50073/rawdata:
total 288
-rw------- 1 splunk splunk    539 Oct 21 06:45 slicesv2.dat
-rw------- 1 splunk splunk     84 Oct 21 06:45 slicemin.dat
-rw------- 1 splunk splunk 275361 Oct 21 06:45 journal.gz
drwx------ 2 splunk splunk   4096 Oct 21 06:45 .
drwx--x--- 3 splunk splunk   4096 Oct 21 06:45 ..
splunk@idx-i-0f8facba5ef9ac497:~$
```

Deep Dive – Eviction

splunk> .conf19

Cache Manager – Eviction Policies

Cached Data

When do we evict?

- Hot buckets are always local
- Warm buckets are not immediately evicted from the cache
  - We look at available storage first, then…
  - Cache manager will attempt to keep buckets that contain events with timestamps from the last 24 hours [hotlist_recency_secs]
- TSIDX and Journals are evicted quicker than other bucket files
  - Cache manager will attempt to keep smaller bucket files for 15 days [hotlist_bloom_filter_recency_hours]

Policy Name	Description
clock	Prefer to evict bucket with the oldest events first, unless it has been accessed recently
lru (default)	Evict the least recently used bucket
random	Randomly evict a bucket
lrult	Evict the bucket with the oldest events first
noevict	Don't evict – This can be used to provide data resiliency instead of indexer clustering

