



Squeezing all the Juice out of Splunk Enterprise Security

Marquis Montgomery, CISSP | Sr. Staff Security Consultant, Splunk

Jae Jung | Professional Services Consultant, Splunk

September 23 – 25, 2017 | Washington, DC

Forward-Looking Statements

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 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.

Splunk, Splunk>, Listen to Your Data, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2017 Splunk Inc. All rights reserved.

1

Introductions and Agenda

Who are these guys, anyway?

Agenda

What will we be talking about today?

ES Under-the-Hood

Checking out the engine



ES Specific Optimizations

Enhancements specific to the ES application




Core Splunk Optimizations

Splunk Enterprise Platform Enhancements

Key Takeaways and Q&A

Tying it all together



2

ES Under-the-Hood

To tune the engine, you need to understand the engine

Security Posture

Edit ↓ ↓ ↕

Overall Security Posture : Key Security Indicators

[Edit](#)

ACCESS NOTABLES

Total Count

390 ↓
-53

ENDPOINT NOTABLES

Total Count

2k ↓
-25

NETWORK NOTABLES

Total Count

2k ↑
+43

THREAT NOTABLES

Total Count


243 ↑
+1

UBA ANOMALIES


Total Count

0 0

Overall Notable Event Occurrence By Urgency



Overall Notable Events Occurrence Trend



Top Notable Events Occurrence

rule_name	sparkline	count
Unroutable Activity Detected		1844
Host With A Recurring Malware Infection		880
Host With Old Infection Or Potential Re-Infection		655
UEBA Threat Detected		241
Default Account Activity Detected		211
Substantial Increase In Events		148
Excessive Failed Logins		109
High Or Critical Priority Host With Malware Detected		93
Host With Multiple Infections		91
Host Sending Excessive Email		55

« prev 1 2 3 next »

Top Notable Event Occurrence by Host

src	sparkline	correlation_search_count	security_domain_count	count
10.1.21.153		4	3	6
10.10.41.200		4	3	6
10.11.36.20		4	3	6
10.1.21.67		3	2	5
10.116.240.105		3	2	5
10.11.36.1		3	2	4
10.11.36.10		3	2	4
10.11.36.12		3	2	4
10.11.36.13		3	2	4
10.11.36.14		3	2	4

« prev 1 2 3 4 5 6 7 8 9 10 next »

Things You Should Know About ES and Performance

- Splunk Enterprise Security is a complex group of Splunk apps that work together, but at its core, it consists of the following components:
 - LOTS of Dashboards
 - Scheduled Searches
 - Correlation Searches
 - Lookup Generator Searches
 - Context Generator Searches
 - Threat Generator Searches
 - Data Model Acceleration
 - Lookup Tables
 - Assets & Identities Tables
 - Trackers
 - KV Store Collections
 - Incident Review
 - Investigations

Key Processes in Enterprise Security

Where can I tune for better performance?

- ▶ Data Models
- ▶ Scheduled Searches
- ▶ Lookup Tables
- ▶ KV Stores



How ES Works

Raw Event is Indexed

Data is generated, forwarded and indexed in Splunk



Data is now available for ES

| tstats queries and dashboards now see data



ES learns about threats and anomalies

ES writes these results in to notable events, summary indexing and data models



Data Model Summary Search Runs

CIM DM normalization is applied, CIM DM field key/value pairs are stored in DM TSIDX

