

Bucket Diversity: Choosing Your Search Mate Wisely

Dean Jackson

Principal Systems Engineer, DELL EMC


Simon O'Brien

Senior Sales Engineer, Splunk

.conf2016

splunk>

WHO WE ARE...



Dean Jackson



Simon O'Brien

Disclaimer

During the course of this presentation, we may make forward looking statements regarding future events or the expected performance 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 could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. 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, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward looking statements we may make. 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 functionality described or to include any such feature or functionality in a future release.


Who Is This Talk For?

About Here




Splunkability Continuum...

You're Expected To Know?



You'll Take Away?



USA

15,000
KM

AUSTRALIA


The Pain...





Buckets Are
Not All
The Same

Bucket Diversity



Bucket Challenge



Bucket Guidance

800 IOPs



“Insufficient disk I/O is the most common limitation found in a Splunk infrastructure.”

We're Here To Help



Choose Your Search Mate Wisely



HOT/WARM
Block
HDD or SSD



COLD
Block or NFS
HDD

How BIG?
SSD or HDD?
How many IOPS?

<https://docs.splunk.com/Documentation/Splunk/latest/Installation/Systemrequirements>


The Big Debate



DAS



SAN / NAS



SDS





20oz.
Pint



15oz.
Schooner



10oz.
Pot
(Middy)



5oz.
Pony

Bucket Capacity

Simple Sizing



1TB/day Ingestion

180 Day Retention

~180TB disk, Right?

Too Easy!

Sizing Up Your Search Mate



Use 50% reduction factor ROT: 15% raw/35% indexes

Multiply your RF / SF if clustered

Leave 10% spare space for index rebuilds, etc.

Hot/warm based on 70% of searches

Cold is your estimated retention

Too Easy!

Same 1TB/day with a 180 day retention requirement.

6 month desired retention

70% searches <7 days

Clustered RF2 /SF2

6 Indexers



$30 \times 1\text{TB} \times 50\% \times 2 + 10\%$
30.8TB
~5TB per Indexer



$150 \times 1\text{TB} \times 50\% \times 2 + 10\%$
165TB
~27.5TB per Indexer

That's a total ~200TB of disk!!

Long Roads, Loads Of Options...



...AND DANGERS



How Do I Make That Work?



Hot/Warm:

homePath = ?
maxDataSize = ?
maxHotSpanSecs = ?
homePath.maxDataSizeMB = ?
maxHotBuckets = ?
maxWarmDBCount = ?



Cold:

coldPath = ?
coldPath.maxDataSizeMB = ?
maxTotalDataSizeMB = ?
frozenTimePeriodInSecs = ?
enableTsidxReduction = ?

The Basics

```
[volume:hotwarm]
path = /mnt/splunk-hotwarm
maxVolumeDataSizeMB = 203400
```

Hot/warm Volume

```
[volume:cold]
path = /mnt/isilon
maxVolumeDataSizeMB = 4096000
```

Cold Volume

```
[main]
repFactor = auto
homePath = volume:hotwarm/defaultdb/db
coldPath = volume:cold/defaultdb/colddb
thawedPath = $SPLUNK_DB/defaultdb/thaweddb
homePath.maxDataSizeMB = 200000
coldPath.maxDataSizeMB = 2000000
```

Index



Things We Forget

Data Model/Report Acceleration

First of all, what's the actual difference?

- Report acceleration and summary indexing *speed up individual searches*, on a report by report basis. They do this by building collections of pre-computed search result aggregates.
- Data model acceleration speeds up reporting for the set of attributes (fields) that you define in a data model.

Acceleration

☒ Accelerate this search

Caution:

* Acceleration may increase search storage and processing costs.

* An accelerated report can return invalid results if it contains tags, event type, or other fields that are not accelerated.