Cache Manager – Eviction Policies

Cached Data

When do we evict?

- As we approach using up all the available cache space.
  - max cache size
  - max volume size
- Clustered target buckets on hot → warm bucket transitions

Eviction

Cached Data

Very similar to CPU Memory Caching!

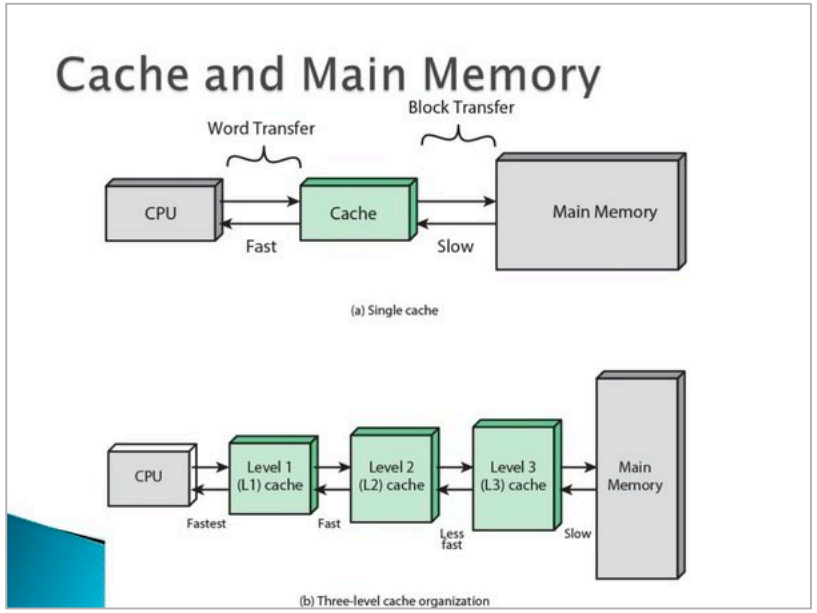

Eviction

Cached Data

Very similar to CPU Memory Caching!

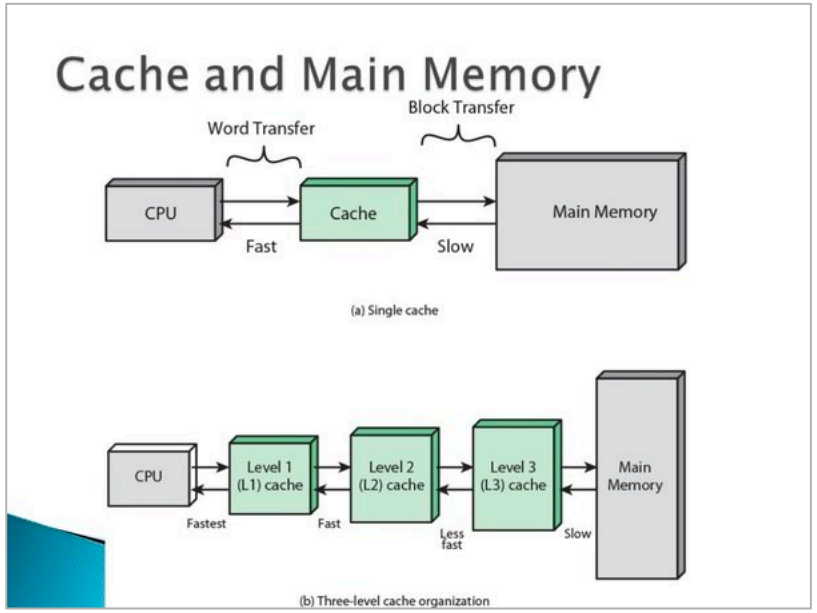

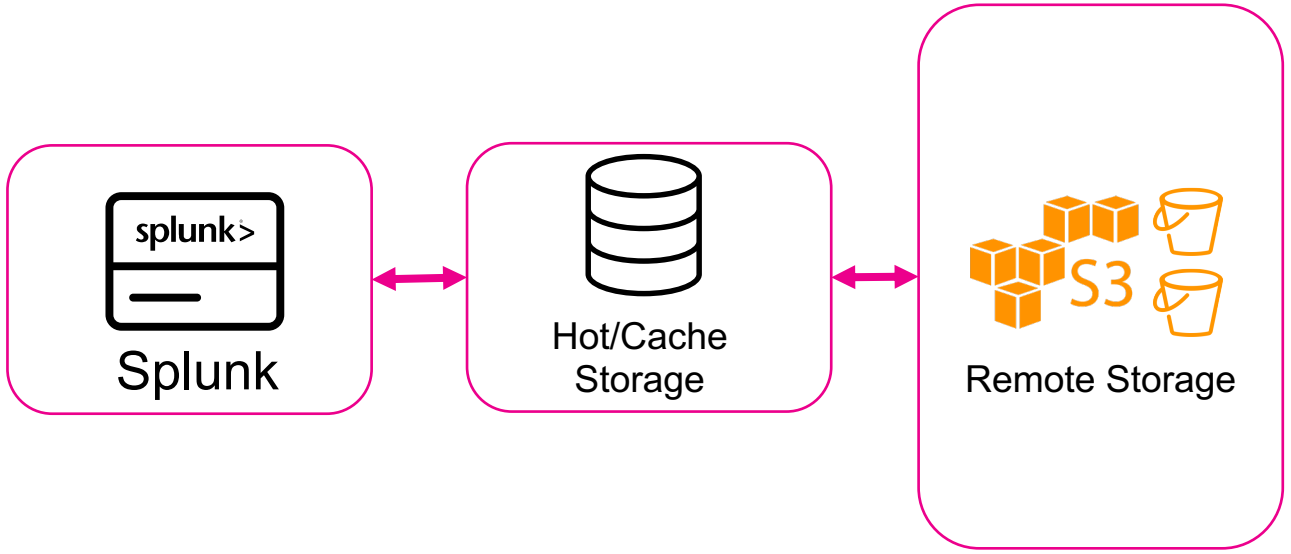

Monitoring Console – Cache Performance

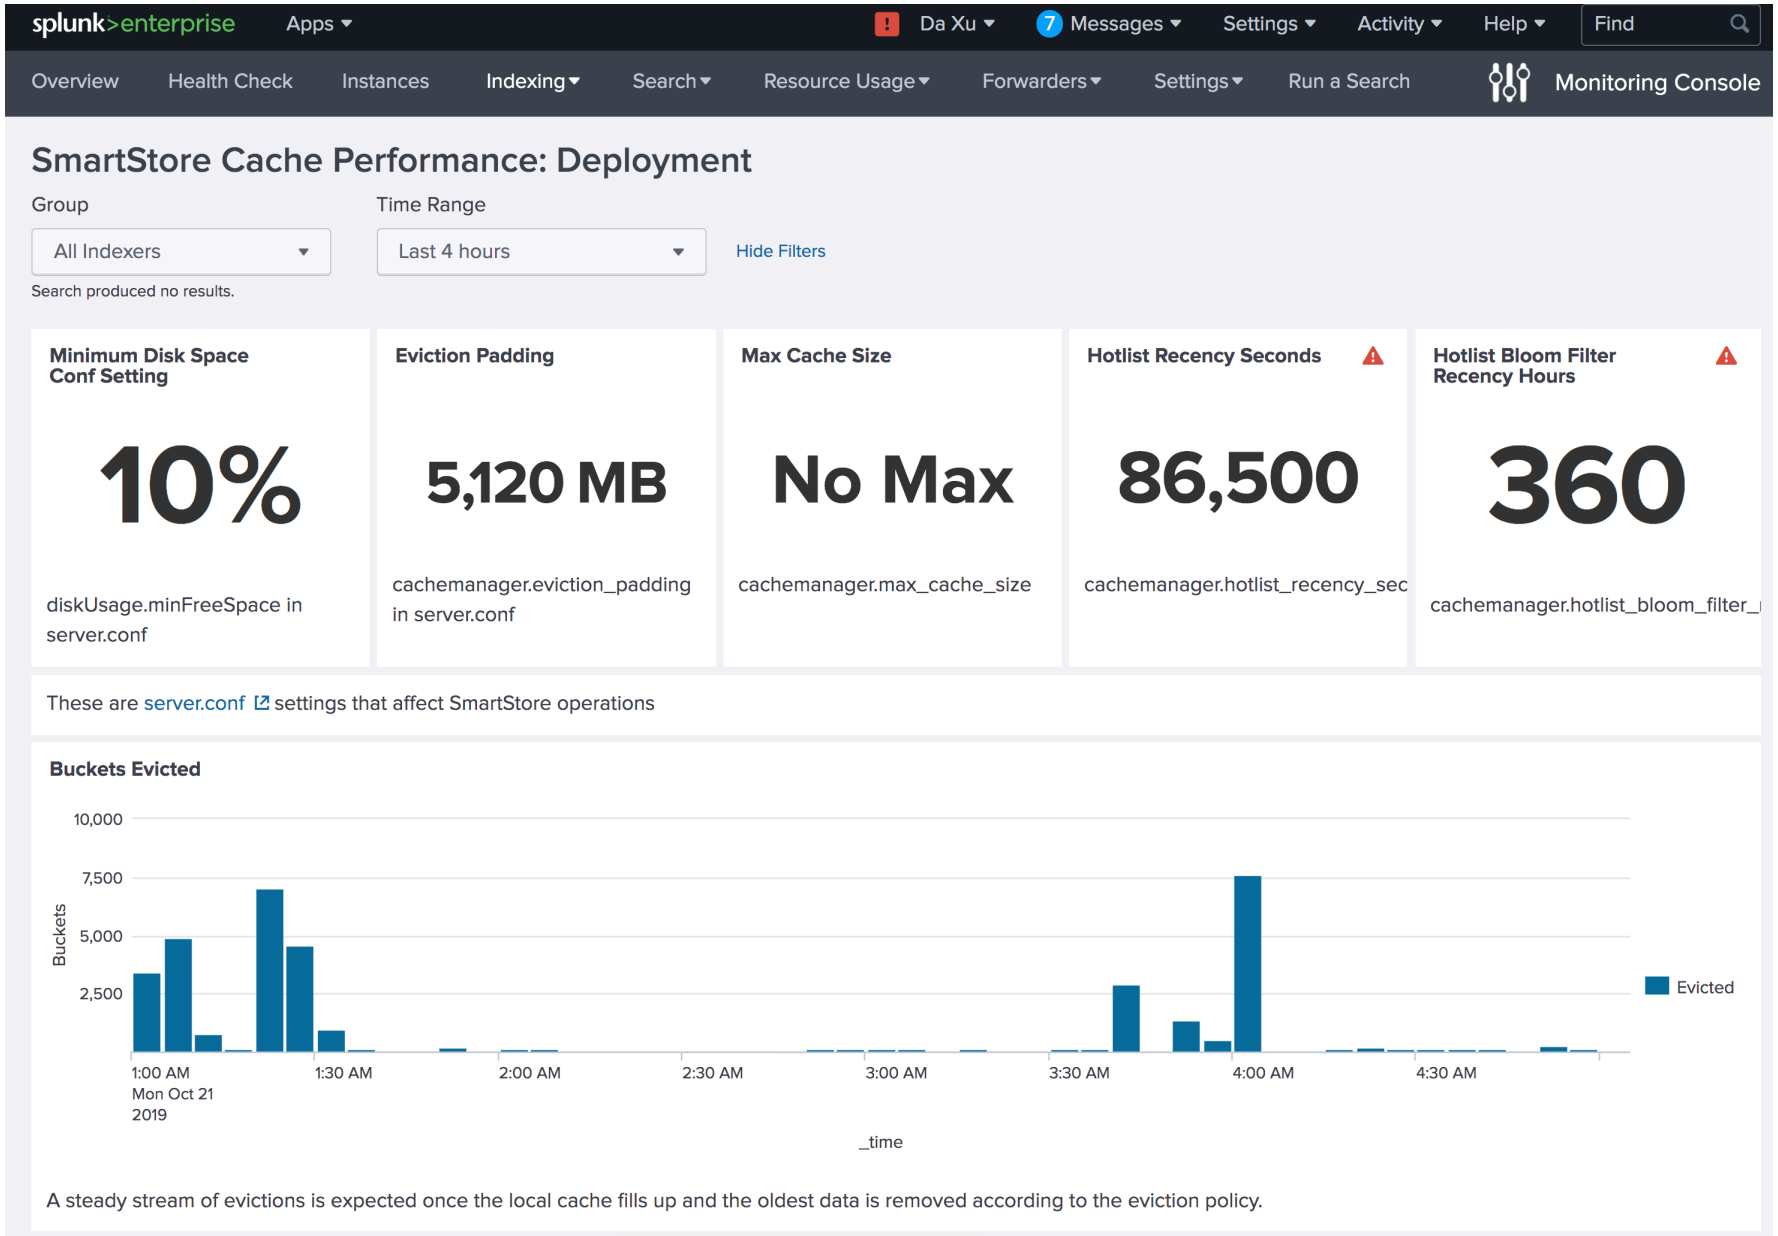

splunk> .conf19

Logs - eviction

