splunk> .conf2017

What's New in Splunk Enterprise and Cloud

Todd Untrecht | VP, Product Management

Sept 26, 2017 | Washington, DC

Action Packed Session!

Live Demos

Audience Challenges

Learn More

Session Name

Day | Time

Speaker Name, Company

Learn More "Pop Outs"

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 Enterprise Releases

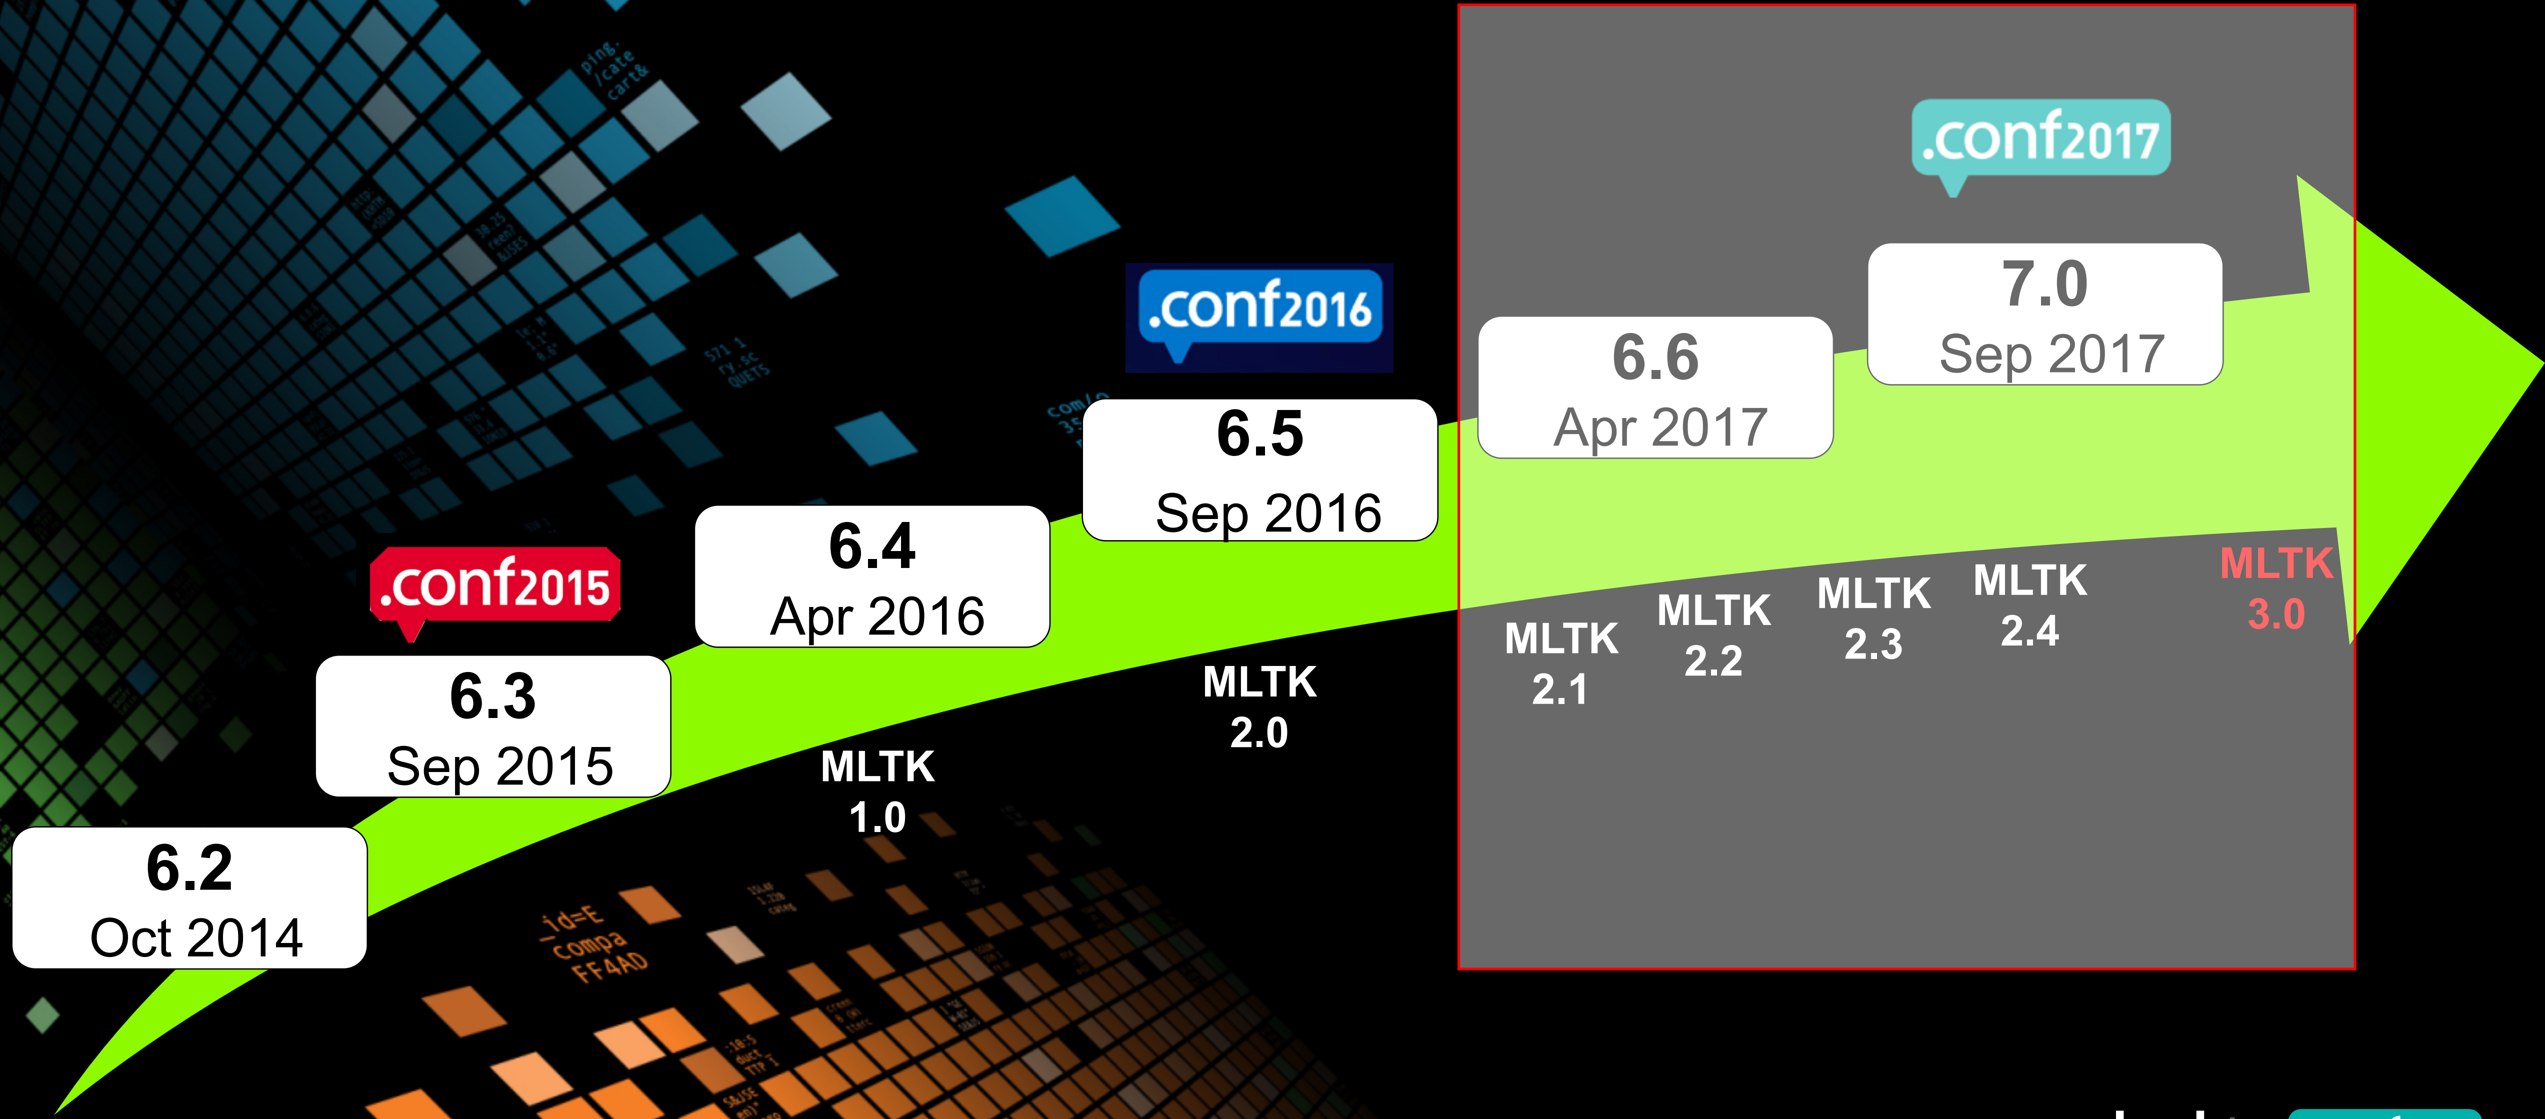

6.2
Oct 2014

.conf2015

6.3
Sep 2015

6.4
Apr 2016

MLTK
1.0

.conf2016

6.5
Sep 2016

MLTK
2.0

6.6
Apr 2017

MLTK
2.1

MLTK
2.2

MLTK
2.3

MLTK
2.4

.conf2017

7.0
Sep 2017

MLTK
3.0

Monitoring & Investigation

Increased productivity for a wider range of users

Solve problems faster

Performance, Scale & Manageability

Faster, more scalable and easier to manage than ever before

Grow more easily

Machine Learning & Intelligence

More powerful and extensible machine learning capabilities

Do less work

Monitoring & Investigation

Increased productivity for a wider range of users

Performance, Scale & Manageability

Faster, more scalable and easier to manage than ever before

Machine Learning & Intelligence

More powerful and extensible machine learning capabilities

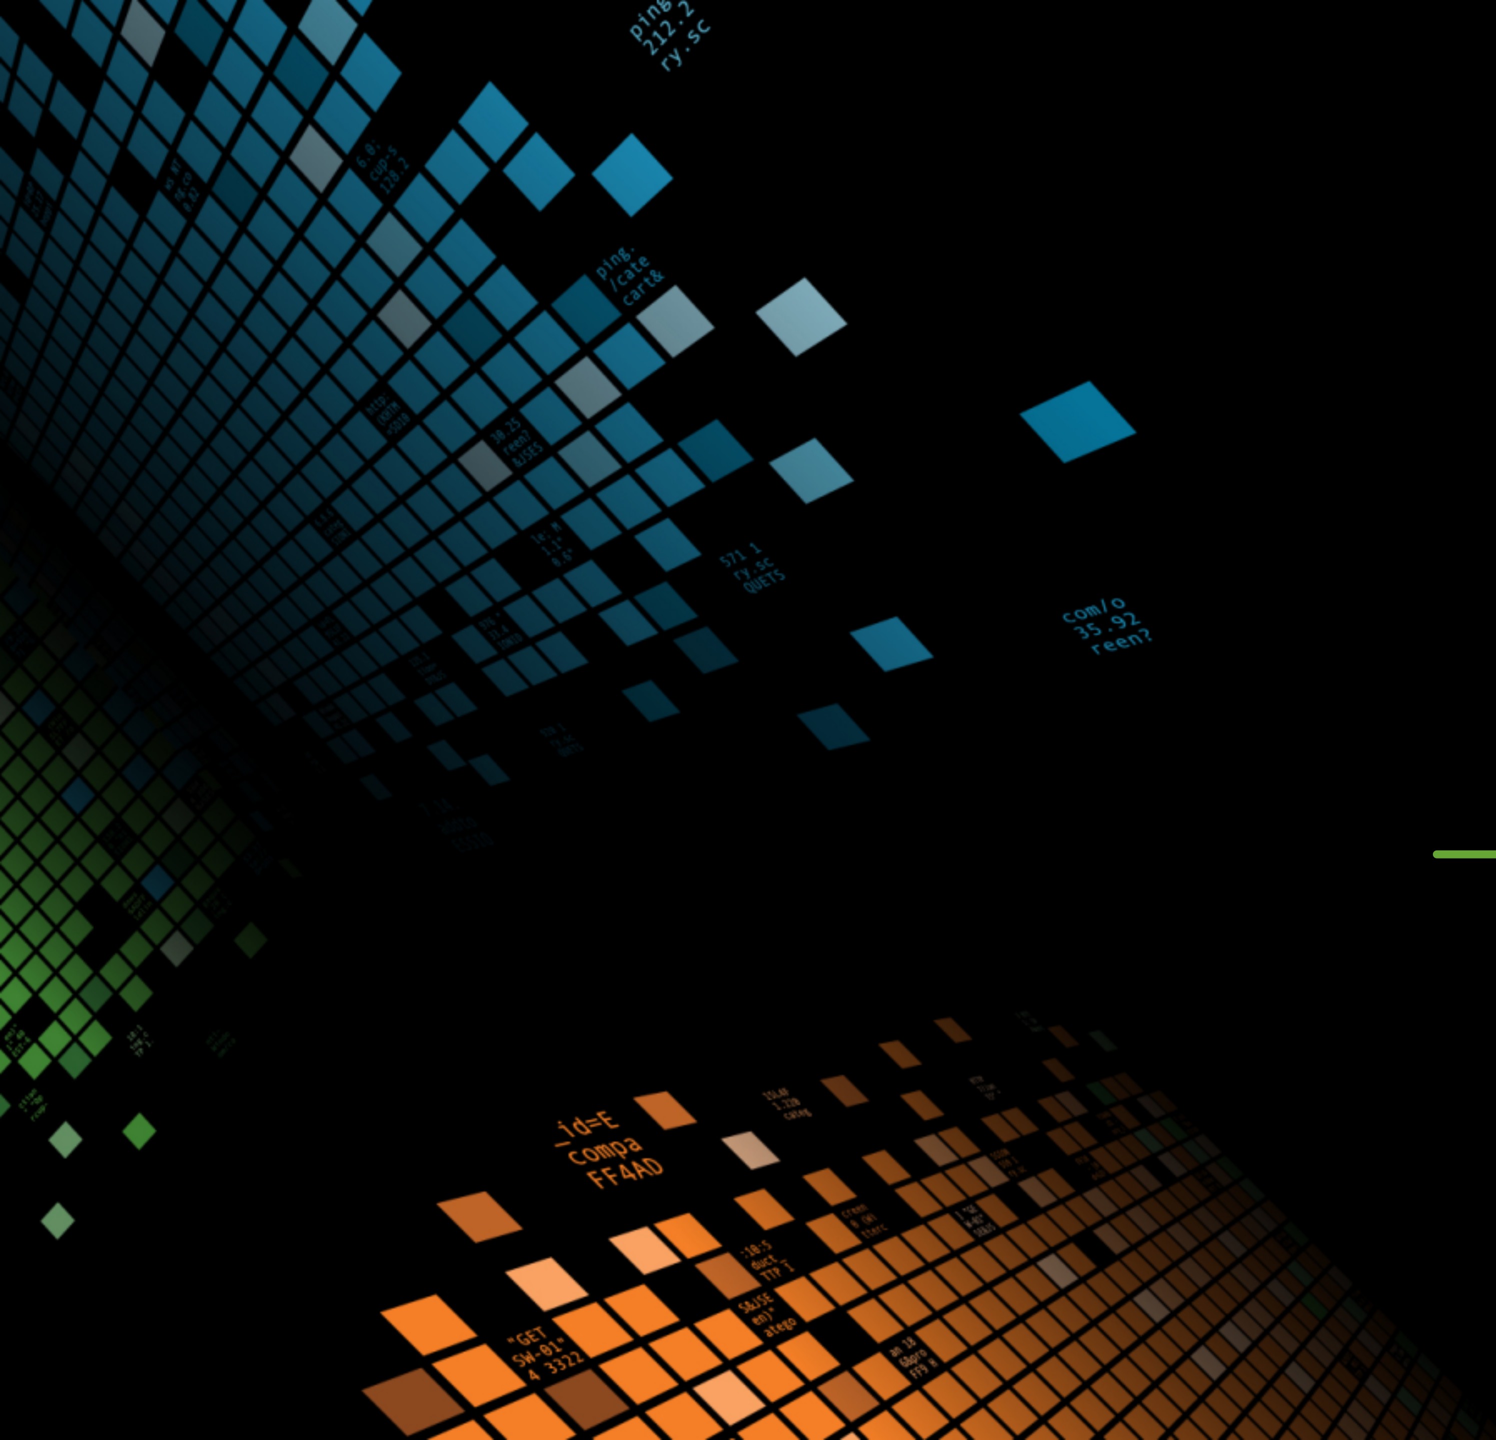
ping
212.2
ry.sc
ping.
/cate
cart&
571 1
ry.sc
QUETS
com/o
35.92
reen?
_id=E
compa
FF4AD
"GET
SW-01"
4 3322

Metrics

7.0

Metrics and Events

New, high-speed metrics engine that integrates seamlessly with events!