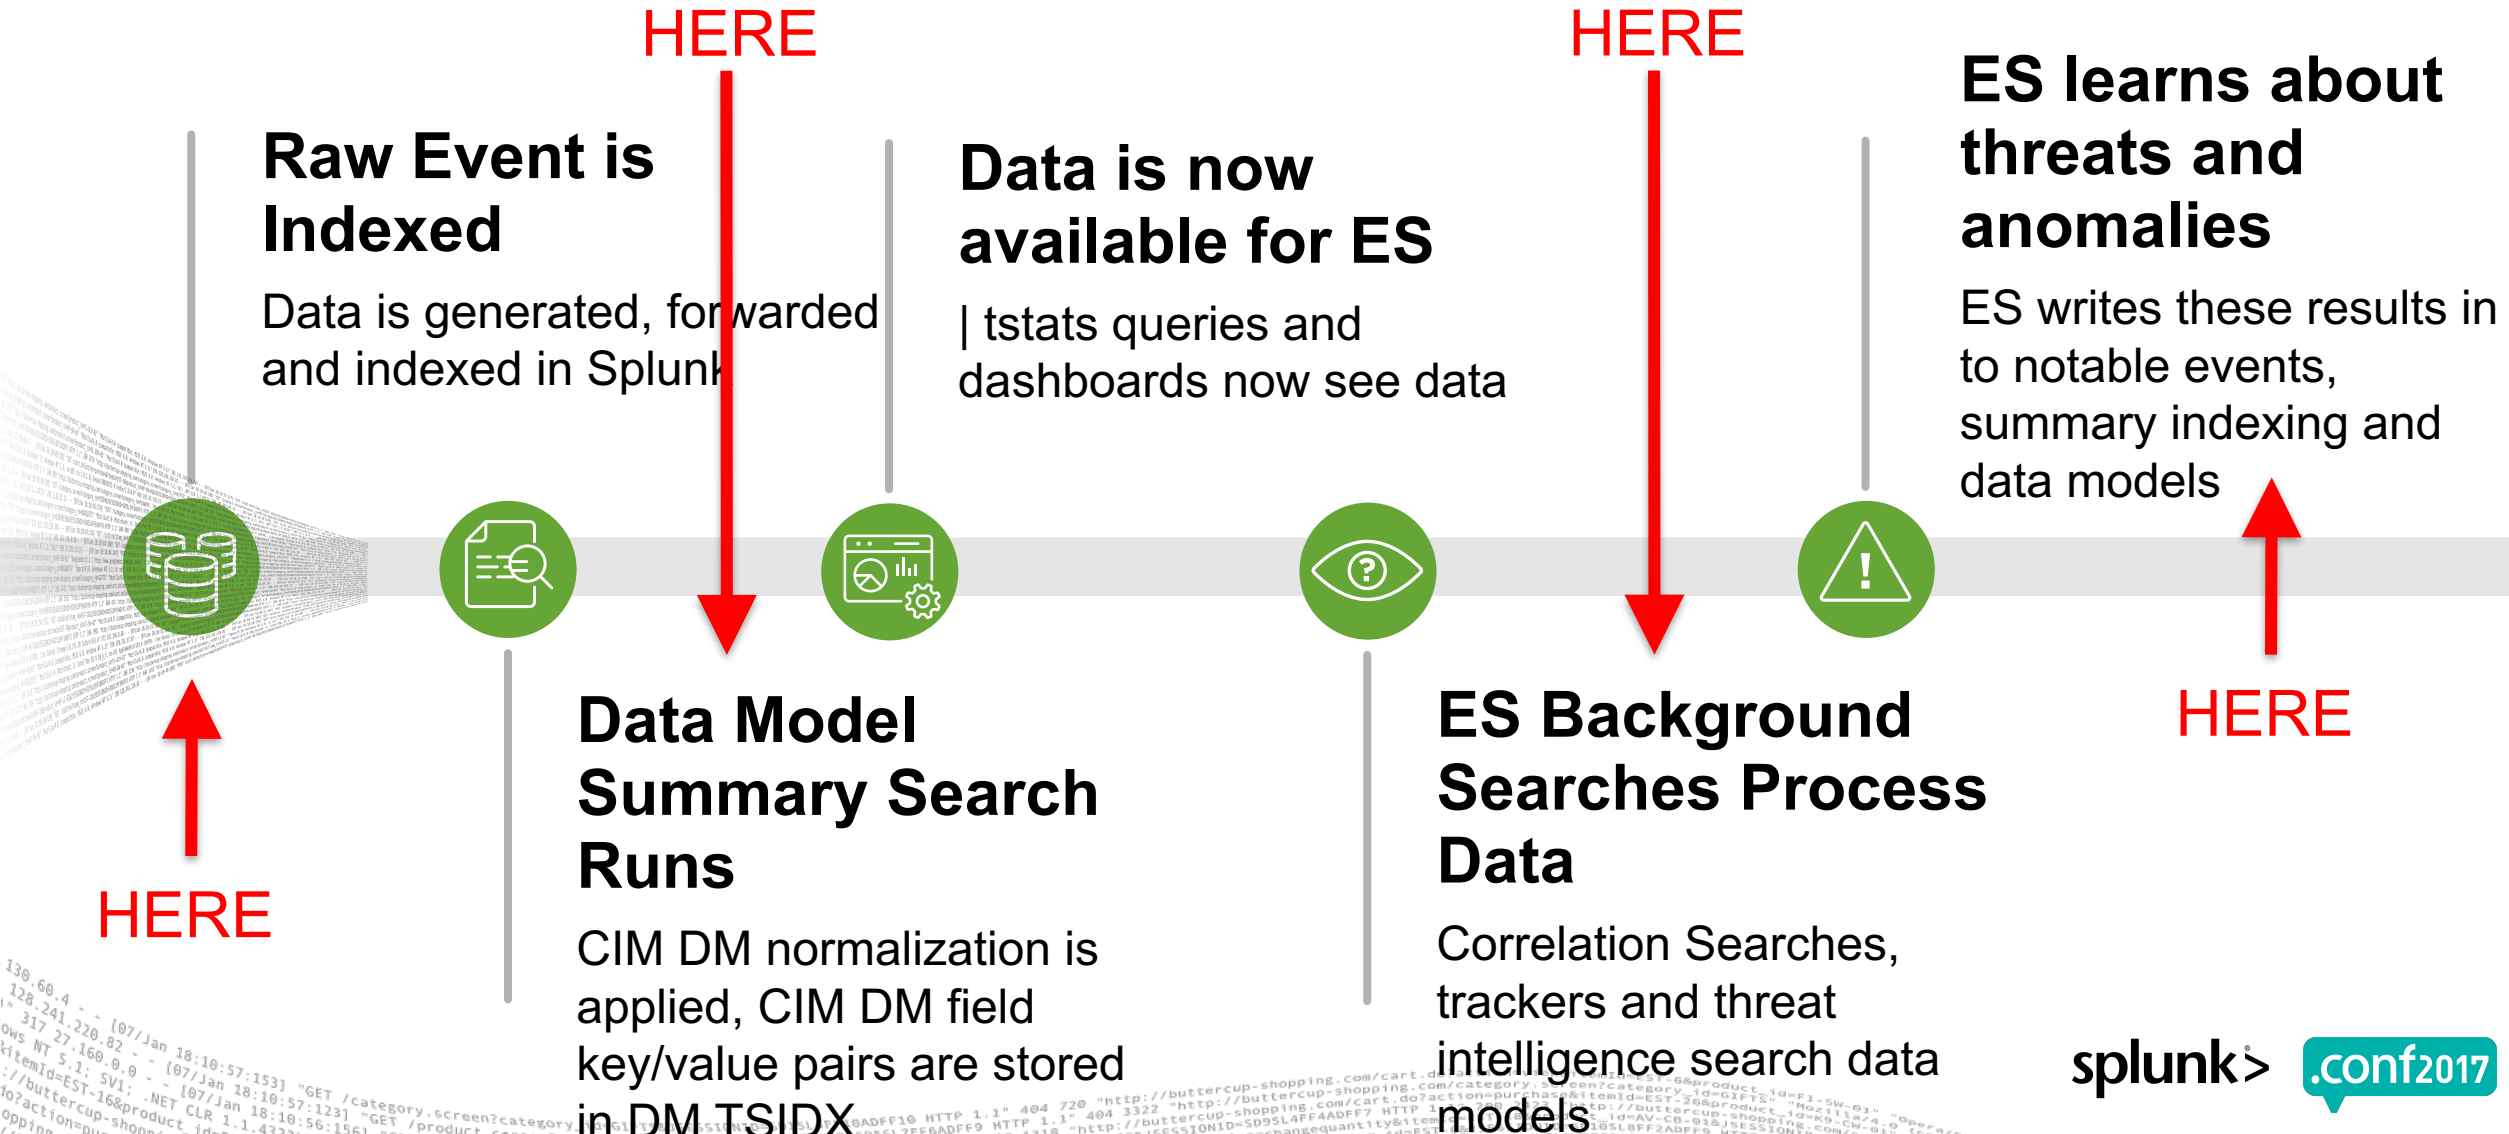


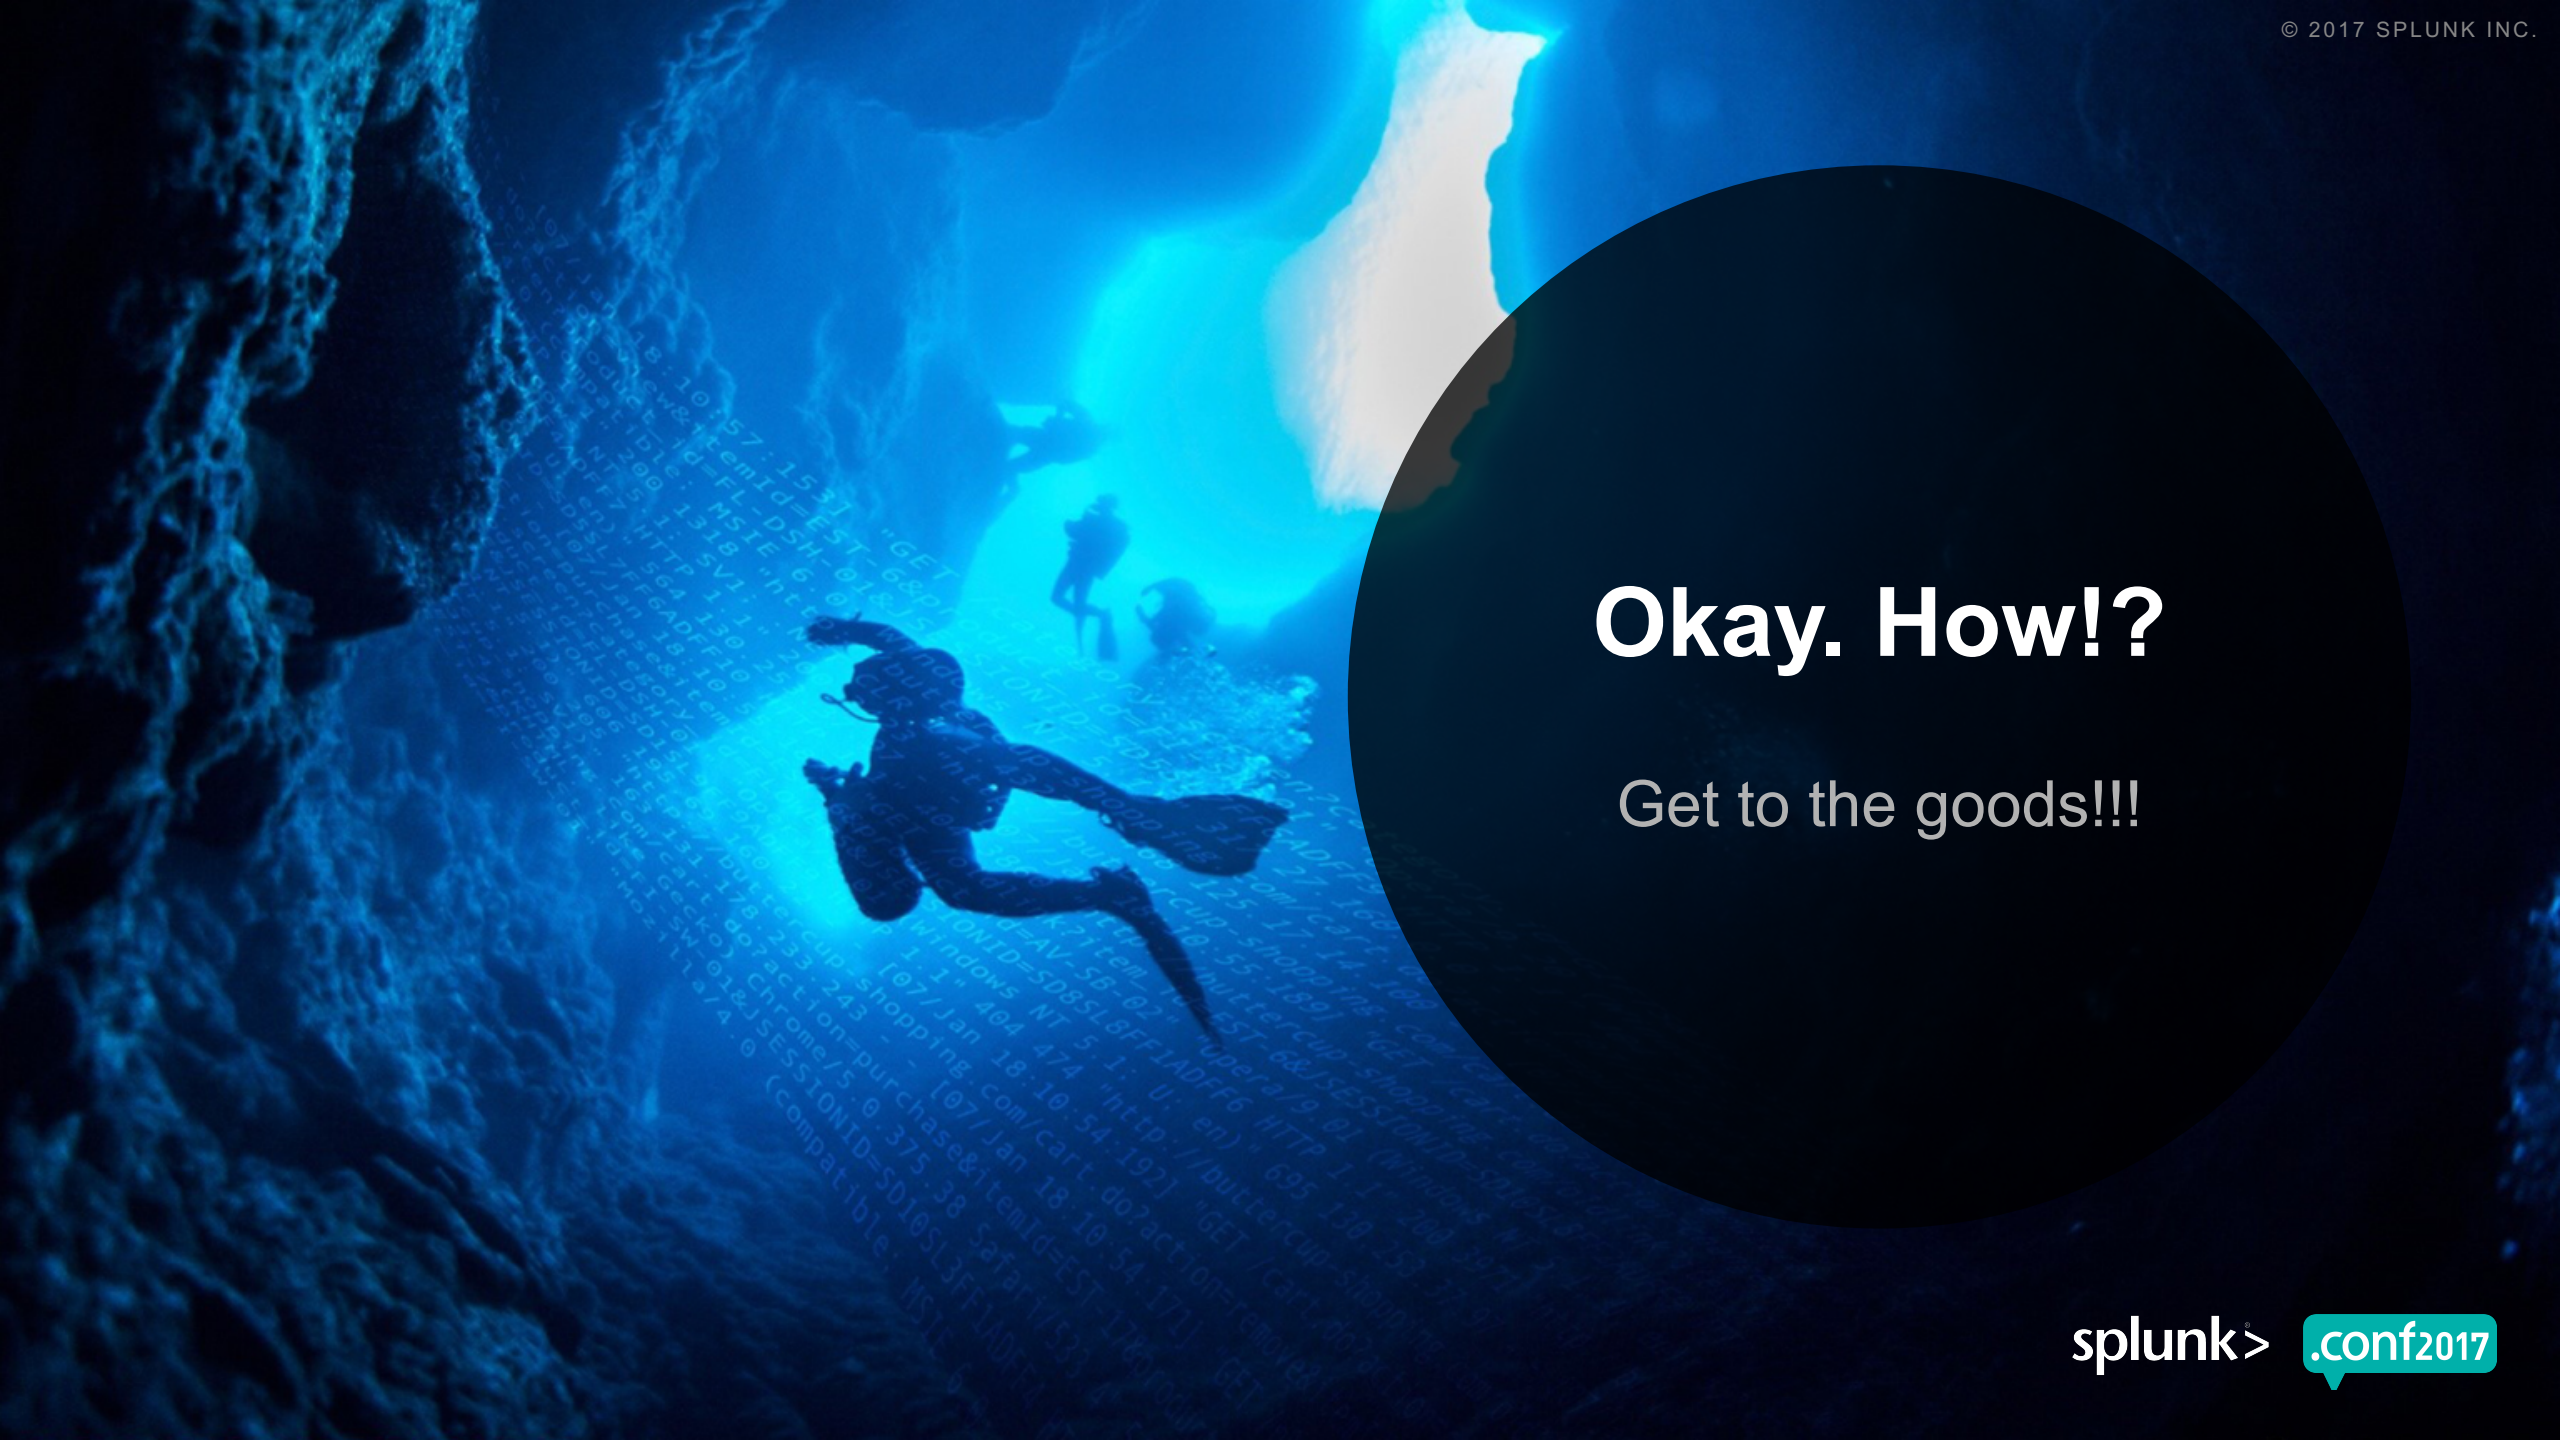
ES Background Searches Process Data

Correlation Searches, trackers and threat intelligence search data models



Places We Can Squeeze More Juice



A diver is swimming in a blue, data-filled underwater environment. The diver is silhouetted against the bright blue light coming from an opening in the cave. The water is filled with floating text, likely representing network traffic or data logs. The overall scene is mysterious and suggests a search for hidden information.

Okay. How!?

Get to the goods!!!

3

Core Splunk Optimizations

The Machine Data Platform

What Are Search Slots

► Remember...


- Correlation Searches
- Lookup Generator Searches
- Context Generator Searches
- Threat Generator Searches