splunk@idx-i-0f8facba5ef9ac497:~/var/log/splunk$ grep -e evicted metrics.log*

metrics.log:10-21-2019 05:36:47.392 +0000 INFO Metrics - group=spacemgr, evict_requested_kb=18908, evicted_kb=778592, elapsed_ms=55, tested=12958, evicted=5, partial_evict=770, insignificant_size=12179, cleaned=7, reserved_bytes=53687091200, free_bytes=428460789760, eviction_runs=1

Monitoring Console - eviction

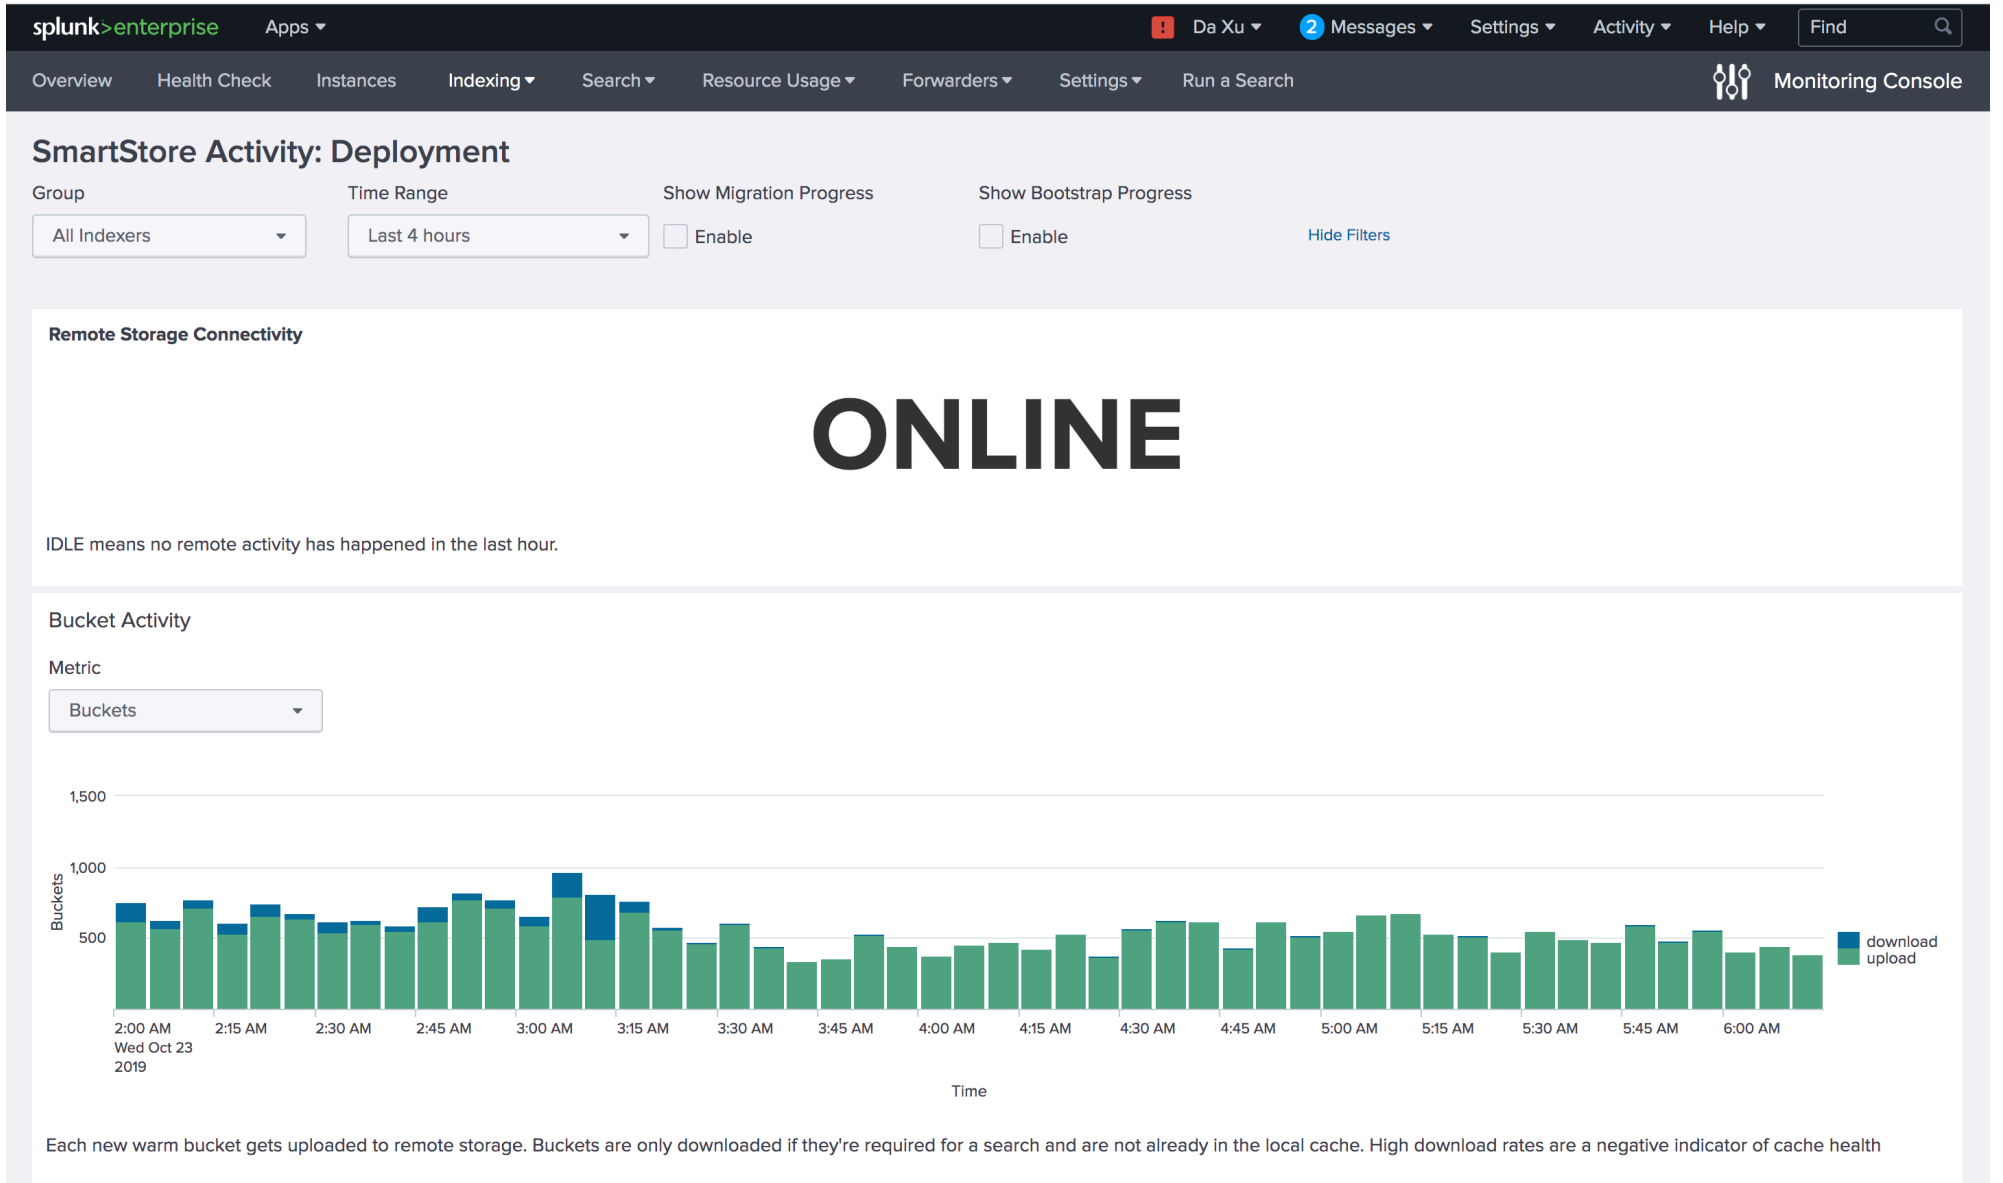

splunk> .conf19

Monitoring Console - cache hits / misses

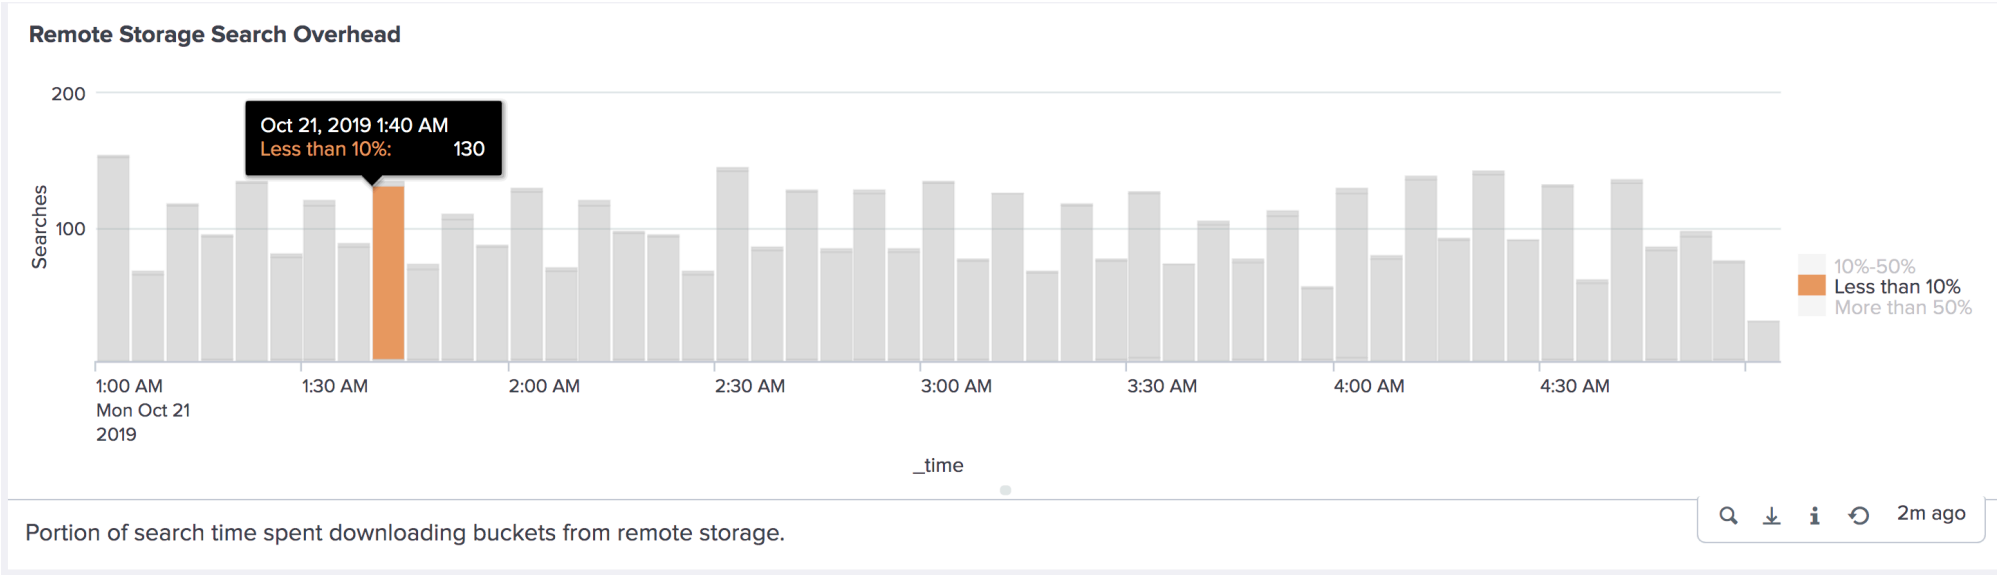

Portion of search time spent downloading buckets from remote storage.