Metrics

- Set of numbers describing a particular process or activity measured over an interval of time - i.e., *time series data*
- Virtually unlimited number of use cases

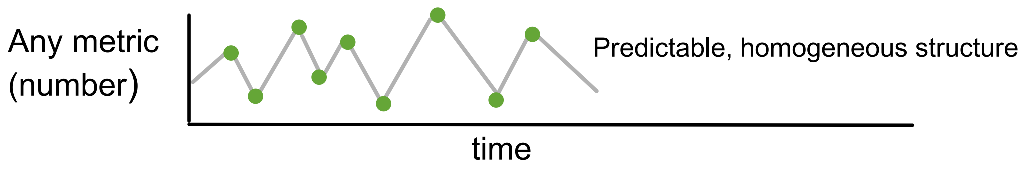

- Common metrics sources:
  - System metrics (CPU, memory, disk), Sensor data (temperature, …)
  - Infrastructure metrics (AWS CloudWatch), Web Tracking (Google Analytics)
  - Application agents (Application Performance Monitoring, error tracking)

Events

- Traditional Splunk - typically text, binary, un/structured data that describe a set of discrete events that happen over time
- Virtually unlimited number of use cases

Traditional Splunk

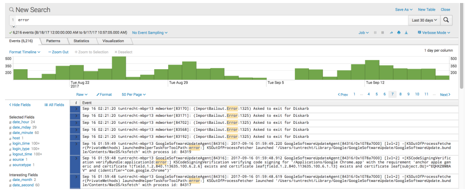

Search: error

✓ 6,216 events (8/18/17 12:00:00.000 AM to 9/17/17 10:57:05.000 AM)

Selected Fields: date_hour 24, date_mday 29, date_minute 60, host 1, login_time 100+, login_type 100+, logout_time 100+, source 1, sourcetype 1

Interesting Fields: date_month 2, date_second 60

```
Sep 16 02:21:20 tuntrecht-mbpr13 mdworker[83170]: (ImportBailout.Error:1325) Asked to exit for Diskarb
Sep 16 02:21:20 tuntrecht-mbpr13 mdworker[83171]: (ImportBailout.Error:1325) Asked to exit for Diskarb
Sep 16 02:21:20 tuntrecht-mbpr13 mdworker[84702]: (ImportBailout.Error:1325) Asked to exit for Diskarb
Sep 16 02:21:20 tuntrecht-mbpr13 mdworker[83568]: (ImportBailout.Error:1325) Asked to exit for Diskarb
Sep 16 02:21:20 tuntrecht-mbpr13 mdworker[83192]: (ImportBailout.Error:1325) Asked to exit for Diskarb
Sep 16 01:59:49 tuntrecht-mbpr13 GoogleSoftwareUpdateAgent[84316]: 2017-09-16 01:59:49.220 GoogleSoftwareUpdateAgent[84316/0x1078a7000] [lvl=2] -[KSOutOfProcessFetcher(PrivateMethods) launchedHelperTaskForToolPath:error:] KSOutOfProcessFetcher launched '/Users/tuntrecht/Library/Google/GoogleSoftwareUpdate/GoogleSoftwareUpdate.bundle/Contents/MacOS/ksfetch' with process id: 84319
Sep 16 01:59:48 tuntrecht-mbpr13 GoogleSoftwareUpdateAgent[84316]: 2017-09-16 01:59:48.912 GoogleSoftwareUpdateAgent[84316/0x1078a7000] [lvl=2] +[KSCodeSigningVerification verifyBundle:applicationId:error:] KSCodeSigningVerification verifying code signing for '/Applications/Google Chrome.app' with the requirement 'anchor apple generic and certificate 1[field.1.2.840.113635.100.6.2.6] exists and certificate leaf[field.1.2.840.113635.100.6.1.13] exists and certificate leaf[subject.OU]="EQHXZ8M8AV" and (identifier="com.google.Chrome")'
Sep 16 01:59:48 tuntrecht-mbpr13 GoogleSoftwareUpdateAgent[84316]: 2017-09-16 01:59:48.619 GoogleSoftwareUpdateAgent[84316/0x1078a7000] [lvl=2] -[KSOutOfProcessFetcher(PrivateMethods) launchedHelperTaskForToolPath:error:] KSOutOfProcessFetcher launched '/Users/tuntrecht/Library/Google/GoogleSoftwareUpdate/GoogleSoftwareUpdate.bundle/Contents/MacOS/ksfetch' with process id: 84317
```

- Common event sources:
  - System and server logs (syslog, journald), APIs (Twitter, Wunderground)
  - Firewall data (Palo Alto Networks, etc)
  - Application, platform and other logs (log4j, log4net, Apache, MySQL, AWS)

Sample Metric

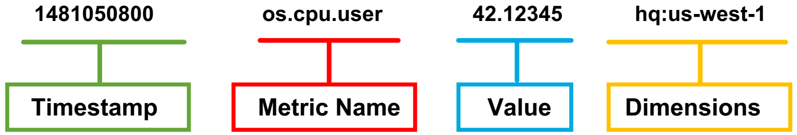

1481050800	os.cpu.user	42.12345	hq:us-west-1
Timestamp	Metric Name	Value	Dimensions

Sample Log

```
[29/Aug/2017 08:47:05:316503] "POST /                a31d69&action=remove&&product_id=BS-
2&JSESSIONID=SD6SAL4FF1ADFF9 H                //www.buttercupenterprises.com/product.screen?
product_id=BS-2" "Mozilla/5.0 (Intel Ma                WebKit/537.36 (KHTML, like Gecko)
Chrome/57.0.2957.0 Safari/537.36" 98
```

Of course can also contain numbers

New Metrics Features in the Platform

Ingest and Query metrics natively with new, high speed metrics engine

SPL

Metric Store

New metrics engine and index to ingest and store metrics *natively* at scale

`mstats`

Query metrics data from metrics indexes (similar to `tstats`)

Metrics Catalog

REST APIs and `mcatalog` to query lists of ingested metrics and dimensions

7.0

Splunk Metrics – Index Creation

New Metrics Index Data Type

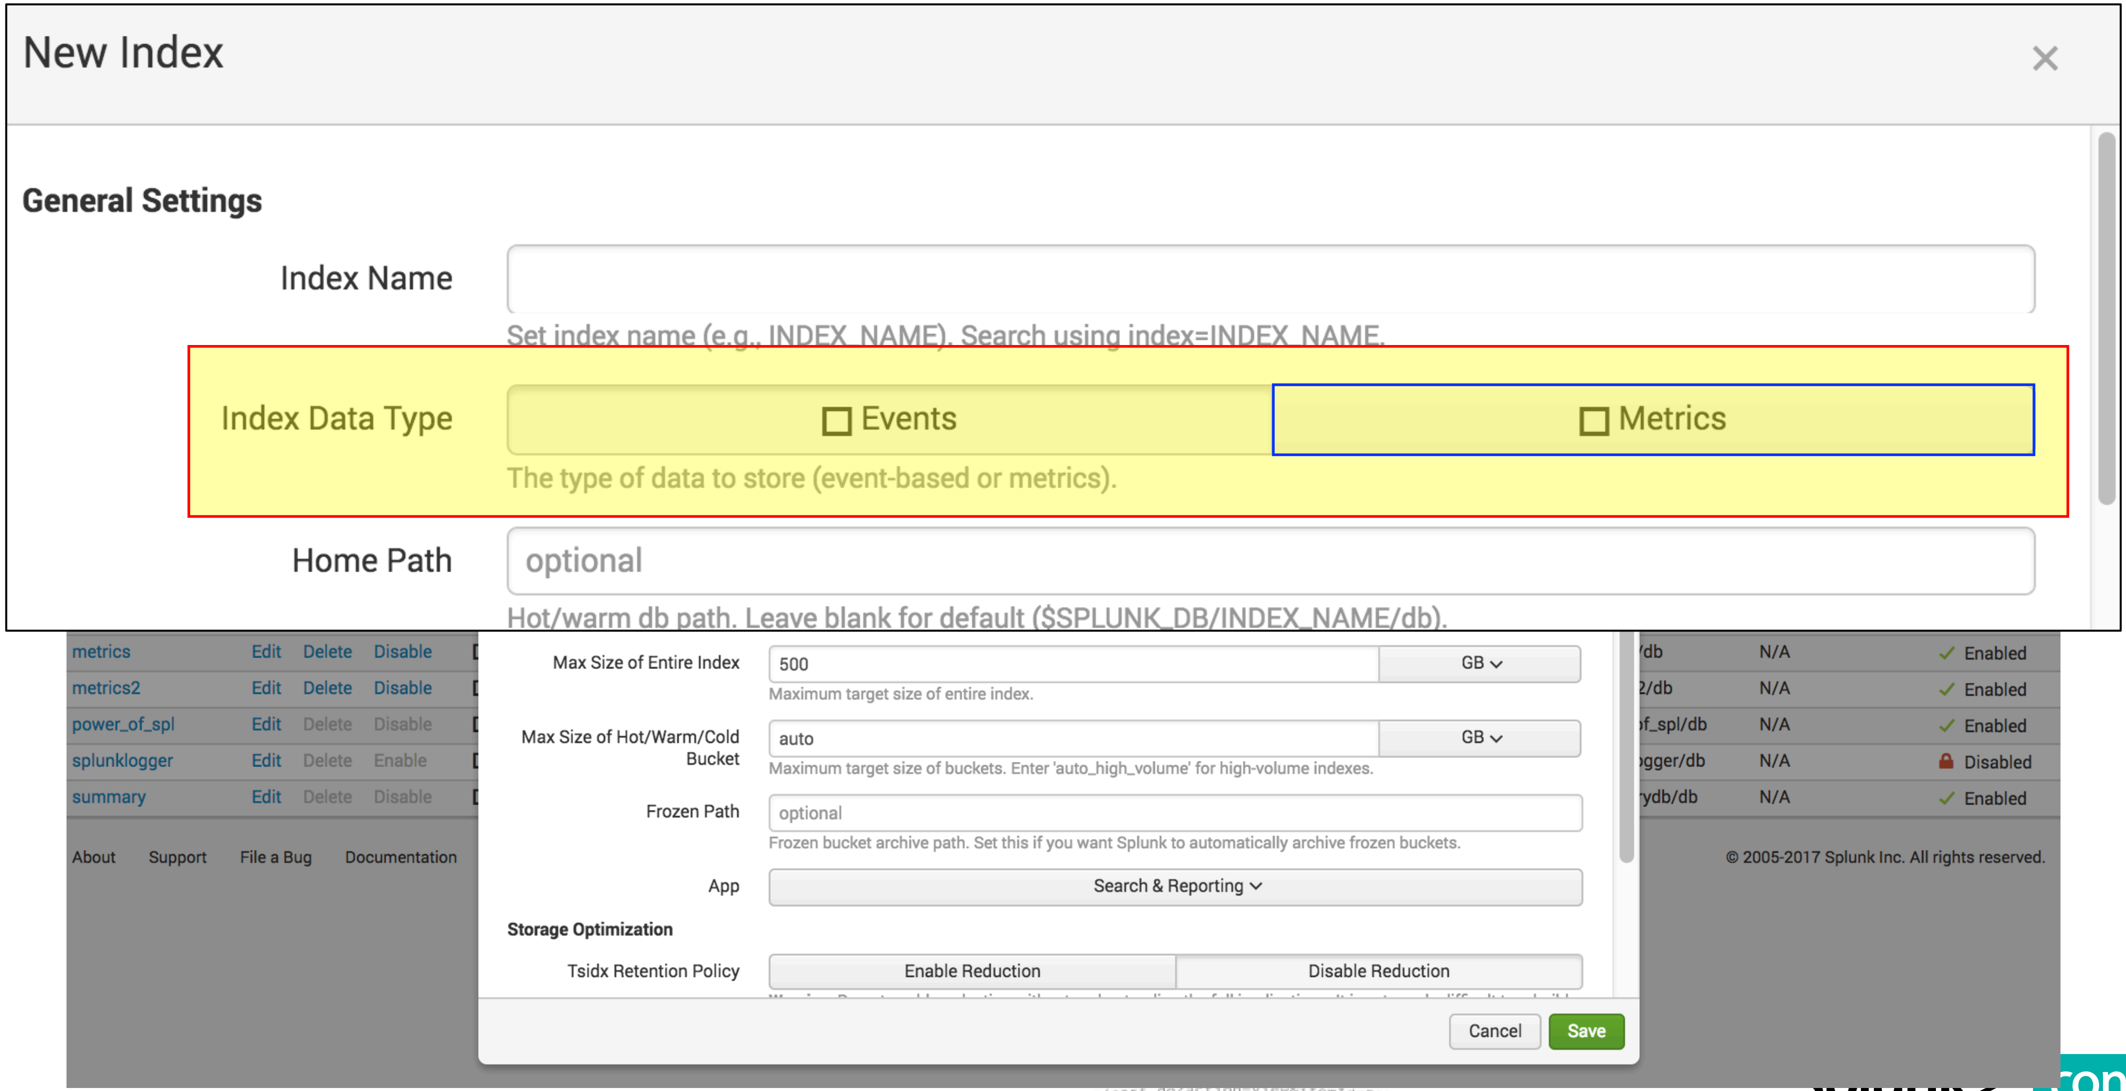

New Index ×

General Settings

Index Name

Set index name (e.g., INDEX_NAME). Search using index=INDEX_NAME.

Index Data Type □ Events □ Metrics

The type of data to store (event-based or metrics).

Home Path: optional

Hot/warm db path. Leave blank for default ($SPLUNK_DB/INDEX_NAME/db).

Max Size of Entire Index: 500 GB

Maximum target size of entire index.

Max Size of Hot/Warm/Cold Bucket: auto GB

Maximum target size of buckets. Enter 'auto_high_volume' for high-volume indexes.

Frozen Path: optional

Frozen bucket archive path. Set this if you want Splunk to automatically archive frozen buckets.

App: Search & Reporting

Storage Optimization

Tsidx Retention Policy: Enable Reduction | Disable Reduction

Cancel | Save

metrics	Edit	Delete	Disable	/db	N/A	✓ Enabled
metrics2	Edit	Delete	Disable	2/db	N/A	✓ Enabled
power_of_spl	Edit	Delete	Disable	f_spl/db	N/A	✓ Enabled
splunklogger	Edit	Delete	Enable	gger/db	N/A	🔒 Disabled
summary	Edit	Delete	Disable	ydb/db	N/A	✓ Enabled

About | Support | File a Bug | Documentation

© 2005-2017 Splunk Inc. All rights reserved.

7.0

Splunk Metrics – Data Ingestion

Supports StatsD, CollectD, and CSV files natively

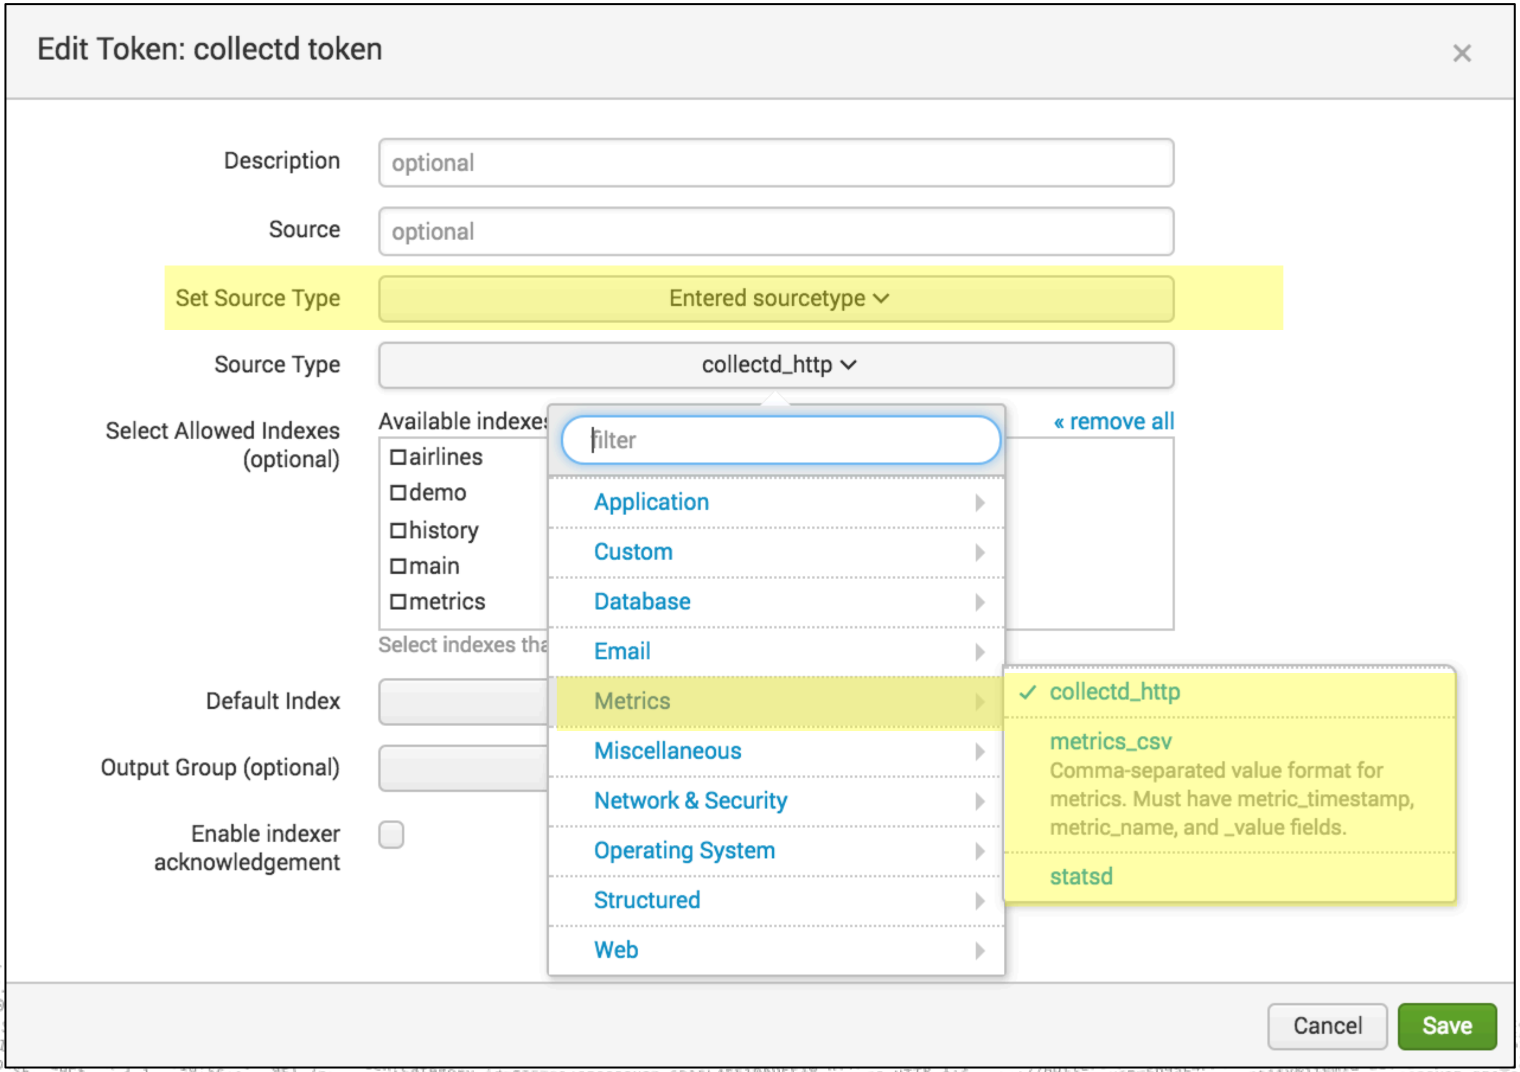

Edit Token: collectd token

Description: optional

Source: optional

Set Source Type: Entered sourcetype

Source Type: collectd_http

Select Allowed Indexes (optional)