- Data Model Acceleration

... are all searches and count against your 22 concurrent searches limit!

Search Scheduler Tuning

How much benefit could we possibly get??

- ▶ Search performance though? Not so great until we spread out the searches to run evenly over time
- ▶ AGGREGATE (Cumulative) search time... 🤔



130.60.4 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D15LAF10ADFF10 HTTP 1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FI-5W-03"
 128.241.220.82 - - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-DSH-01&JSESSIONID=5D35L7FF6ADFF0 HTTP 1.1" 404 3322 "http://buttercup-shopping.com/cart.do?action=purchase&itemId=EST-268product_id=KQ-CW-01"
 317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /oldlink?item_id=EST-26&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 200 1318 "http://buttercup-shopping.com/cart.do?action=changequantity&itemId=EST-18&product_id=AV-CB-01&JSESSIONID=5D55L9FF1ADFF3"
 10.2.1.1 - - [07/Jan 18:10:56:156] "GET /oldlink?item_id=EST-6&JSESSIONID=5D15L8FF2ADFF9 HTTP 1.1" 200 3865 "http://buttercup-shopping.com/cart.do?action=remove&itemId=EST-14"
 10.2.1.1 - - [07/Jan 18:10:56:156] "GET /category.screen?category_id=FLOWERS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 200 3865 "http://buttercup-shopping.com/cart.do?action=remove&itemId=EST-14"
 10.2.1.1 - - [07/Jan 18:10:56:156] "GET /category.screen?category_id=FLOWERS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 200 3865 "http://buttercup-shopping.com/cart.do?action=remove&itemId=EST-14"


Data Balancing

Use the resources at our disposal