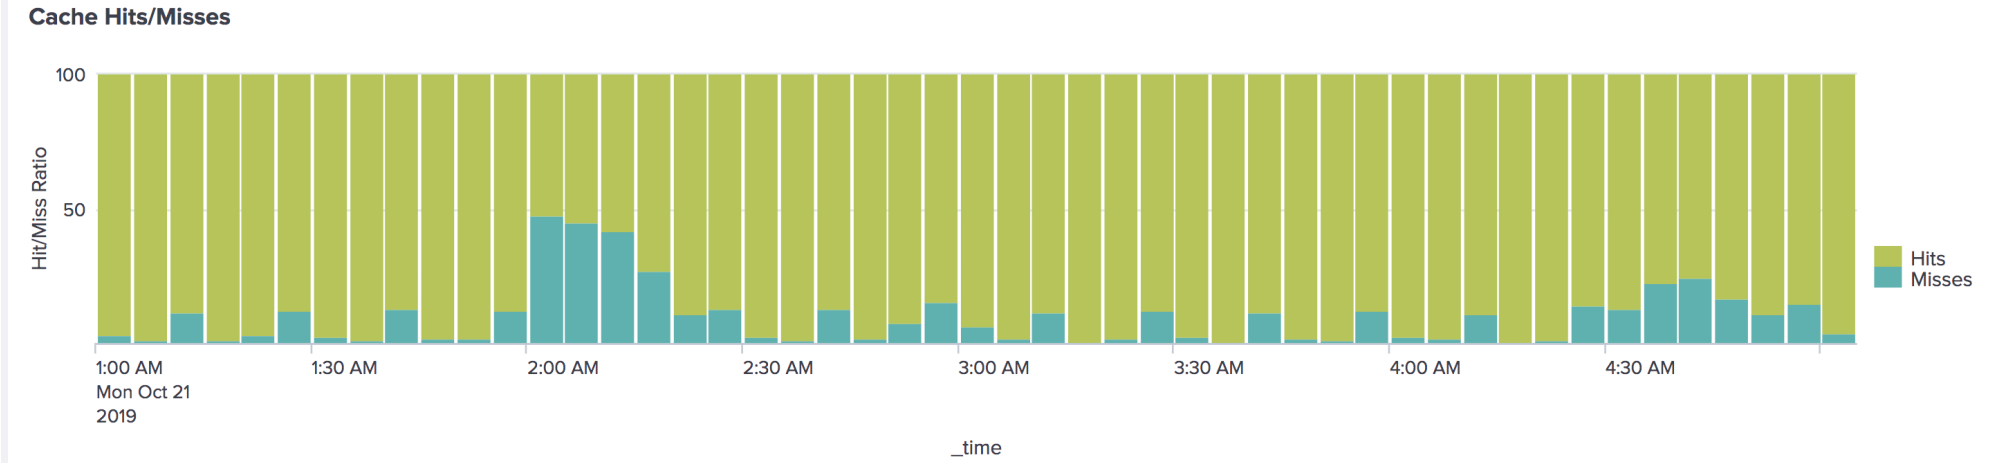

When a search needs a bucket and it's already in cache, that's a hit. If the bucket needs to be downloaded, that's a miss.

Deep Dive – Migration

splunk> .conf19

Migrations

Migration steps:

- Smart store should be enabled cluster-wide simultaneously
  1. Shut down the cluster
  2. Start the Cluster Master
  3. Enable Smart Store
  4. Push a bundle – this will finish instantly when no indexers are connected
  5. Bring up the indexers
  6. Migration will start and run in the background

Migrations – in the background

How S2 migration upload works

- Peers will upload all their searchable copies to remote storage
- Only one copy will “win” and remain on remote storage
- Restarting an indexer before S2 migration completes causes the indexer to resume migration

Migrations – in the background

Any limitations? Caveats? Scalability concerns?

- Initialization is resource heavy
- Uploading 10K's of buckets may take several hours
- Search is impacted during this time

Any capacity considerations eg. disk usage, cpu?

- Greatest impact is on the networking
- About 30% performance degradation on searching during S2 migration

Migrations – in the background

Tracking start (one entry per index):