Available indexes

□airlines
□demo
□history
□main
□metrics

Select indexes tha

Default Index

Output Group (optional)

Enable indexer acknowledgement □

✓ collectd_http

metrics_csv
Comma-separated value format for metrics. Must have metric_timestamp, metric_name, and _value fields.

statsd

Supports other Metrics Clients via manual configuration

```
# props.conf.example

[graphite_plaintext]
TIME_PREFIX = \s(\d{0,10})$
TIME_FORMAT =  %s
NO_BINARY_CHECK = true
SHOULD_LINEMERGE = false
pulldown_type = 1
TRANSFORMS-graphite-host = graphite_host
TRANSFORMS-graphite-metricname = graphite_metric_name
TRANSFORMS-graphite-metricvalue = graphite_metric_value
category = Metrics
```

DEMO - Splunk Metrics

What is the CPU impact of plugging in my phone and playing a video?

Metrics

- Ingest CPU metrics data from my laptop
  (using CollectD into HTTP Event Collector)
- Take a video of all you (everyone wave)!
- Plug in my phone, Play the video
- Watch what happens to my CPU

Events

- Monitor system.log for events that signal I plugged in my phone
- Look for events with “connected apple”
- Run the query
- Save the query to use later

Combine Metrics and Event data
using new Event Annotation

7.0

Splunk Metrics

Taking the *meh* out of metrics

- New, high speed metrics engine and index
- Get all the benefits of the Splunk platform:
  - Visualizations and alerting
  - Easy Data ingestion
  - Clustering, Scaling
  - Leverage open source for existing sourcetypes (*statsd, collectd*)
- **20x** and beyond performance improvement for monitoring and alerting use cases with metrics data
- Integrates beautifully with traditional Splunk events

Metrics–great for high volume data, large # of searches in one dashboard

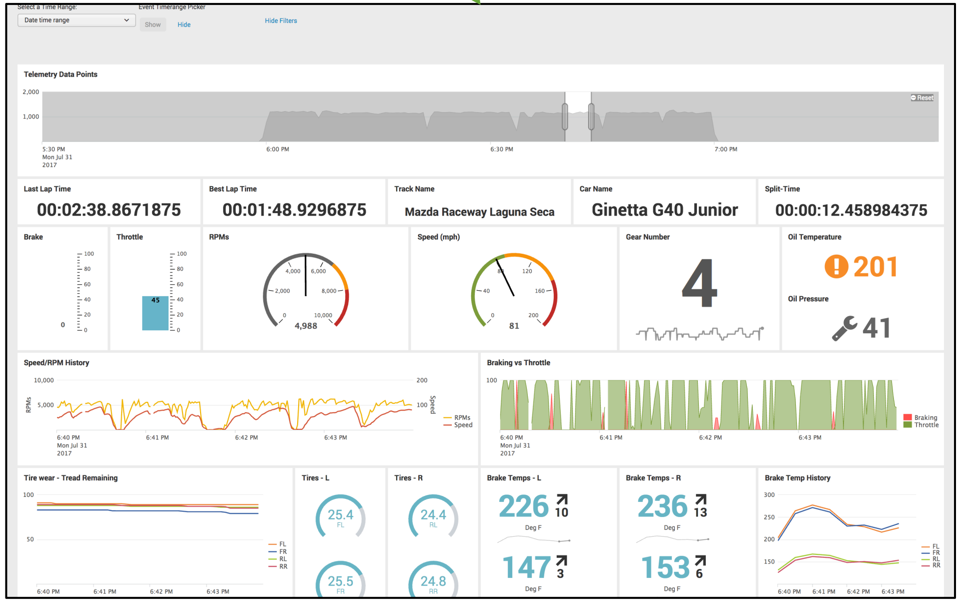

Select a Time Range: Date time range

Event Timerange Picker

Last Lap Time: 00:02:38.8671875

Best Lap Time: 00:01:48.9296875

Track Name: Mazda Raceway Laguna Seca

Car Name: Ginetta G40 Junior

Split-Time: 00:00:12.458984375

…and we're just getting started…lots more to come!

7.0

Splunk Event Annotation

Surfaces more insights from your data

- Doesn’t just correlate logs and metrics in a single view
- Adds context to any time chart (e.g., line, column, area)
- Under the covers, it’s just an embedded search that enables you to pull markers and labels from many different sources
  - (e.g., log data, lookup files, or external sources)

Event Annotation

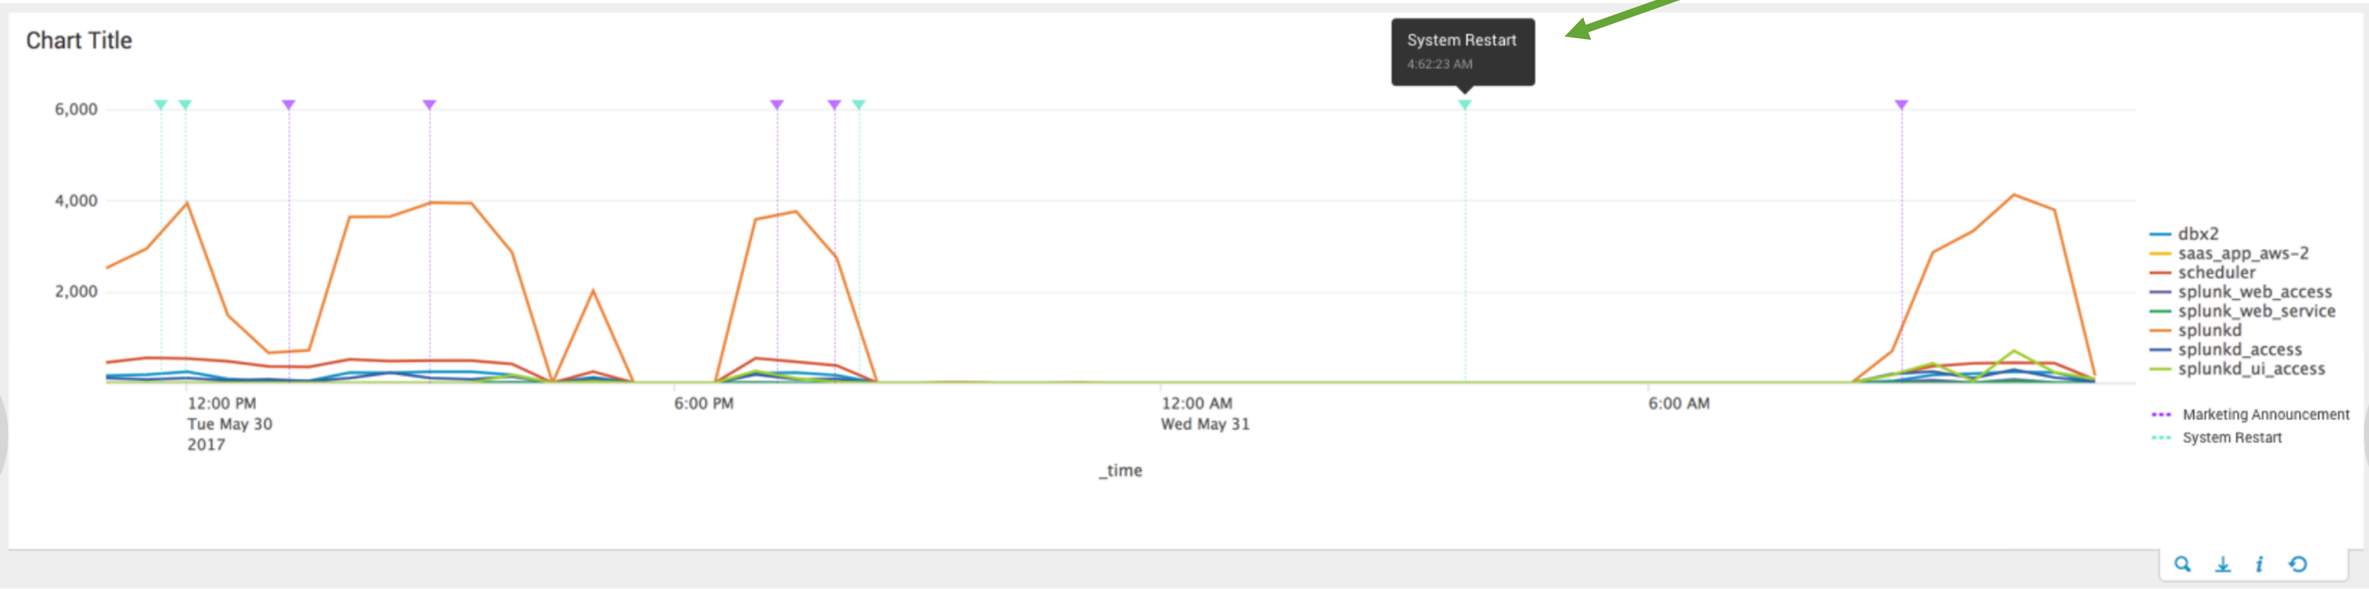

Chart Enhancements & Report Actions

7.0

DEMO

Chart Enhancements

Perfect for metric and multi-series monitoring use cases

Change Line Width

View and Compare Values in a Chart Legend

Change Line Styles (dashes, dots etc)