[Learn More](#)

Summary range

7 Days

Data Model Acceleration → By Default

What impact does enabling these features have on my storage?!

```
Splunk-index11:/mnt/splunk-hotwarm/isilon# ls  
Colddb datamodel_summary db thaweddb
```

High performance analytics store = .tsidx files

Ad-hoc DM's are stored in dispatch on SH's – persistent DM's on IDX's

Updated every 5 minutes, cleaned every 30 minutes

Private data models cannot be accelerated

Amount of storage used is relative to:

Number of events ++ summary range

Cardinality of DM attributes – something like clientip may be high

<http://docs.splunk.com/Documentation/Splunk/latest/Knowledge/Acceleratedatamodels>

Data Model Acceleration Config

If you are setting size limits for your indexes, and not your DM's, you're doing it wrong.

```
[volume:dmsummaries]
path = /mnt/dmsummaries
maxVolumeDataSizeMB = 203400
```

```
[main]
repFactor = auto
homePath = volume:hotwarm/defaultdb/db
coldPath = volume:cold/defaultdb/colddb
tstatsHomePath = volume:dmsummaries/defaultdb/datamodel_summary
```


TSTATS && ACCELERATED DM's – Use Them!

Tstats can search distributed .tsidx files

Use the search term – FROM
datamodel=<datamodelname>

```
| tstats avg(foo) FROM  
datamodel=buttercup_games WHERE bar=valuex
```

You should expect dramatically faster search results using this method



**FAST, like something
out of Mad Max...**

6.4 Index Reduction

My Splunk SE told me it would reduce my disk utilization by 66%**
Sweet!!

Enabled Tsidx reduction
on Cluster Master

Set Cold to 1 week
before reduction
(6+ months in indexes)

Push...

“Minify”


```
[main]
repFactor = auto
homePath = volume:hotwarm/main/db
coldPath = volume:cold/main/colddb
thawedPath = $SPLUNK_DB/main/thaweddb
maxDataSize = 10000
maxHotBuckets = 10
enableTsidxReduction = true
timePeriodInSecBeforeTsidxReduction = 604800
```

Summary of Bucket Events					
_time	index	Bucket ID	Transition	Bucket Name	Reason
2016-08-25 07:14:011	windows	347	Minify	rb_1471467861_1471448246_347	Bucket Passed Minify Age
2016-08-25 07:14:12.474	perfmon	456	Minify	rb_1471467948_1471450024_456	Bucket Passed Minify Age
2016-08-25 07:04:14.165	windows	497	Minify	db_1471467499_1471452669_497	Bucket Passed Minify Age
2016-08-25 06:44:12.459	perfmon	621	Minify	db_1471466465_1471457307_621	Bucket Passed Minify Age


<http://docs.splunk.com/Documentation/Splunk/latest/Indexer/Reducetsidxdiskusage>

6.4 Index Reduction

Coldddb Before
1,552 GB




**Clustered
Indexer #1**




**Clustered
Indexer #2**


**Reduced
By 46%**



Coldddb After
837 GB



**Clustered
Indexer #1**




**Clustered
Indexer #2**

<http://docs.splunk.com/Documentation/Splunk/latest/Indexer/Reducetsidxdiskusage>

Storage Data Reduction

My storage vendor told me it would reduce my capacity by more than 100%
CRIKEY!!





Bucket Performance

How Much To Spend On Your Mate



HOT/WARM

High Performance



COLD or **HOT/WARM**

Average Performance



COLD or **FROZEN**

Slow Performance

Flashy Search Mate



1000's IOPs

FAST response time = FAST searches

Smaller footprint


Lower Power



100's IOPs

SLOWER response time = SLOWER searches


Flashy Buckets – Sweet As!!



>2X SPM




Rare Searches




Bucket Io: Indexing

4KB and 512KB IOs




Bucket Io: Searches



Indexer Bucket Sizing



7-30 days
Typical



100 IOPs per
core


2 x 12 core CPUs

2400 IOPs
5TB hot/warm
20TB cold

100-150GB/day




90 -365 days
Typical



200GB per core
hot/warm

800GB per core
cold



Putting It All Together



Dean & Simon's Bucket List



Use flash drives for hot/warm

If using cold, HDDs are good

100 IOPs per core

1TB per core

TSIDX reduction may be better than Frozen

70% of searches should be on hot/warm

$\frac{1}{2}$ Daily Ingestion x RF/SF +10% = Min. Capacity Req'd

Don't forget to account for data model acceleration



WHY?

Questions??



<https://community.emc.com/community/connect/anz/asc>

<https://bigdatabeard.com/>

THANK YOU

.conf2016