- 03-29-2017 21:31:15.178 -0700 INFO DatabaseDirectoryManager - Remote storage migration needed for idx=foo for a bucket count=9

Tracking end of migration (all indexes):

- 03-29-2017 21:31:24.417 -0700 INFO CacheManager - Remote storage migration of buckets and summaries completed (duration_sec=9 upload_jobs=9)

Migrations – in the background

Endpoint that tracks status (in progress):

```
./master/bin/splunk search "|rest /services/admin/cacheman/_metrics |fields splunk_server
migration.*"
```

splunk_server	migration.current_job	migration.start_epoch	migration.status	migration.total_jobs
fool13-peer	3	1485306460	running	43
fool15-peer	6	1485306476	running	46
fool14-peer	5	1485306468	running	44

Other Items

- Must be using a single object store.
- If multi-region in AWS, the endpoint must be against a specific region’s S3 bucket
- No support yet for non gzip compression, tsidx minification

SmartStore at Lockheed Martin

Bill Ern Enterprise Splunk Product Owner

splunk> .conf19

Background

Interfaces Across 4 Main Categories

Splunk Enterprise Security™

Splunk IT Service Intelligence™

- Enterprise Splunk offering since 2016
- Splunk infrastructure is built in AWS
- Interfaces: ~27,394
  - Agent Based
  - Agentless
  - API-Based
  - IoT Sensors (IAI Premium Module)
- Users
  - 1,266 Users
  - 166 Power Users
- Daily Data Ingest