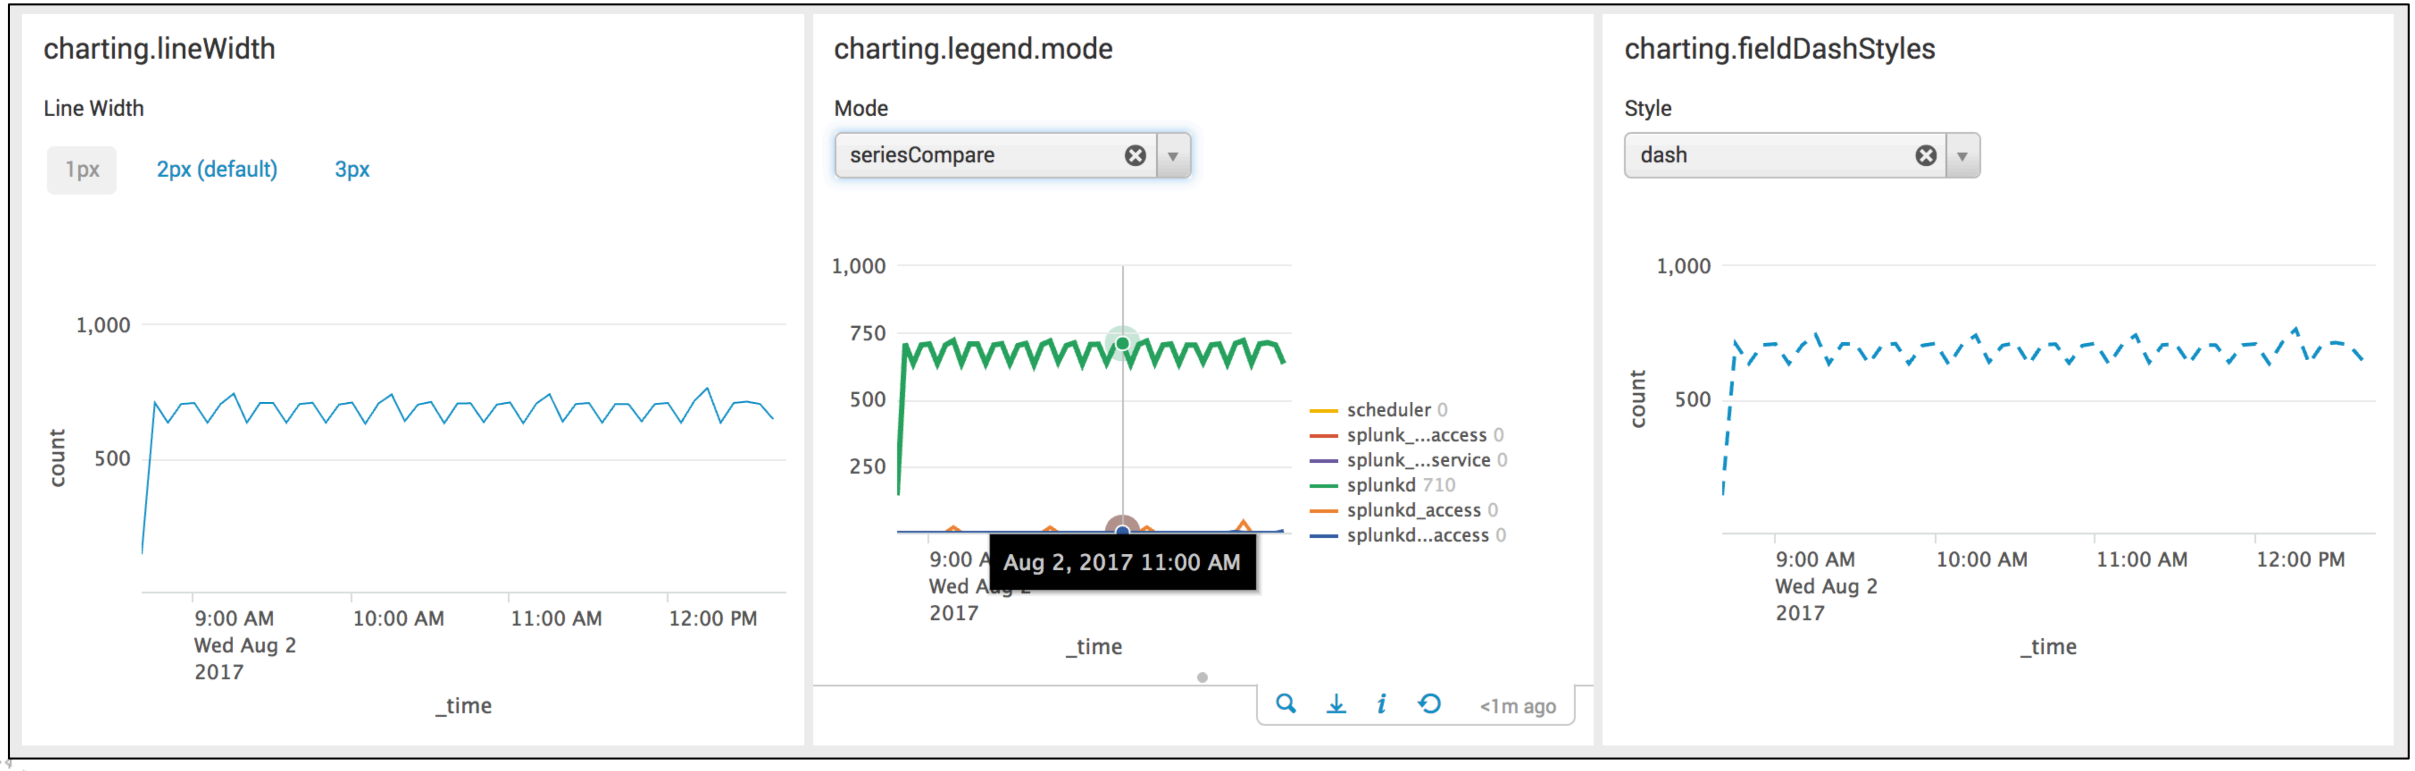

charting.lineWidth

Line Width: 1px | 2px (default) | 3px

charting.legend.mode

Mode: seriesCompare

charting.fieldDashStyles

Style: dash

Report Actions

Ability to set any action when scheduling a report

- Use any action that is installed and enabled in your alert actions framework

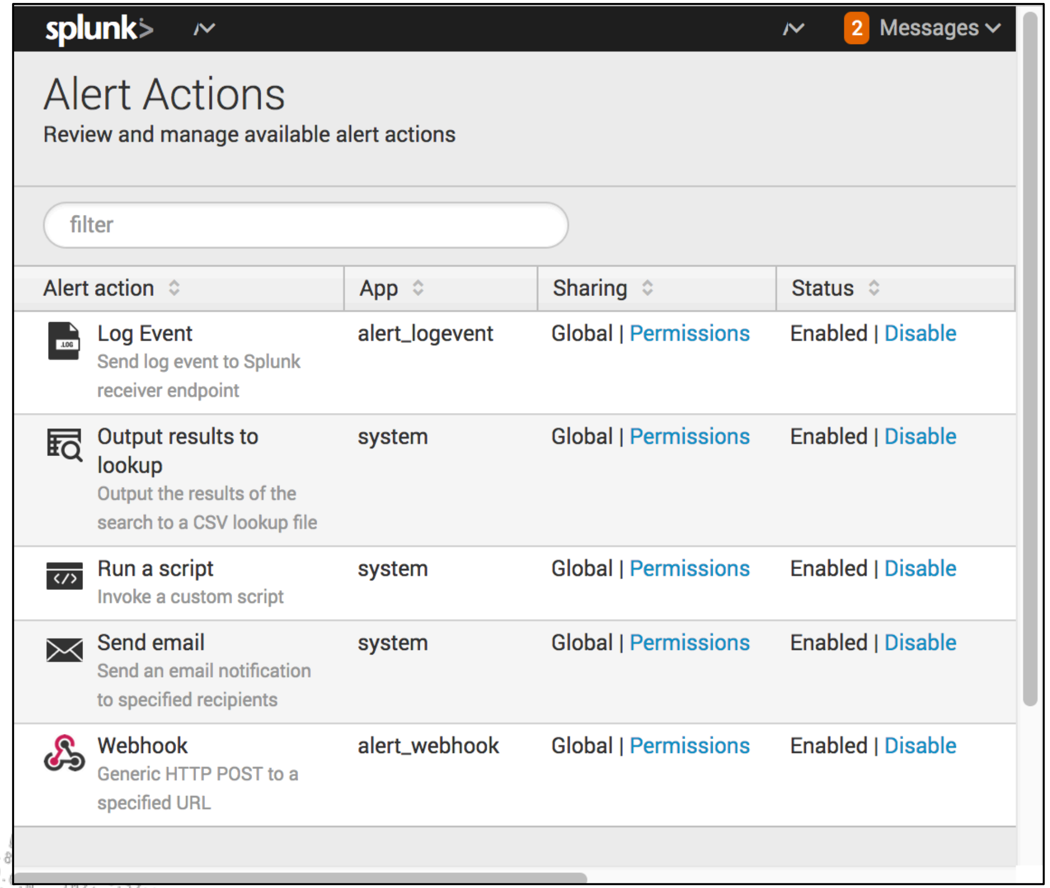

Alert Actions

Review and manage available alert actions

filter

Alert action	App	Sharing	Status
Log Event Send log event to Splunk receiver endpoint	alert_logevent	Global | Permissions	Enabled | Disable
Output results to lookup Output the results of the search to a CSV lookup file	system	Global | Permissions	Enabled | Disable
Run a script Invoke a custom script	system	Global | Permissions	Enabled | Disable
Send email Send an email notification to specified recipients	system	Global | Permissions	Enabled | Disable
Webhook Generic HTTP POST to a specified URL	alert_webhook	Global | Permissions	Enabled | Disable

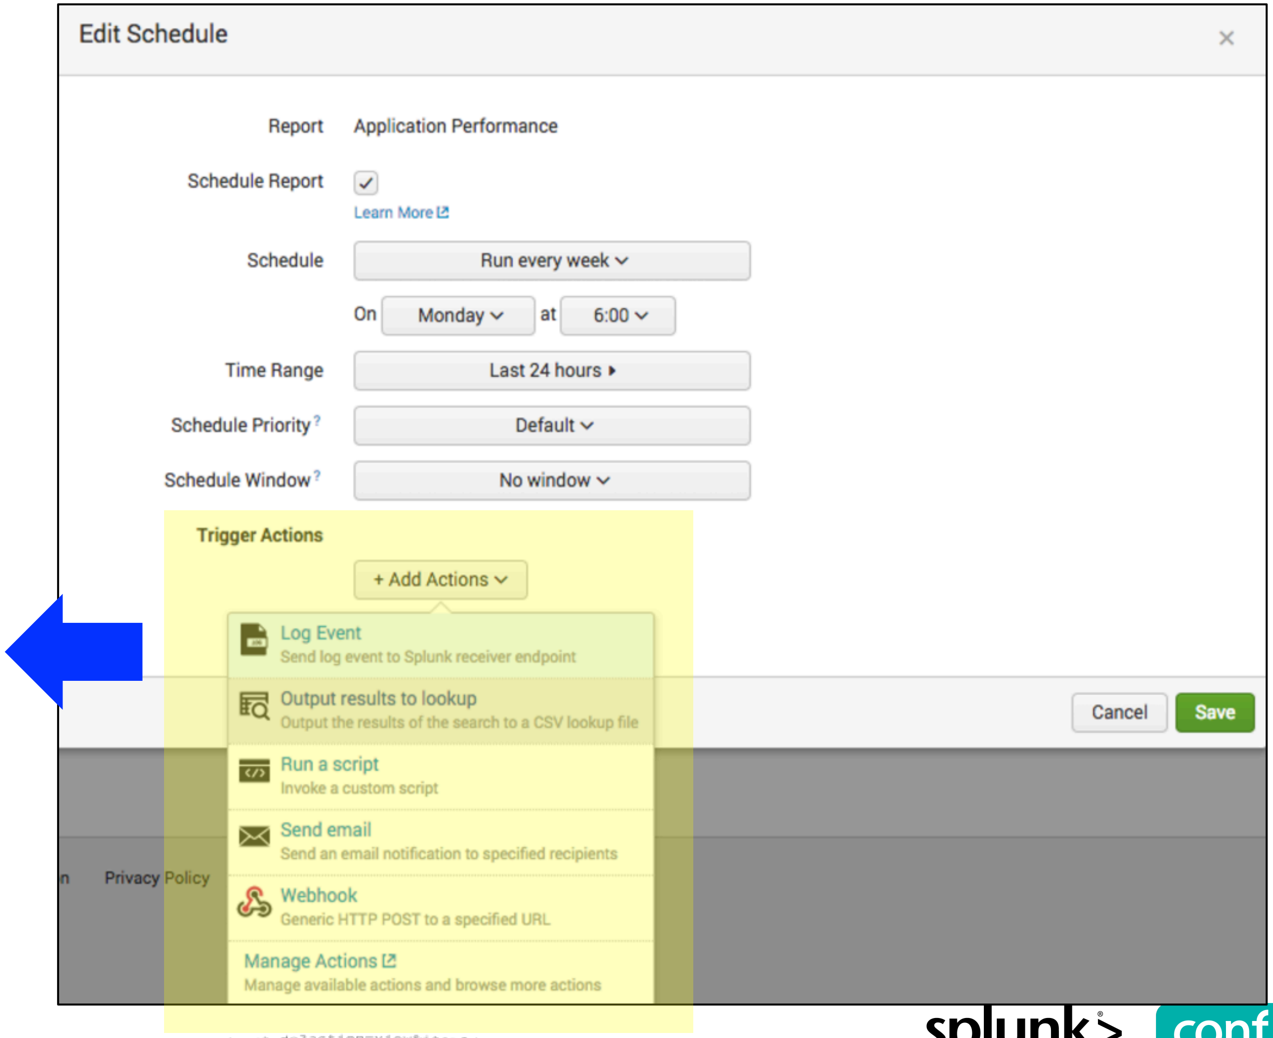

Edit Schedule

Report: Application Performance

Schedule Report: ☑ Learn More

Schedule: Run every week

On Monday at 6:00

Time Range: Last 24 hours

Schedule Priority: Default

Schedule Window: No window

Trigger Actions

+ Add Actions

- Log Event: Send log event to Splunk receiver endpoint
- Output results to lookup: Output the results of the search to a CSV lookup file
- Run a script: Invoke a custom script
- Send email: Send an email notification to specified recipients
- Webhook: Generic HTTP POST to a specified URL
- Manage Actions: Manage available actions and browse more actions

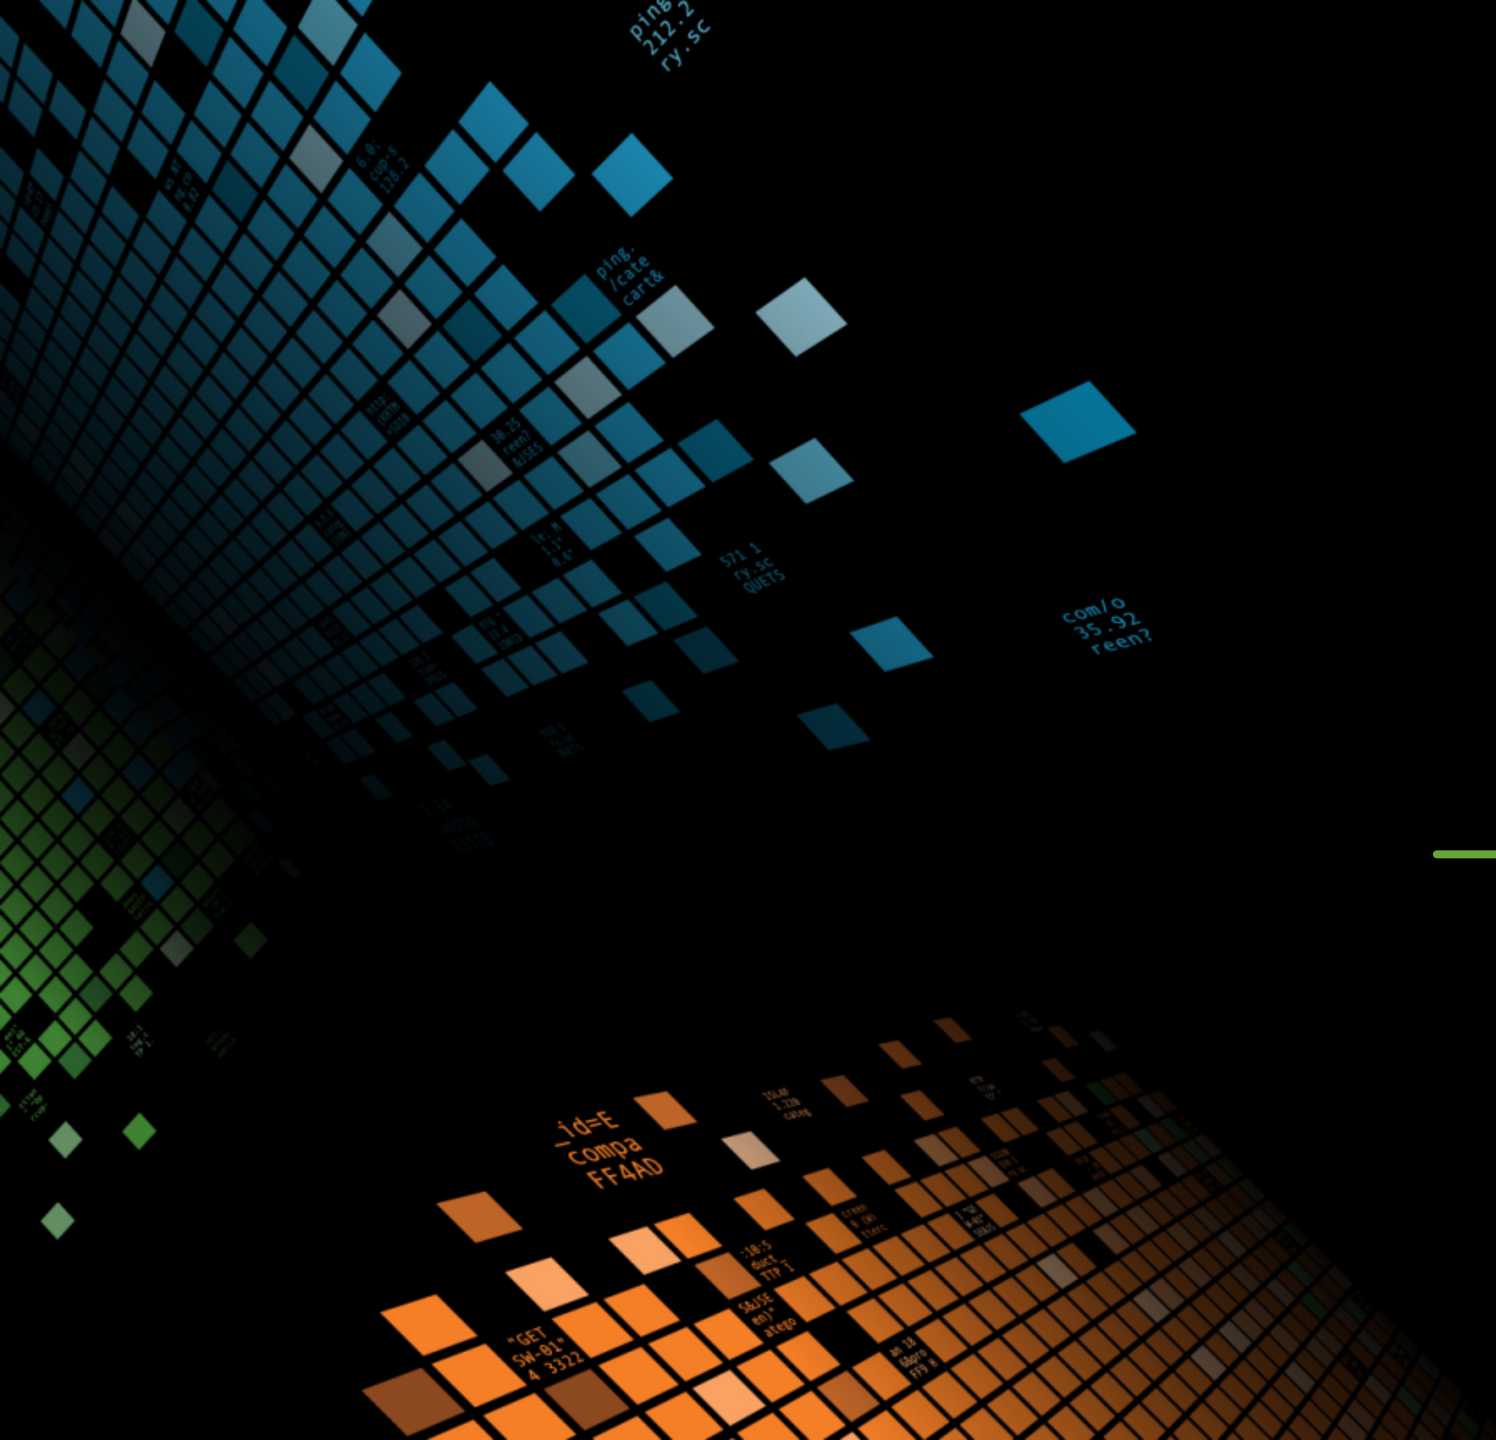

Trellis Layout

6.6

Trellis Layout

Go from one chart to multiple charts with a single click of a button!