- ▶ Even data distribution is crucial in parallel computing
 - We have powerful indexers at our disposal, we should be using them
- ▶ Ways to improve data distribution:
 - Enable parallel pipelines on intermediate forwarders (UF and HF)(In server.conf)
 - Route directly from Universal Forwarders to Indexers where possible
 - Consider the following changes to forwarders' outputs.conf:
 - forceTimebasedAutoLB = true
 - autoLBFrequency
 - autoLBVolume (6.6 only)

Data Balancing

Use the resources at our disposal




- ▶ | tstats summariesonly=t count WHERE index=* by splunk_server _time | timechart span=5m sum(count) by splunk_server

```
130.60.4 - - [07/Jan 18:10:57:123] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
128.241.220.82 - - [07/Jan 18:10:57:123] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
317.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
130.60.4 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
128.241.220.82 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
317.27.160.0 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=5D55L9FF1ADFF3 HTTP 1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-26&product_id=RP-LI-02" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_2; rv:53.0) Gecko/20100801 Firefox/53.0
```


Remove Unnecessary TAs

- ▶ Splunk ES makes use of tagged eventtypes within applications to generate syntax for searches and data models
- ▶ An excessive amount of tags will add to execution time of searches and data model acceleration time
- ▶ **ADVANCED** Tip: Disable eventtypes that will not actually reference any data in your environment, ever