  - ~10.4 TB of data
- Version as of the Migration was 7.2.3

Before - High Level Indexer Tier Architecture

Before the Migration to SmartStore

Indexer
Cluster
Master

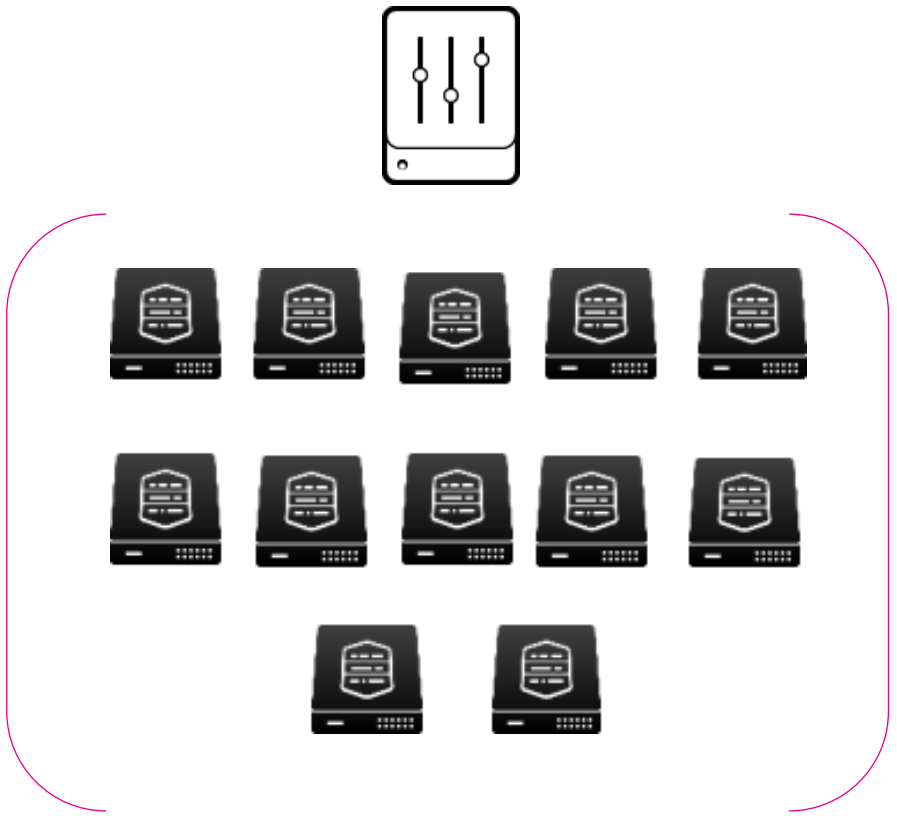

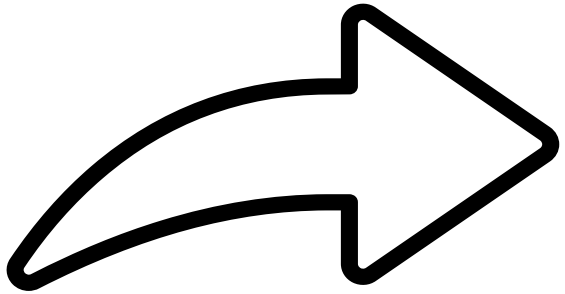

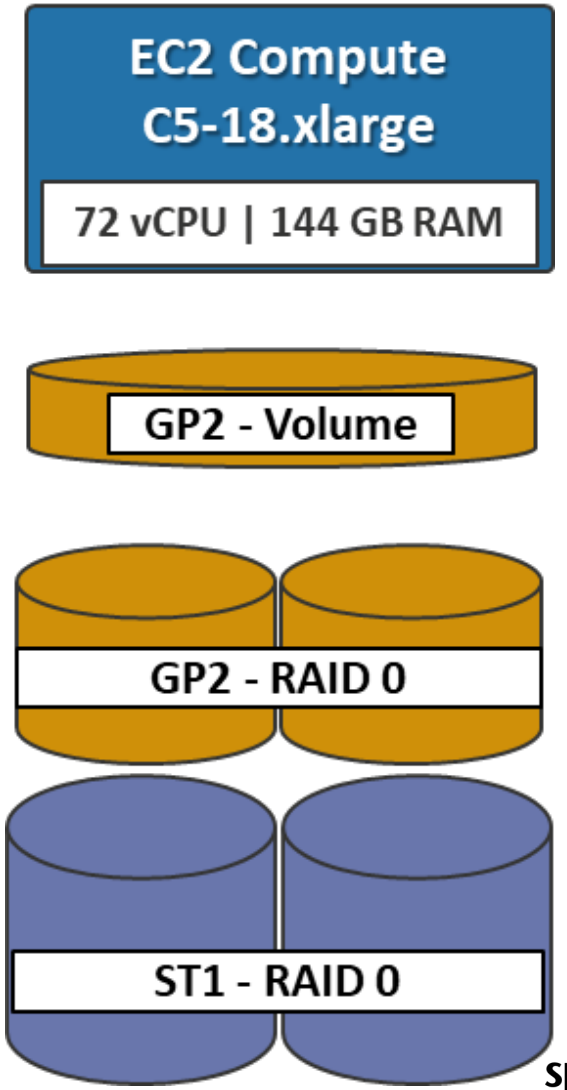

Overall Goals

Target of reducing the AWS cost by 30%

1. Change the Indexing Architecture to SmartStore

amazon web services™ | S3

2. Move to I3 Reserved Instances (RI) 3 Year convertible

amazon web services™

EC2 Reserved Instances

3. Use “Native AWS Services”

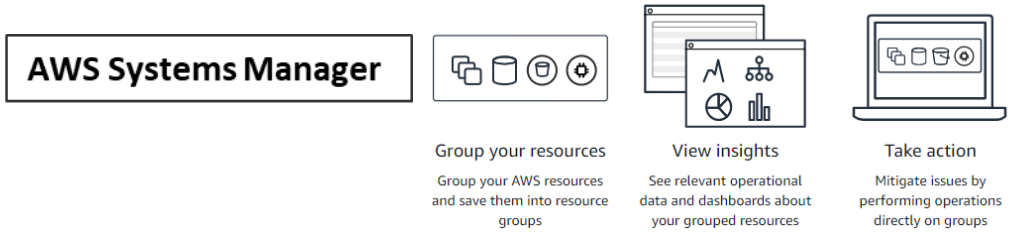

Moving to S2 Architecture

Moving to S2 Architecture

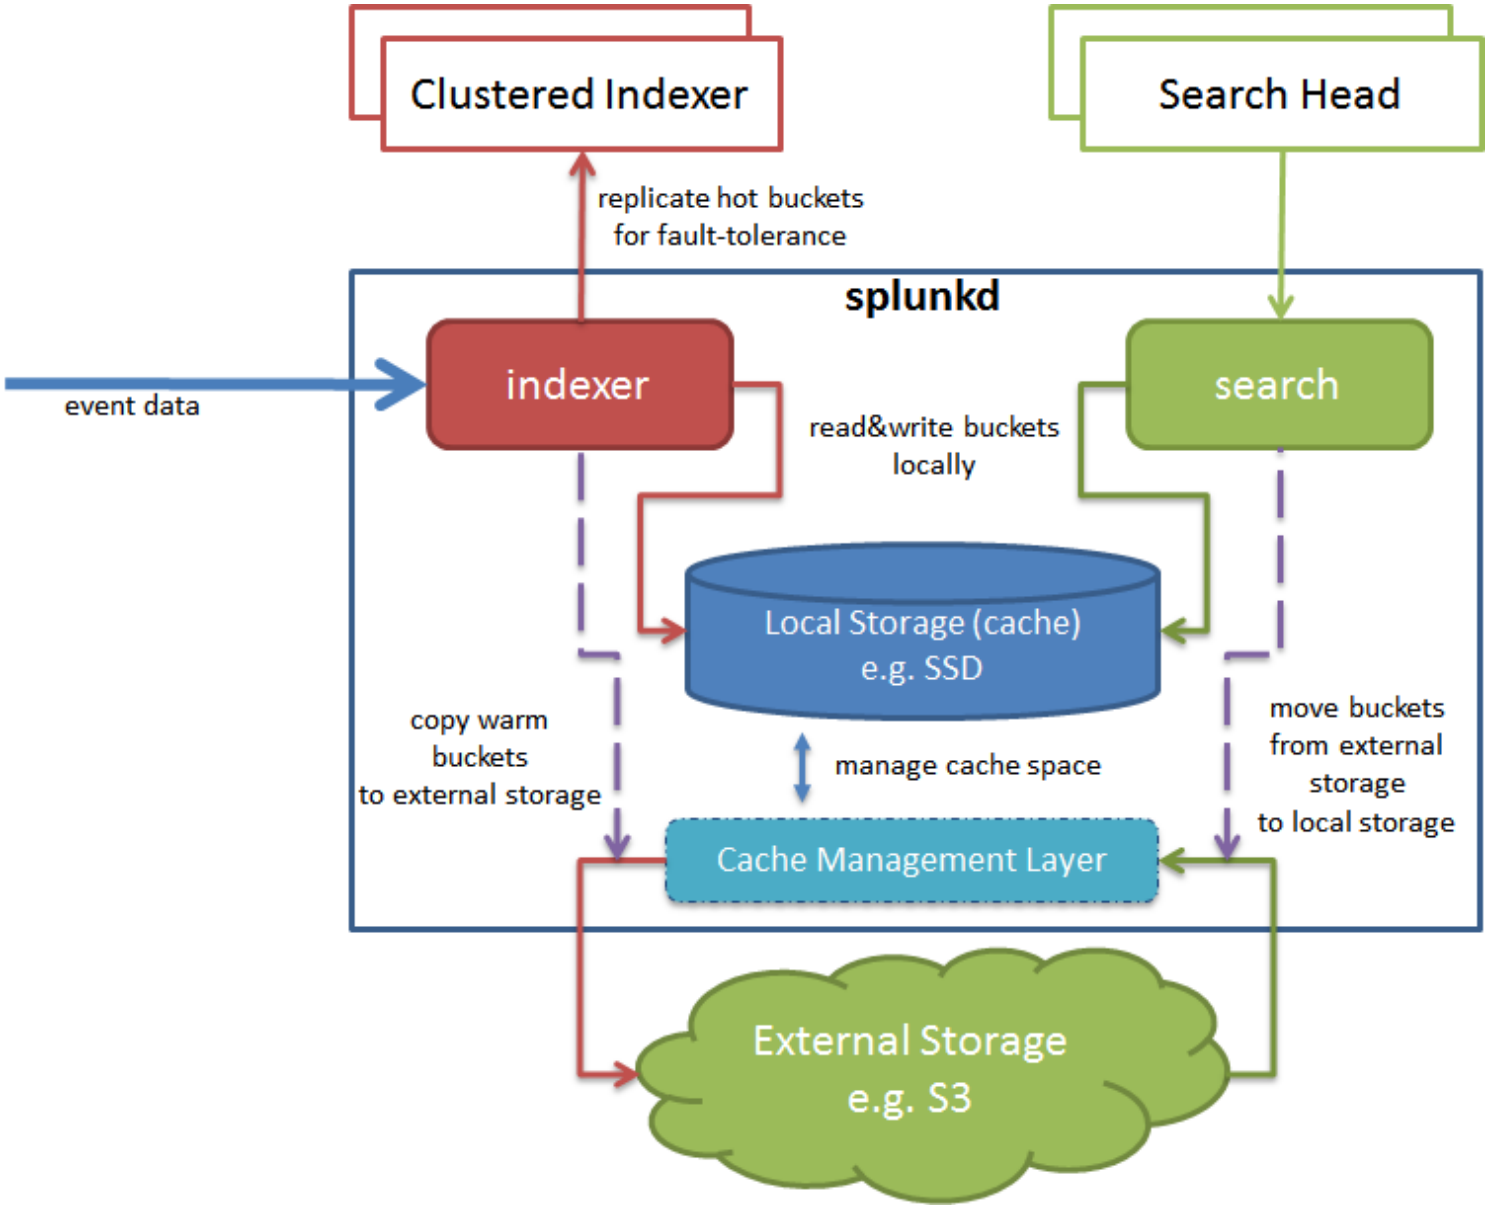

After - High Level Indexer Tier Architecture

After the Migration to Smart Store

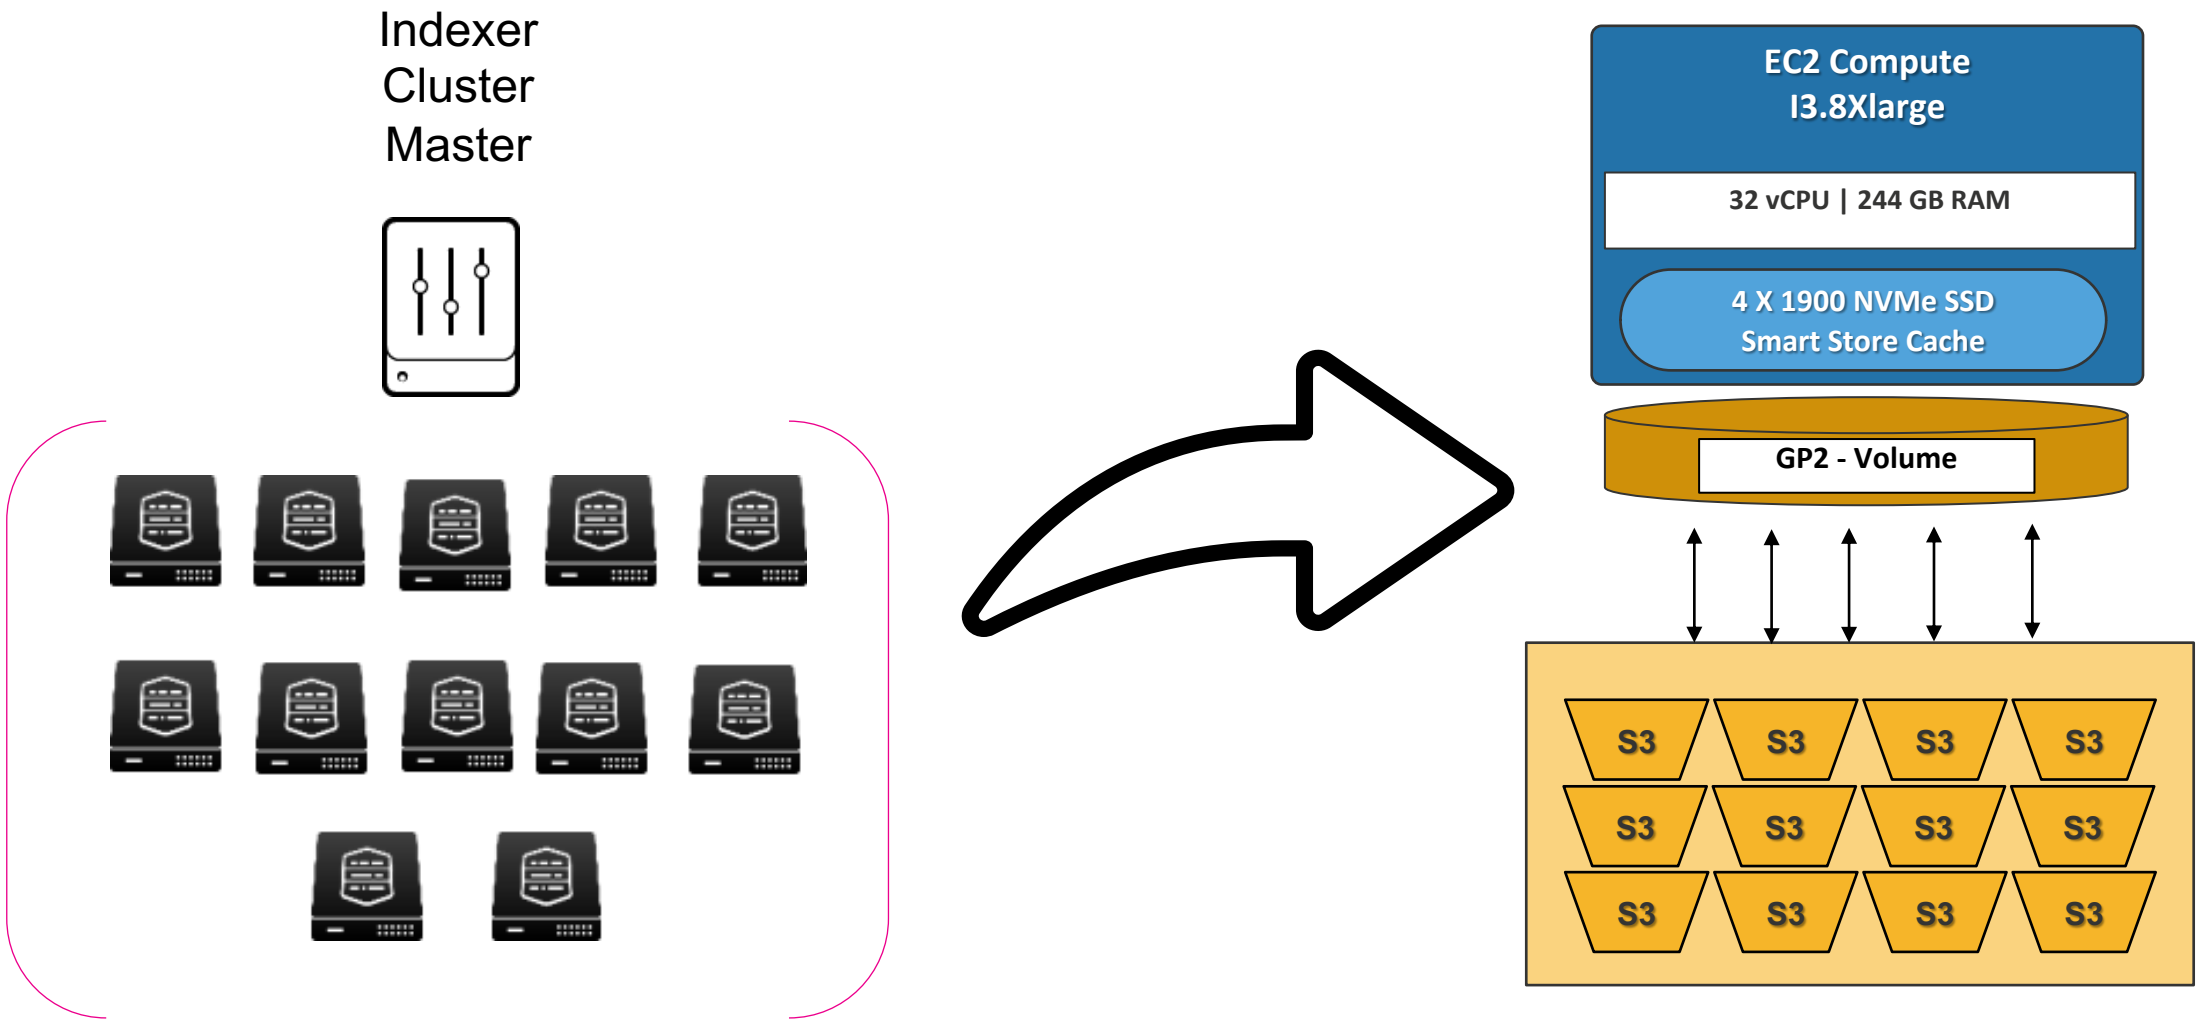

Monitoring the Migration Process

During Migration

1. Enterprise Splunk Version 7.3 Supports SmartStore Metrics Views

2. Continuously monitor indexer tier

```
$ splunk search "|rest /services/admin/cacheman/_metrics |fields splunk_server migration.*" -au
th admin:passwd
```

```
splunk_server      migration.current_job migration.start_epoch migration.status migration.total
_jobs
---------------  ---------------------- -------------------- -------------- ---------------
-----