- Perfect for comparing information across dimensions
- Super easy to use via new Trellis icon directly in the UI
- Super efficient at runtime as it executes the base search only once
- Uses results from `chart`, `stats`, and `timechart` commands
- Works in Search, Reports, Dashboards

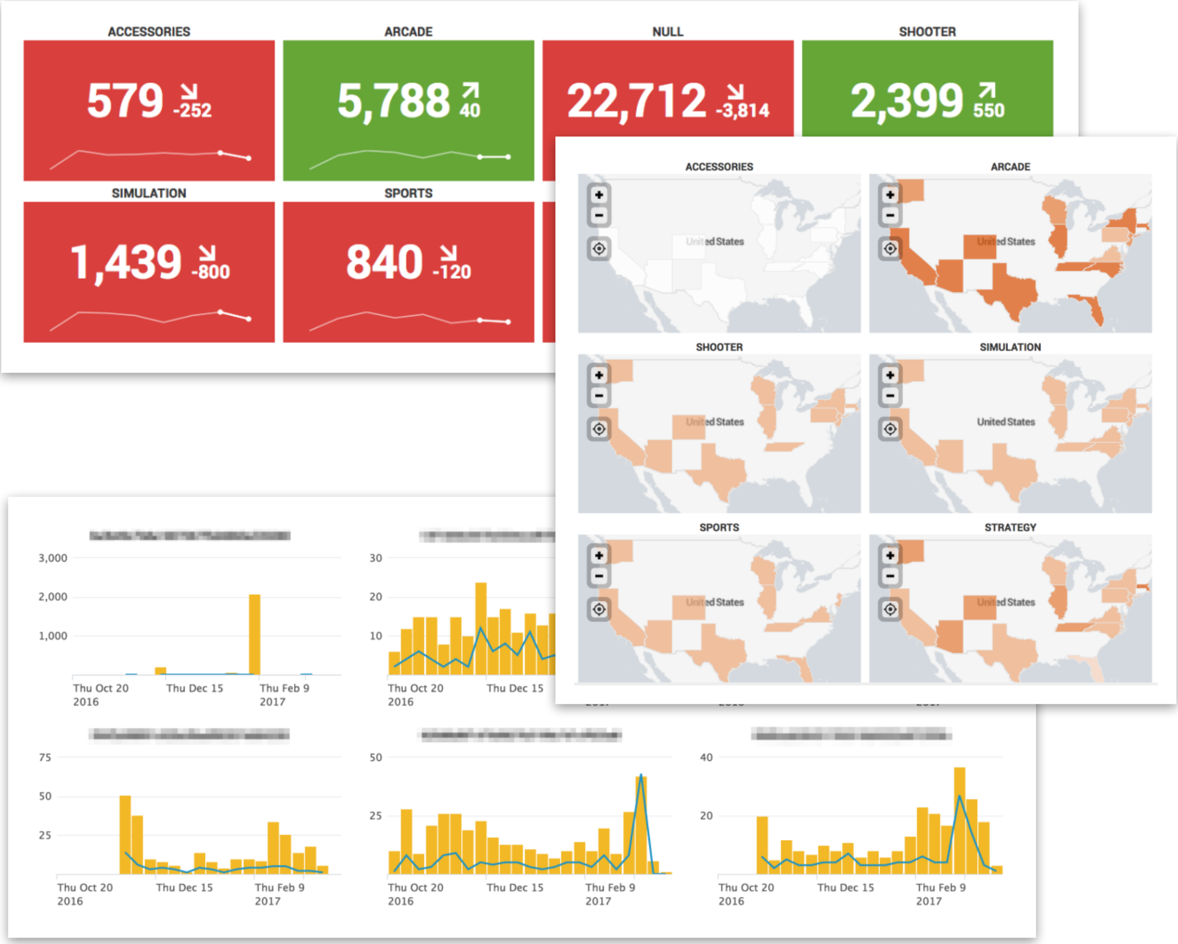

splunk> .conf2017

DEMO - Trellis Layout

Buttercup Games - Show me the count of all website actions across product categories

```
buttercupgames action!=NULL | chart count by categoryId, action
```

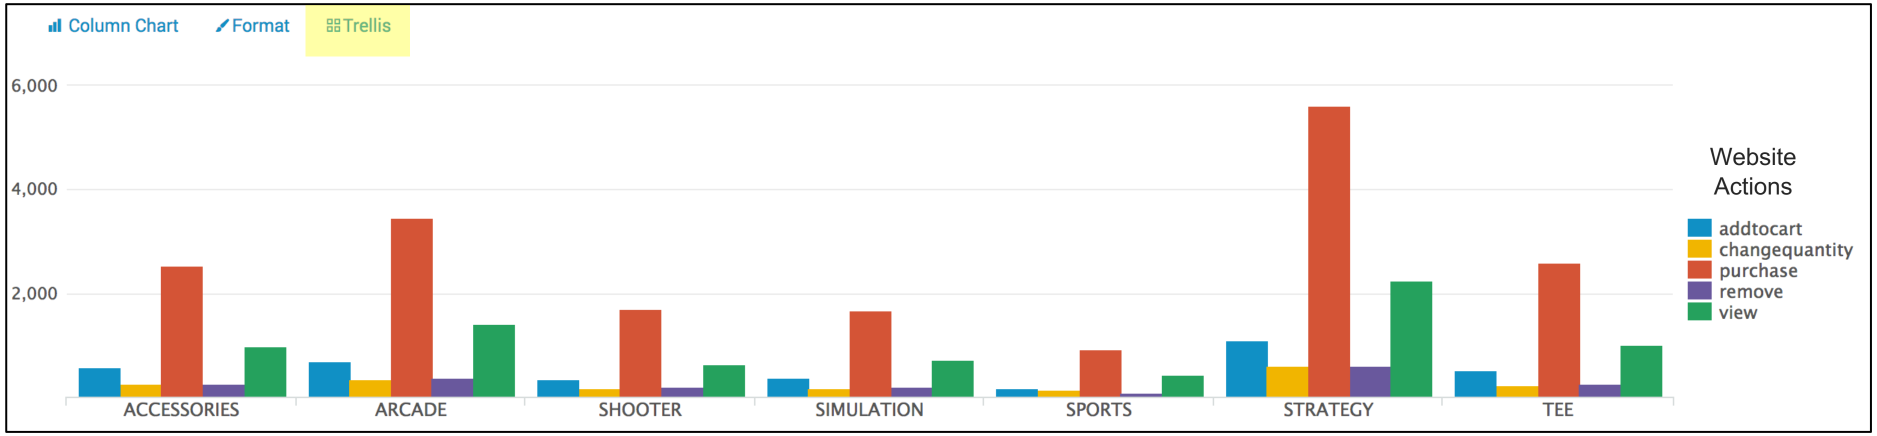

- But what if I want to see multiple charts and compare across website actions?
- And I don’t want to manually create new queries and charts….
- Simple example of how Trellis can help!

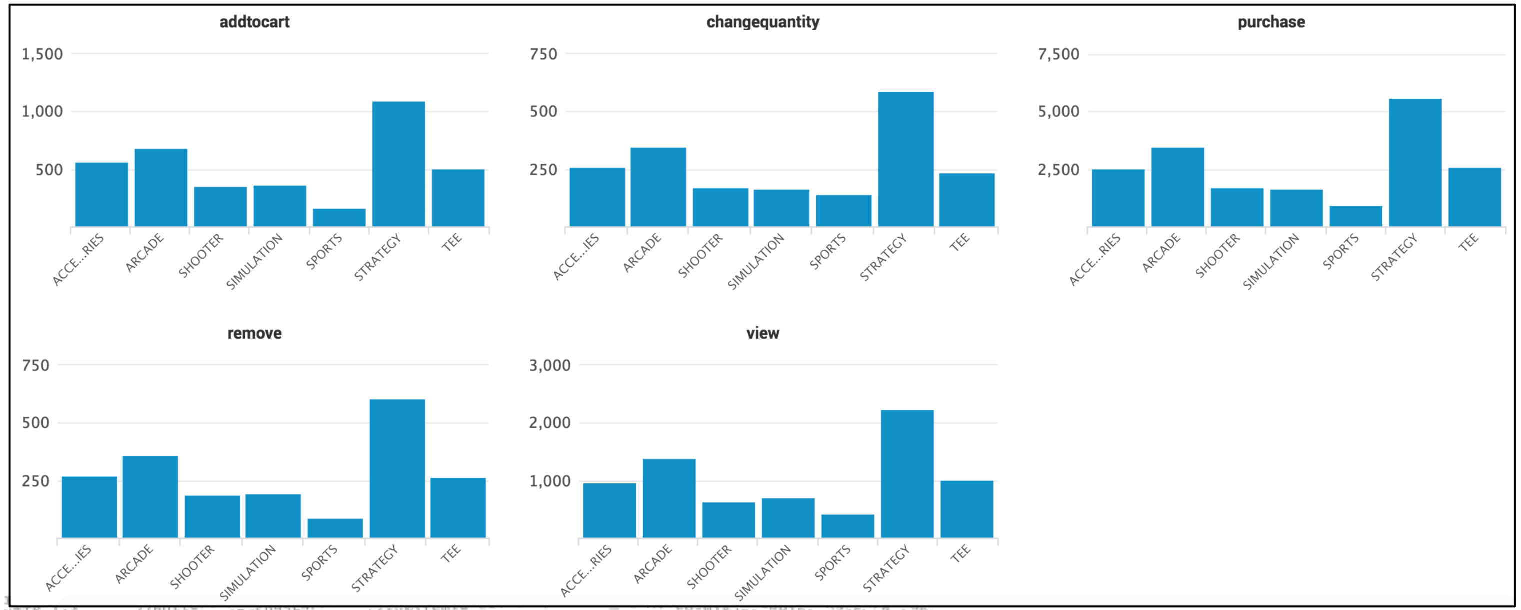

6.6

Trellis Example

Splunk Blog – Episode 6

May 2017, Stephen Luedtke, Splunk

Traveling on Time with Trellis

- Airlines with the most cancelled flights (vs previous week)
- Airlines with the most delays (vs previous week)
- Leveraging simple Single Value visualization

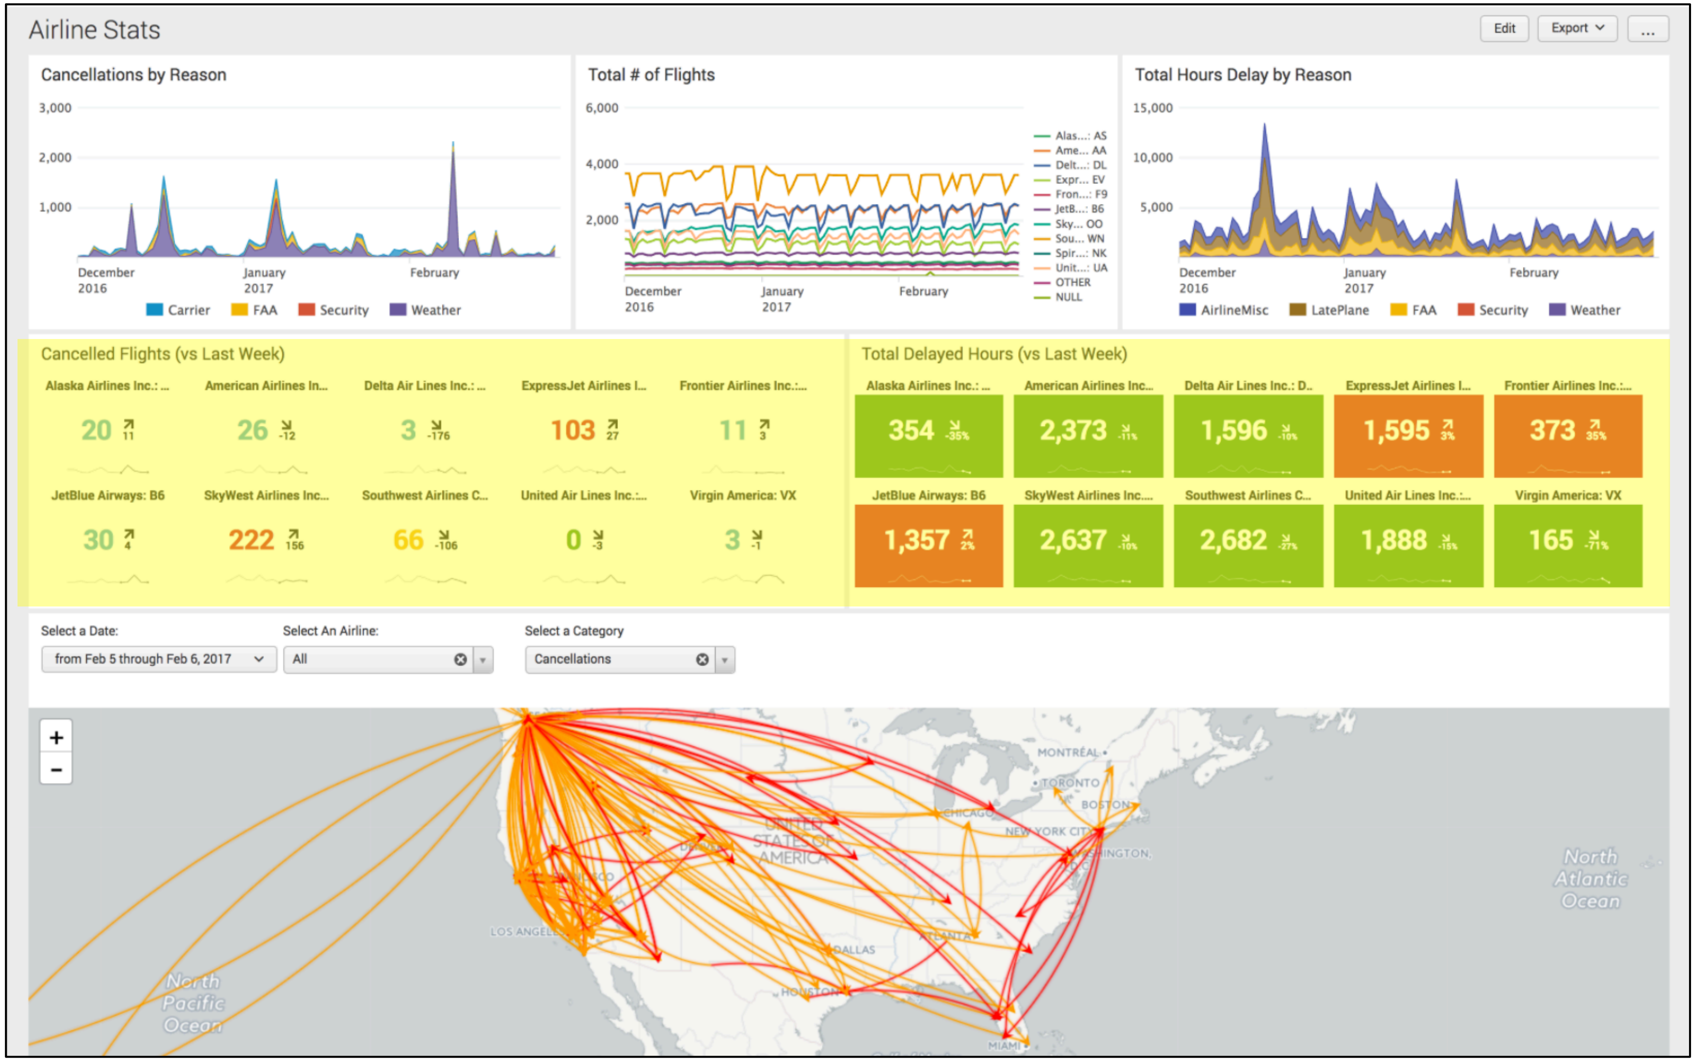

Airline Stats

Cancelled Flights (vs Last Week)

Alaska Airlines Inc.: ... 20 ↗ 11
American Airlines In... 26 ↘ -12
Delta Air Lines Inc.: ... 3 ↘ -176
ExpressJet Airlines I... 103 ↗ 27
Frontier Airlines Inc.:... 11 ↗ 3
JetBlue Airways: B6 30 ↗ 4
SkyWest Airlines Inc... 222 ↗ 156
Southwest Airlines C... 66 ↘ -106
United Air Lines Inc.:... 0 ↘ -3
Virgin America: VX 3 ↘ -1

Total Delayed Hours (vs Last Week)

Alaska Airlines Inc.: ... 354 ↘ -35%
American Airlines Inc... 2,373 ↘ -11%
Delta Air Lines Inc.: D.. 1,596 ↘ -10%
ExpressJet Airlines I... 1,595 ↗ 3%
Frontier Airlines Inc.:... 373 ↗ 35%
JetBlue Airways: B6 1,357 ↗ 2%
SkyWest Airlines Inc.... 2,637 ↘ -10%
Southwest Airlines C... 2,682 ↘ -27%
United Air Lines Inc.:... 1,888 ↘ -15%
Virgin America: VX 165 ↘ -71%

Select a Date: from Feb 5 through Feb 6, 2017

Select An Airline: All

Select a Category: Cancellations

Trellis Contest Winner

Matt Zerfas, SONIFI Solutions

- Shows Wi-Fi Signal Strength across 40 floors of a hotel
- Uses Trellis with a customer choropleth map
- Measures wi-fi signal strength in dBm (decibel milliwatts)