Bundle Size Matters

- ▶ Search performance at the SH and IDX tier is greatly impacted by the bundle
 - The larger it is, the greater the impact
 - Large bundles over WAN links (such as indexers in the cloud) simply exacerbate the problem
- ▶ Bundle size blowouts can be caused by a number of factors
 - Large lookups
 - "backups" of configuration changes
 - Core dumps
 - Sneaky files like .git versioning metadata that could be included in automation process
 - Support files used in complex apps (DBX or Tripwire)

Bundle Size Matters


Contents of
\$SPLUNK_HOME/var/run

```
[root@master run]#
[root@master run]#
[root@master run]#
[root@master run]#
[root@master run]#
[root@master run]# pwd
/opt/splunk/var/run
[root@master run]# ls -lah
total 943M
drwx--x--x  5 root root  4.0K Aug  8 21:41
drwx--x--x  7 root root  4.0K Oct 25 2016
-rw-----  1 root root   50K Aug  1 03:29 master.splunktools.com-1501554602-1501
-rw-----  1 root root 235M Aug  1 03:29 master.splunktools.com-1501558190.bundle
-rw-----  1 root root   45 Aug  1 03:29 master.splunktools.com-1501558190.bundle
-rw-----  1 root root 235M Aug  8 21:37 master.splunktools.com-1502228244.bundle
-rw-----  1 root root   44 Aug  8 21:37 master.splunktools.com-1502228244.bundle
drwx-----  7 root root  4.0K Jun 27 10:16 searchpeers
-rw-----  1 root root   299 Jul 11 18:56 serverclass.xml
-rw-----  1 root root 237M Jun 23 16:52 sh.splunktools.com-1498236732.bundle
-rw-----  1 root root   44 Jun 23 16:52 sh.splunktools.com-1498236732.bundle.i
-rw-----  1 root root 237M Jun 23 16:56 sh.splunktools.com-1498236990.bundle
-rw-----  1 root root   45 Jun 23 16:56 sh.splunktools.com-1498236990.bundle.i
drwx--x--x 13 root root  4.0K Aug  9 21:08 splunk
drwx-----  3 root root  4.0K Oct 15 2015 tmp
[root@master run]#
```


4

ES Optimizations

Data Model Tuning



- ▶ ES utilizes several Data Models from the Splunk Common Information Model.
- ▶ Data Model Acceleration summarizes all events in scope down to key value pairs of specific fields, as defined in the Data Model.
- ▶ By default, Splunk searches all indexes for data relevant to a particular data model, and is normally filtered by special tags.
- ▶ Data Models can be **tuned to specific indexes for each data model**, resulting in better efficiency in summarizing the key value pairs needed for the Data Model.

CIM Setup

[Apps](#) » Splunk_SA_CIM

Splunk Common Information Model Add-on Set Up

By default a datamodel will search across all indexes. Use the configuration panel below to constrain data model searches to specific indexes.

Data Models

Alerts

No restriction

Application State

No restriction

Authentication

Restricted to: main, risk, twitter

Certificates

No restriction

Change Analysis

No restriction

Compute Inventory

No restriction

Databases

No restriction

DLP

No restriction

Email

No restriction

Indexes

Settings

Indexes

[Edit Manually](#)

[Learn More](#)

Filter

	Name ^	App ⇅	Current Size ⇅
<input checked="" type="checkbox"/>	main	org_all_indexes	5,081 MB
<input checked="" type="checkbox"/>	risk	org_all_indexes	3 MB
<input checked="" type="checkbox"/>	twitter	org_all_indexes	1 MB
<input type="checkbox"/>	_audit	org_all_indexes	520 MB
<input type="checkbox"/>	_internal	org_all_indexes	2,435 MB
<input type="checkbox"/>	_introspection	system	1,161 MB
<input type="checkbox"/>	_telemetry	system	1 MB
<input type="checkbox"/>	_thefishbucket	org_all_indexes	1 MB
<input type="checkbox"/>	add_on_builder_index	splunk_app_addon-builder	1 MB
<input type="checkbox"/>	cim_modactions	org_all_indexes	3 MB
<input type="checkbox"/>	cim_summary	org_all_indexes	1 MB
<input type="checkbox"/>	endpoint_summary	org_all_indexes	1 MB
<input type="checkbox"/>	history	org_all_indexes	1 MB
<input type="checkbox"/>	ioc	org_all_indexes	1 MB