cluster1-master                                                 not_started
peer1.ajax.com              8                 1484942186        running                35
peer2.ajax.com              7                 1484942190        running                37
peer2.ajax.com              5                 1484942194        running                36
```

After the Migration

Observations

1. Performance issues with the C5.8Xlarge instances
2. Understand the Cache Hits/Misses, bucket evictions for performance
3. Upgraded Enterprise Splunk to version 7.3 – ability to troubleshoot with the MC
4. Migrated the indexer tier to I3.8xlarge
5. Started to add additional indexers to the tier

Trouble

Troubleshooting Methods

1. Use the Monitoring Console
2. Log files
3. CLI commands
4. Rest endpoints

Lessons Learned

Take Away

1. You cannot revert to non-SmartStore after you migrate – Can be in a mixed mode
2. Read the documentation to understand the Prerequisites
3. Test the configuration on a standalone instance
4. Look at the documentation for common issues

What were the Results

Cost Reduction Efforts

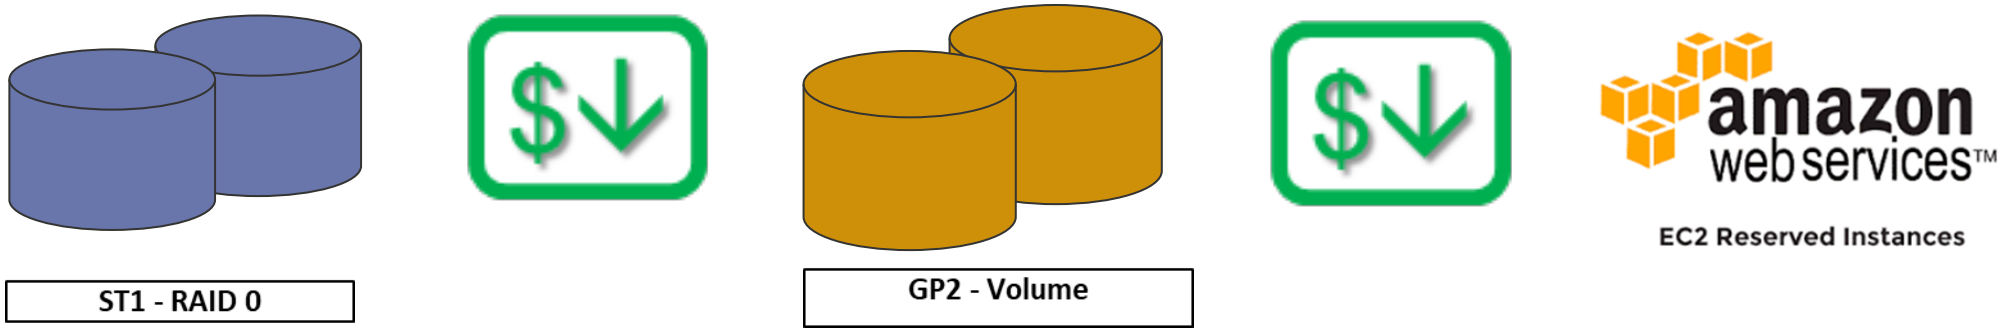

ST1 Storage

~15% Savings

GP2 Storage

~12% Savings

Cheaper I3 RI

~5% Savings

.conf19

splunk>

Thank You!

Go to the .conf19 mobile app to

RATE THIS SESSION