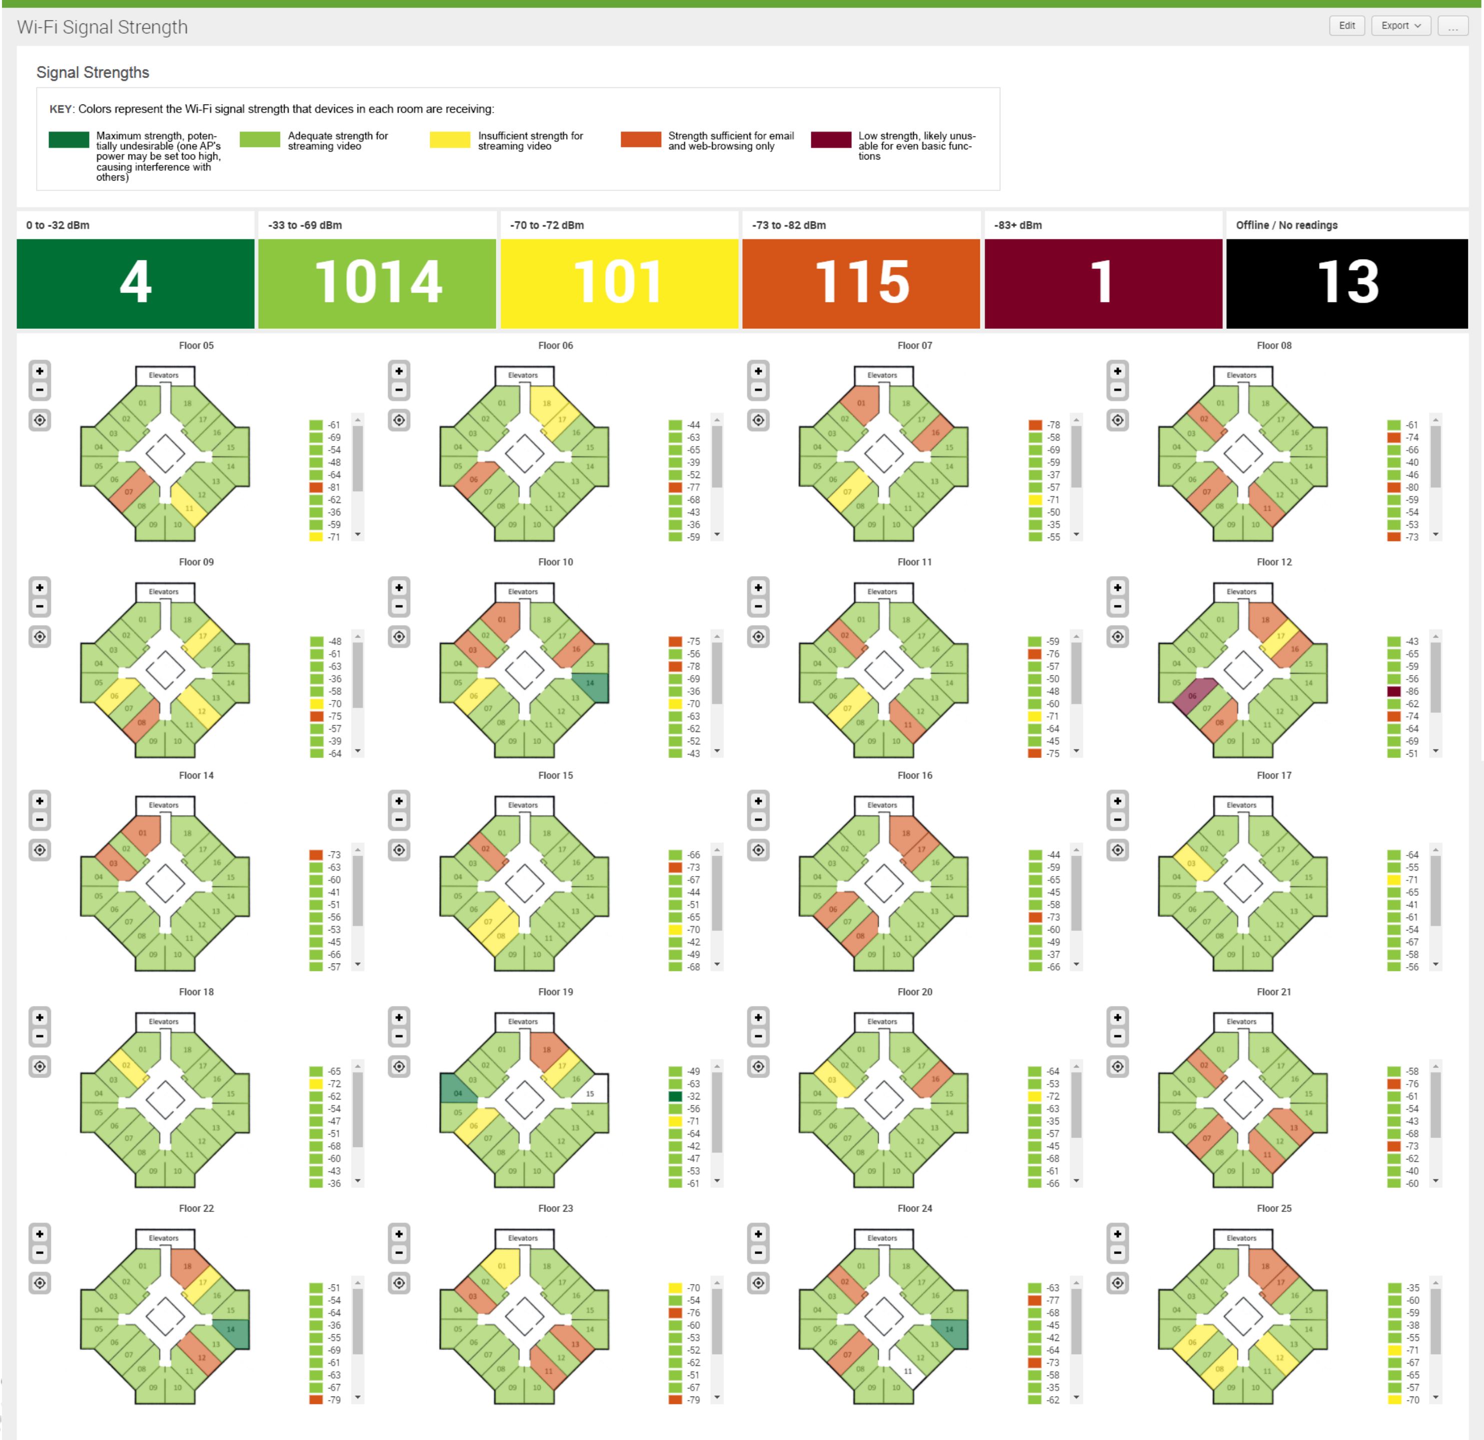

Wi-Fi Signal Strength

Signal Strengths

KEY: Colors represent the Wi-Fi signal strength that devices in each room are receiving:

🟩	Maximum strength, potentially undesirable (one AP's power may be set too high, causing interference with others)
🟩	Adequate strength for streaming video
🟨	Insufficient strength for streaming video
🟧	Strength sufficient for email and web-browsing only
🟥	Low strength, likely unusable for even basic functions

0 to -32 dBm	-33 to -69 dBm	-70 to -72 dBm	-73 to -82 dBm	-83+ dBm	Offline / No readings
4	1014	101	115	1	13

Floor 05: -61 -69 -54 -48 -64 -81 -62 -36 -59 -71

Floor 06: -44 -63 -65 -39 -52 -77 -68 -43 -36 -59

Floor 07: -78 -58 -69 -59 -37 -57 -71 -50 -35 -55

Floor 08: -61 -74 -66 -40 -46 -80 -59 -54 -53 -73

Floor 09: -48 -61 -63 -36 -58 -70 -75 -57 -39 -64

Floor 10: -75 -56 -78 -69 -36 -70 -63 -62 -52 -43

Floor 11: -59 -76 -57 -50 -48 -60 -71 -64 -45 -75

Floor 12: -43 -65 -59 -56 -86 -62 -74 -64 -69 -51

Floor 14: -73 -63 -60 -41 -51 -56 -53 -45 -66 -57

Floor 15: -66 -73 -67 -44 -51 -65 -70 -42 -49 -68

Floor 16: -44 -59 -65 -45 -58 -73 -60 -49 -37 -66

Floor 17: -64 -55 -71 -65 -41 -61 -54 -67 -58 -56

Floor 18: -65 -72 -62 -54 -47 -51 -68 -60 -43 -36

Floor 19: -49 -63 -32 -56 -71 -64 -42 -47 -53 -61

Floor 20: -64 -53 -72 -63 -35 -57 -45 -68 -61 -66

Floor 21: -58 -76 -61 -54 -43 -68 -73 -62 -40 -60

Floor 22: -51 -54 -64 -36 -55 -69 -61 -63 -67 -79

Floor 23: -70 -54 -76 -60 -53 -52 -62 -51 -67 -79

Floor 24: -63 -77 -68 -45 -42 -64 -73 -58 -35 -62

Floor 25: -35 -60 -59 -38 -55 -71 -67 -65 -57 -70

Dashboard, Search, and Datasets

User Productivity Improvements

splunk> .conf2017

6.6

New Dashboard Drilldown Editor

Build custom interactivity in your dashboards without the need to learn XML

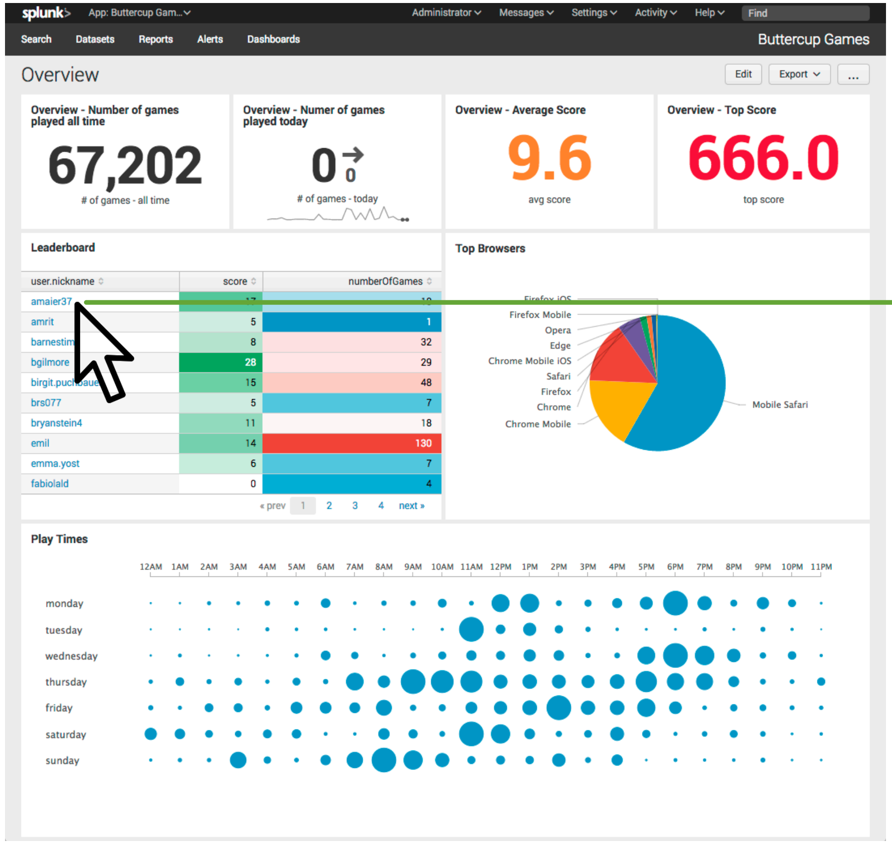

Leaderboard

user.nickname	score	numberOfGames
amaier37	[illegible]	[illegible]
amrit	5	1
barnestim	8	32
bgilmore	28	29
birgit.puchbaue	15	48
brs077	5	7
bryanstein4	11	18
emil	14	130
emma.yost	6	7
fabiolald	0	4

« prev 1 2 3 4 next »

- Link to search
- Link to dashboard
- Link to report
- Link to custom URL
- Manage tokens

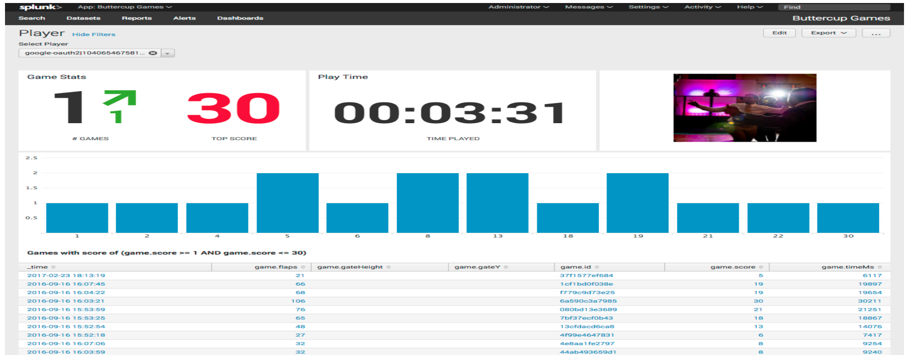

Using a simple gui-based drilldown editor

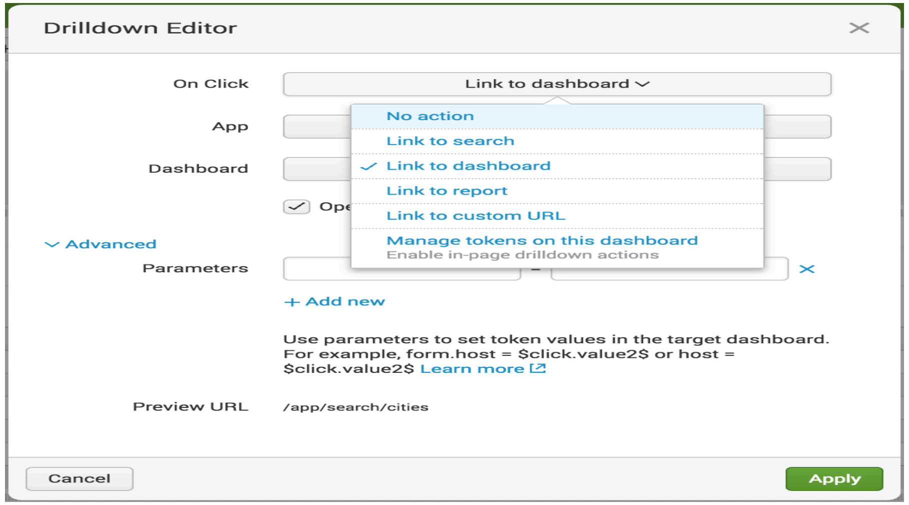

On Click: Link to dashboard

App

Dashboard

☑ Ope

Advanced

Parameters

Use parameters to set token values in the target dashboard. For example, form.host = $click.value2$ or host = $click.value2$ Learn more

Preview URL: /app/search/cities

splunk> .conf2017

New Dashboard Search Editor

Improved search editing experience on dashboards

- Added Search Assistance to Dashboard Search Editor
  - Syntax Highlighting
  - Keyboard Shortcuts
  - Compact Search Assistant

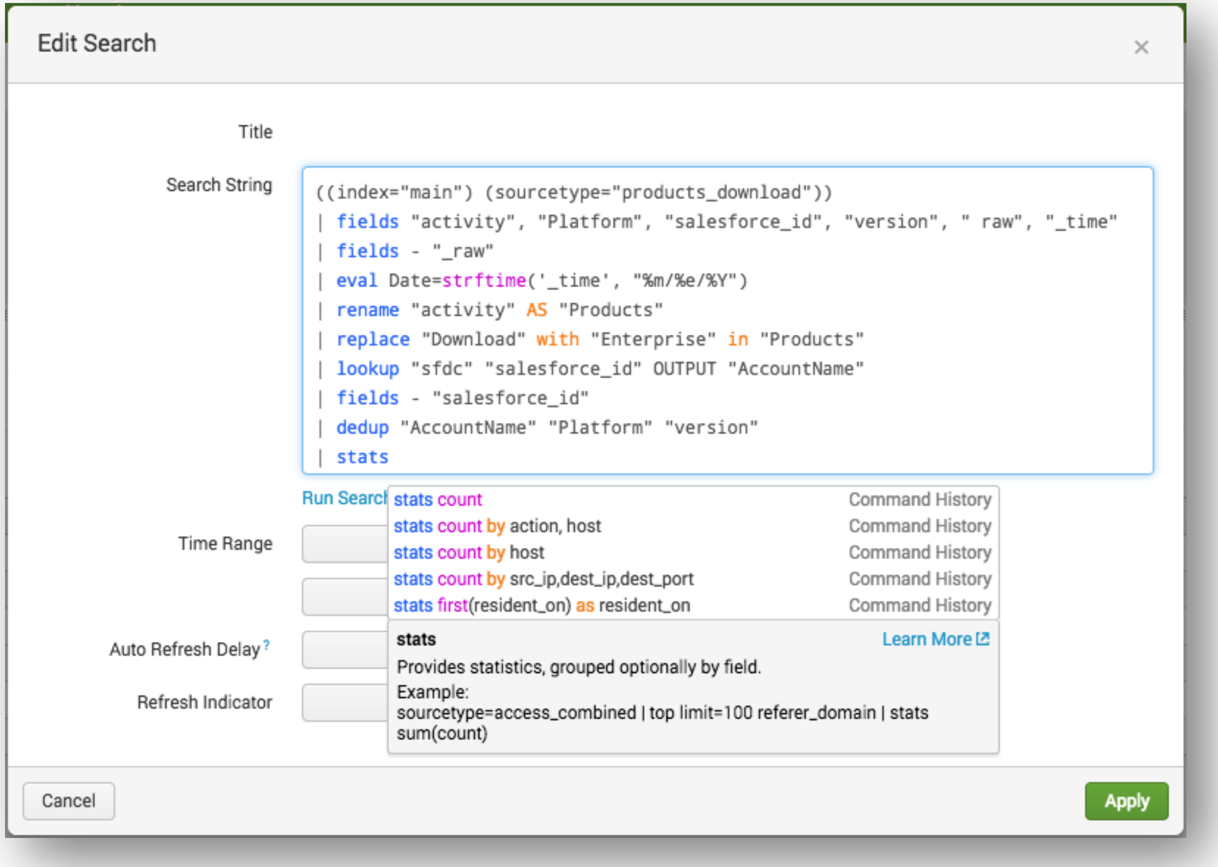

```
((index="main") (sourcetype="products_download"))
| fields "activity", "Platform", "salesforce_id", "version", " raw", "_time"
| fields - "_raw"
| eval Date=strftime('_time', "%m/%e/%Y")
| rename "activity" AS "Products"
| replace "Download" with "Enterprise" in "Products"
| lookup "sfdc" "salesforce_id" OUTPUT "AccountName"
| fields - "salesforce_id"
| dedup "AccountName" "Platform" "version"
| stats
```

```
stats count                                      Command History
stats count by action, host                      Command History
stats count by host                              Command History
stats count by src_ip,dest_ip,dest_port          Command History
stats first(resident_on) as resident_on          Command History
```

stats Learn More

Provides statistics, grouped optionally by field.