Data Model Tuning

Use the Configure > CIM Setup menu in ES

datamodels.conf acceleration.backfill_time

Limit the impact of Data Model Acceleration

- ▶ Data Model Activity consumes search slots that you may need for ad hoc search
- ▶ Sometimes, its better to not backfill old data model summaries all at once
- ▶ You can limit how far back Splunk attempts to summarize datamodels with backfill_time in datamodels.conf



```
130.60.4 - - [07/Jan 18:10:57:153] "GET /category.screen?category_id=GIFTS&JSESSIONID=SD1SLAFF10ADFF10 HTTP/1.1" 404 720 "http://buttercup-shopping.com/cart.do?action=view&itemId=EST-6&product_id=FL-SW-01" Operate 20
128.241.220.82 - - [07/Jan 18:10:57:123] "GET /product.screen?product_id=FL-DISH-01&JSESSIONID=SD5SL7FF6ADFF9 HTTP/1.1" 404 3322 "http://shopping.com/cart.do?action=purchase&itemId=EST-20&product_id=K9-CU-01" Compute 20
317.27.160.0.0 - - [07/Jan 18:10:56:156] "GET /oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
.../buttercup-shopping.com/oldlink?item_id=EST-16&product_id=RP-LI-02" 468 125.17.14.11 link?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3 HTTP/1.1" 200 1318 "http://buttercup-shopping.com/oldlink?item_id=EST-26&JSESSIONID=SD5SL9FF1ADFF3" Compute 20
```

datamodels.conf acceleration.backfill_time

- ▶ acceleration.backfill_time = <relative-time-str>
- ▶ * ADVANCED: Specifies how far back in time the Splunk software should create its column stores. * ONLY set this parameter if you want to backfill less data than the retention period set by 'acceleration.earliest_time'. You may want to use this parameter to limit your time window for column store creation in a large environment where initial creation of a large set of column stores is an expensive operation.
- ▶ * WARNING: Do not set 'acceleration.backfill_time' to a narrow time window. If one of your indexers is down for a period longer than this backfill time, you may miss accelerating a window of your incoming data.
- ▶ * MUST be set to a more recent time than 'acceleration.earliest_time'. For example, if you set 'acceleration.earliest_time' to '-1y' to retain your column stores for a one year window, you could set 'acceleration.backfill_time' to '-20d' to create column stores that only cover the last 20 days. However, you cannot set 'acceleration.backfill_time' to '-2y', because that goes farther back in time than the 'acceleration.earliest_time' setting of '-1y'. * Defaults to empty string (unset).
- ▶ When 'acceleration.backfill_time' is unset, the Splunk software always backfills fully to 'acceleration.earliest_time.'

Assets and Identities Table Lookup Performance


- ▶ ES carries along with it a number of lookup tables, two of which could become very large.
- ▶ The process of “indexing” large lookups could slow down ES
- ▶ If you see a long period of time in Job Inspector for search.command.lookups, preventing indexing of large lookups may provide a performance improvement.
- ▶ limits.conf tweak `max_memtable_bytes` slightly larger than your assets/identities
 - `max_memtable_bytes = <integer>`
 - * Maximum size, in bytes, of static lookup file to use an in-memory index for.
 - * Lookup files with size above `max_memtable_bytes` will be indexed on disk
 - * A large value results in loading large lookup files in memory leading to bigger process memory footprint.
 - * Caution must be exercised when setting this parameter to arbitrarily high values!
 - * Default: 10000000 (10MB)

Assets and Identities Table Lookup Performance

search.command.lookups


2.69	command.search.filter	64	-	-
1.51	command.search.calcfields	64	193,555	193,555
0.54	command.search.fieldalias	64	193,555	193,555
0.00	command.search.index.usec_1_8	363,541	-	-
0.00	command.search.index.usec_4096_32768	11	-	-
0.00	command.search.index.usec_512_4096	6,808	-	-
0.00	command.search.index.usec_64_512	22,853	-	-
0.00	command.search.index.usec_8_64	31,369	-	-
59.73	command.search.rawdata	64	-	-
5.22	command.search.kv	64	-	-
4.51	command.search.lookups	64	193,555	193,555
0.15	command.search typer	64	2,460	2,460
0.06	command.search.tags	64	2,460	2,460
0.02	command.search.summary	71	-	-
0.00	dispatch.check_disk_usage	2	-	-
0.02	dispatch.createdSearchResultInfrastructure	1	-	-
0.17	dispatch.evaluate	1	-	-
0.17	dispatch.evaluate.search	1	-	-
18.98	dispatch.fetch	72	-	-
20.30	dispatch.finalizeRemoteTimeline	1	-	-
0.07	dispatch.parserThread	71	-	-

Key Takeaways

- 
- ▶ Getting more “juice” out of Enterprise Security is really about Splunk optimization.
 - ▶ Understanding the under-the-hood inner workings make ES easier to tune and optimize.
 - ▶ There are a few easy knobs you can turn that drastically impact performance – make one change at a time and test!

5

Q&A



Making Machine Data Accessible, Usable And Valuable To Everyone.

Thank You

Don't forget to **rate this session** in the
.conf2017 mobile app

splunk® **.conf2017**