Example:
sourcetype=access_combined | top limit=100 referer_domain | stats sum(count)

6.6

Search Productivity Enhancements

Continuing to implement the most asked for Search IDE productivity features

Line Numbers, Auto-formatting, and Dark Mode

Macro Expansions and Keyboard Shortcuts

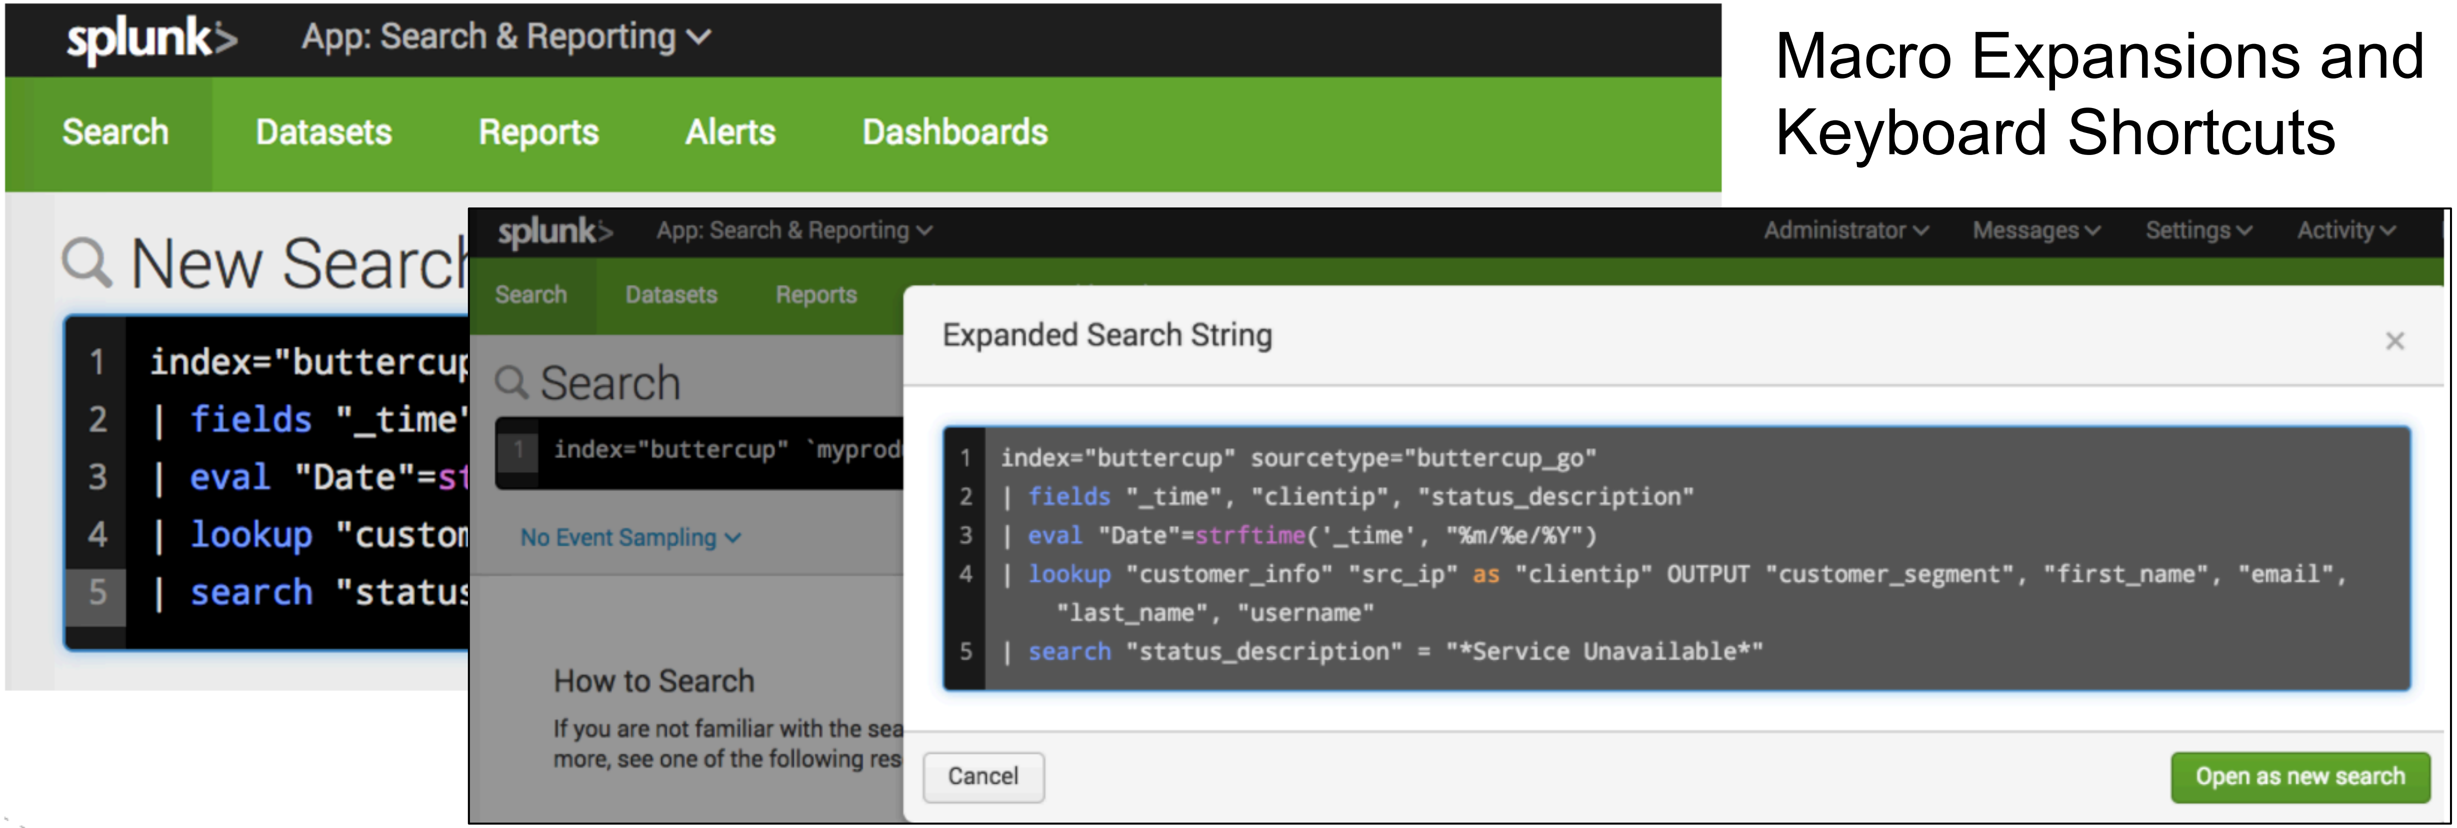

```
index="buttercup
| fields "_time'
| eval "Date"=st
| lookup "custom
| search "status
```

```
index="buttercup" `myprod
```

If you are not familiar with the sea more, see one of the following res

```
index="buttercup" sourcetype="buttercup_go"
| fields "_time", "clientip", "status_description"
| eval "Date"=strftime('_time', "%m/%e/%Y")
| lookup "customer_info" "src_ip" as "clientip" OUTPUT "customer_segment", "first_name", "email",
    "last_name", "username"
| search "status_description" = "*Service Unavailable*"
```

Table Datasets

Turn raw, unstructured events into structured tables without writing SPL

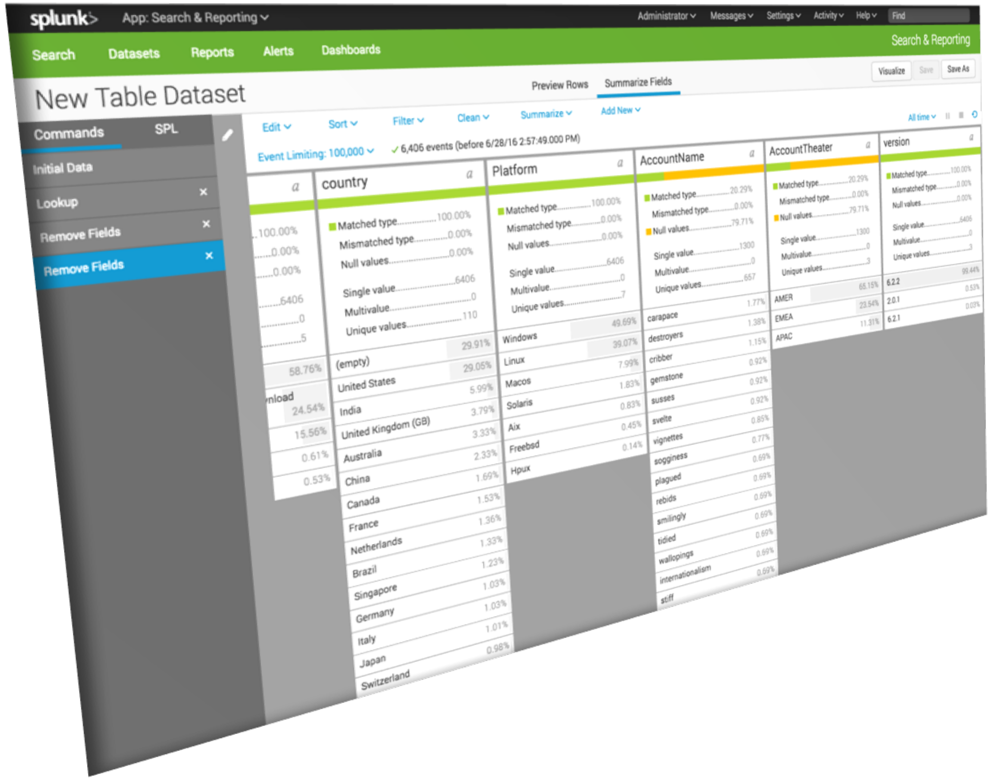

Event Limiting: 100,000 ✓ 6,406 events (before 6/28/16 2:57:49.000 PM)

country
Matched type 100.00%
Mismatched type 0.00%
Null values 0.00%
Single value 6406
Multivalue 0
Unique values 110

country	
(empty)	29.91%
United States	29.05%
India	5.99%
United Kingdom (GB)	3.79%
Australia	3.33%
China	2.33%
Canada	1.69%
France	1.53%
Netherlands	1.36%
Brazil	1.33%
Singapore	1.23%
Germany	1.03%
Italy	1.03%
Japan	1.01%
Switzerland	0.98%

Platform
Matched type 100.00%
Mismatched type 0.00%
Null values 0.00%
Single value 6406
Multivalue 0
Unique values 7

Platform	
Windows	49.69%
Linux	39.07%
Macos	7.99%
Solaris	1.83%
Aix	0.83%
Freebsd	0.45%
Hpux	0.14%

AccountName
Matched type 20.29%
Mismatched type 0.00%
Null values 79.71%
Single value 1300
Multivalue 0
Unique values 657

AccountName	
carapace	1.77%
destroyers	1.38%
cribber	1.15%
gemstone	0.92%
susses	0.92%
svelte	0.92%
vignettes	0.85%
sogginess	0.77%
plagued	0.69%
rebids	0.69%
smilingly	0.69%
tidied	0.69%
wallopings	0.69%
internationalism	0.69%
stiff	0.69%

AccountTheater
Matched type 20.29%
Mismatched type 0.00%
Null values 79.71%
Single value 1300
Multivalue 0
Unique values 3

AccountTheater	
AMER	65.15%
EMEA	23.54%
APAC	11.31%

version
Matched type 100.00%
Mismatched type 0.00%
Null values 0.00%
Single value 6406
Multivalue 0
Unique values 3

version	
6.2.2	99.44%
2.0.1	0.53%
6.2.1	0.03%

- Launched last year with Enterprise 6.5
- Available from Datasets Listings page
- Lots of new capabilities in 6.6…

6.6

Table Datasets

Turn raw, unstructured events into structured tables without writing SPL

Accelerate Datasets

Select Time Range

Improved Editing

Schedule Reports

Export Results

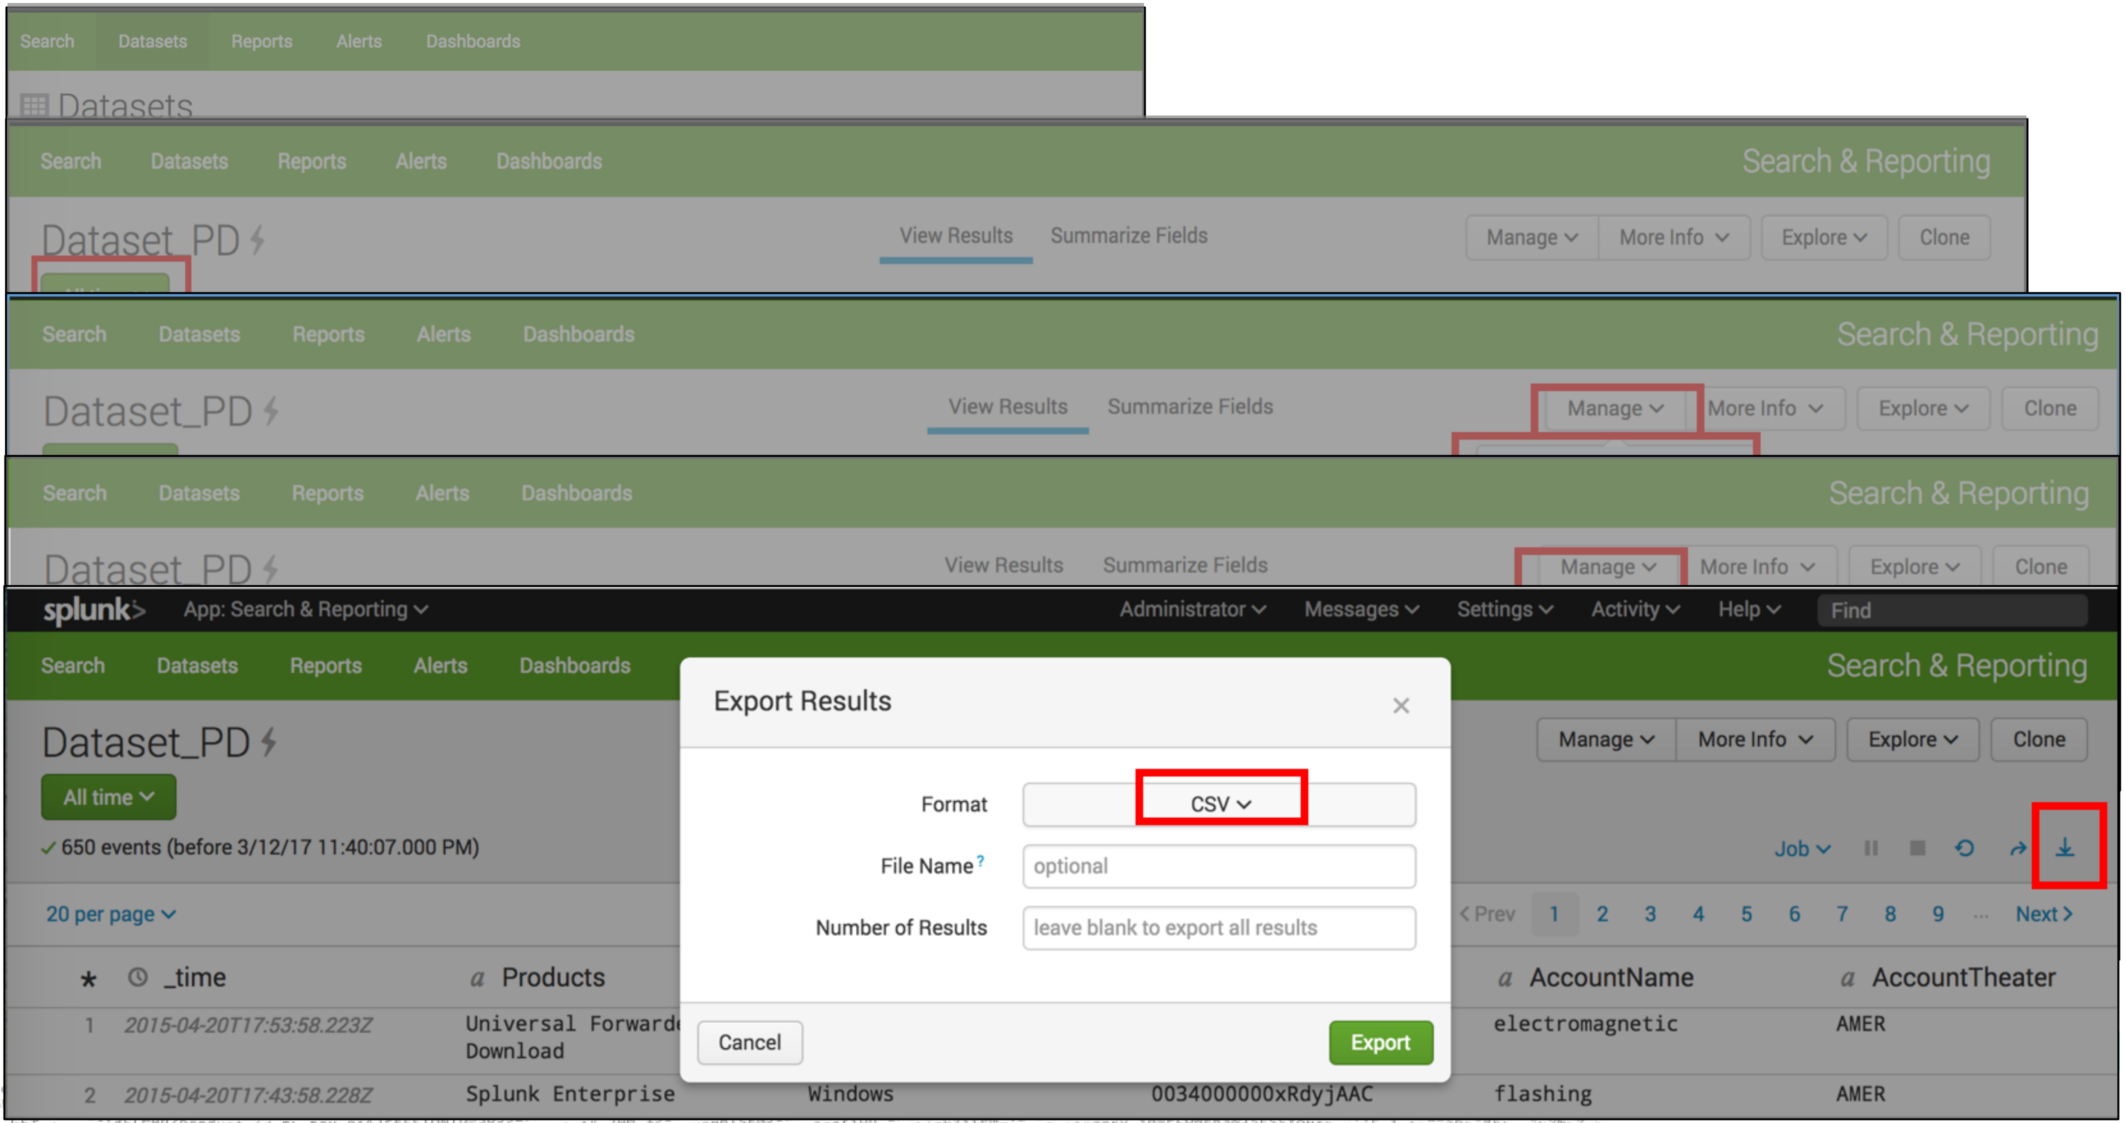

Export Results

Format: CSV

File Name: optional

Number of Results: leave blank to export all results

Cancel

Export

✓ 650 events (before 3/12/17 11:40:07.000 PM)

	_time	Products			AccountName	AccountTheater
1	2015-04-20T17:53:58.223Z	Universal Forwarde Download			electromagnetic	AMER
2	2015-04-20T17:43:58.228Z	Splunk Enterprise	Windows	0034000000xRdyjAAC	flashing	AMER

Amazon Kinesis Firehose

Beta

Amazon Kinesis Firehose Support (Private Beta)

Push data from Kinesis Firehose into Splunk

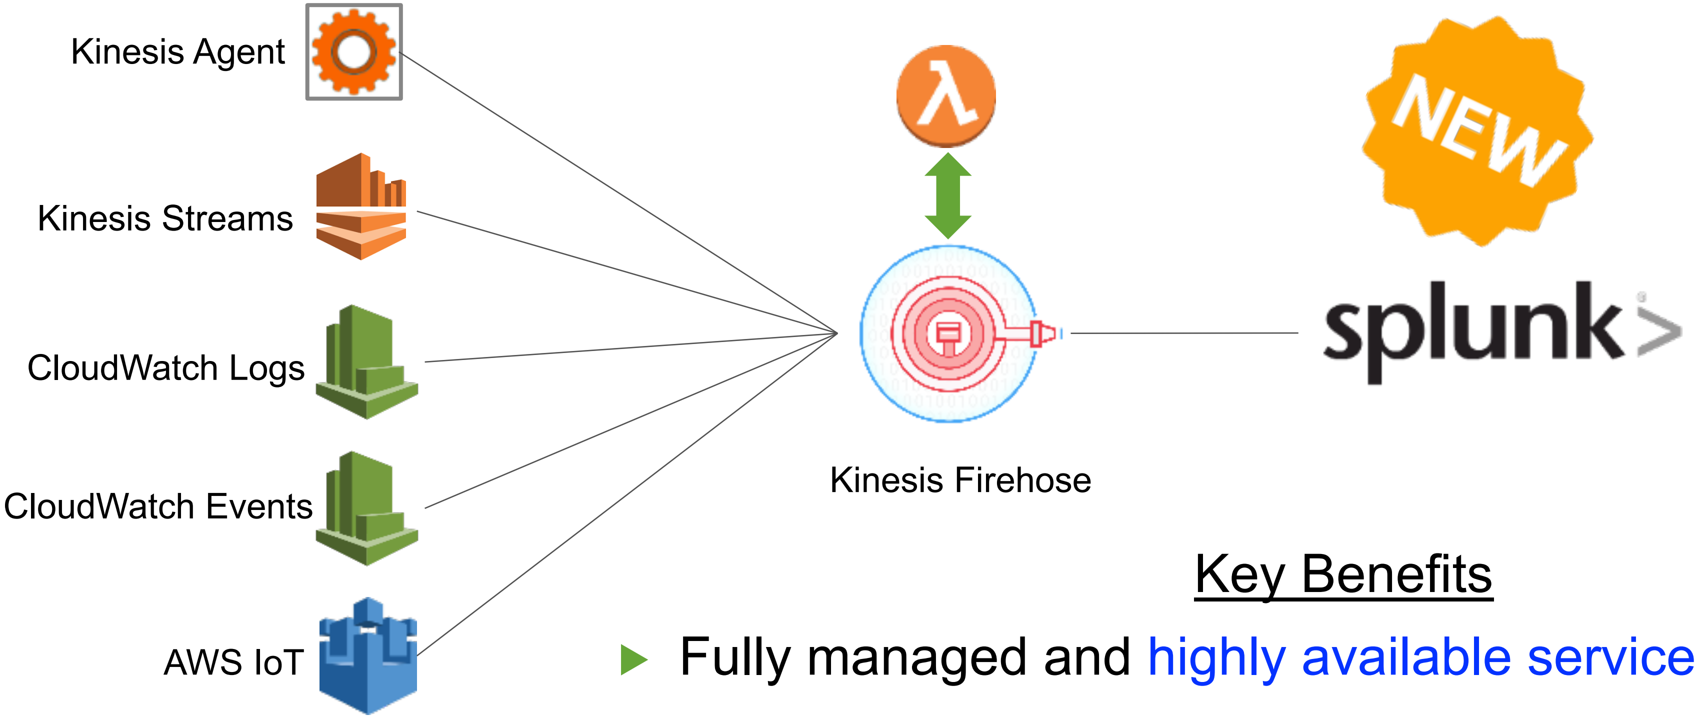

Key Benefits

- Fully managed and highly available service for getting data into Splunk
- Well integrated with various data sources
- Bypass the need for setting up and managing heavy weight forwarder
- Ability to transform raw data prior to sending it to Splunk
- Easy to use with no programming requirement

Monitoring & Investigation

Increased productivity for a wider range of users

Performance, Scale & Manageability

Faster, more scalable and easier to manage than ever before

Machine Learning & Intelligence

More powerful and extensible machine learning capabilities

Improved Search Performance & Scale

7.0

Splunk Metrics Performance

Monitoring and Alerting use cases with metrics improved 20X+

Racing Data (4MM+ Events)

- Chart average speed over 12 hours
- Top panel is Standard Splunk Event Index
- Bottom panel is new Metrics Index
- Watch how quickly the Metrics index (bottom) finishes

Play Video →

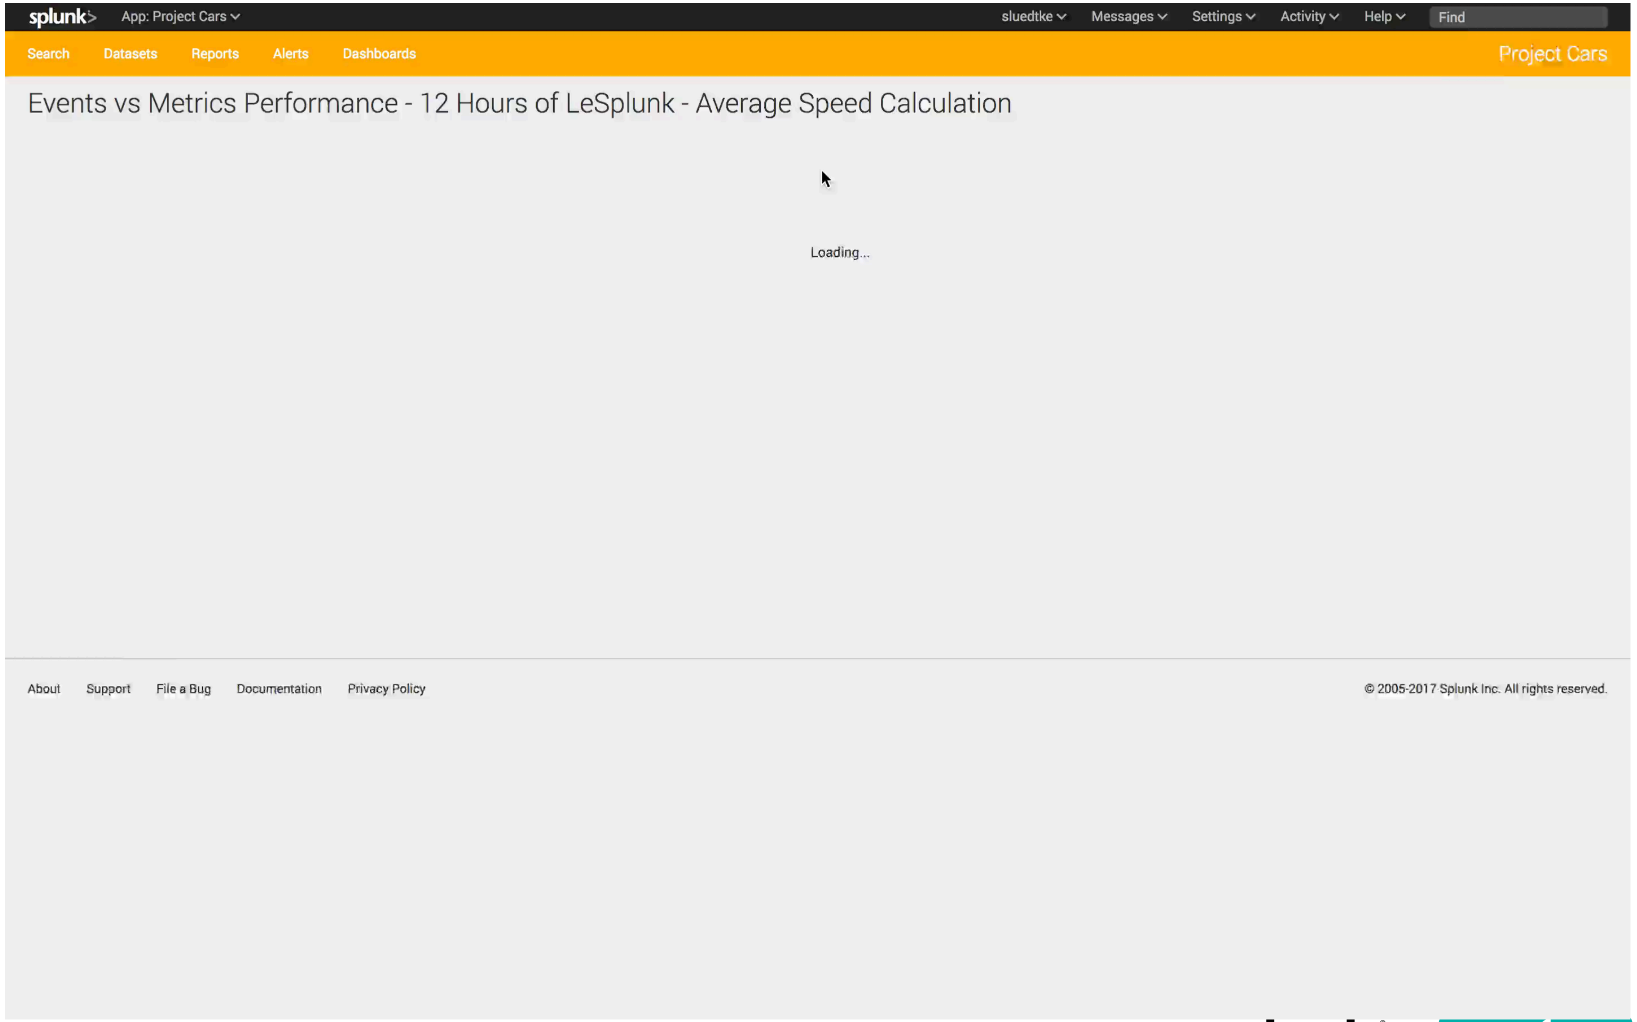

Events vs Metrics Performance - 12 Hours of LeSplunk - Average Speed Calculation

Loading...

© 2005-2017 Splunk Inc. All rights reserved.

splunk> .conf2017

Demo – SPL Optimizer

Buttercup Games

If we raise prices 15%, how many of the Strategy products could we sell for over $20?

Can you make this search more efficient?

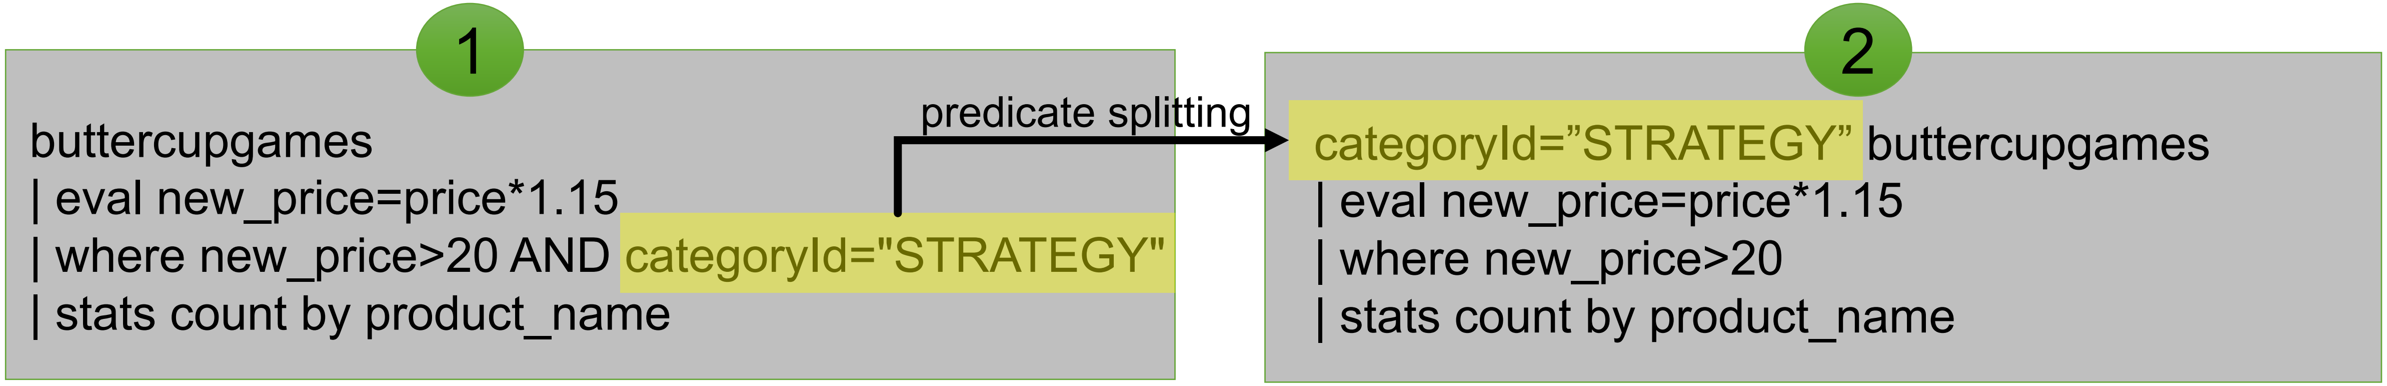

1

```
buttercupgames
| eval new_price=price*1.15
| where new_price>20 AND categoryId="STRATEGY"
| stats count by product_name
```

predicate splitting

2

```
categoryId=”STRATEGY” buttercupgames
| eval new_price=price*1.15
| where new_price>20
| stats count by product_name
```

…search result out of 232,630 events

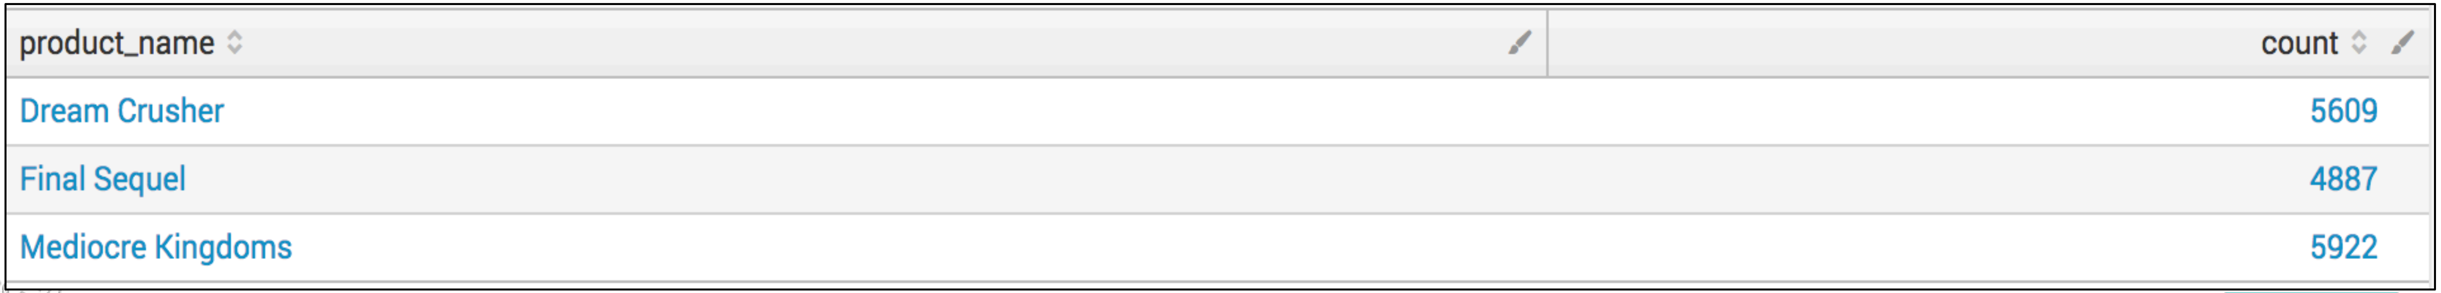

product_name	count
Dream Crusher	5609
Final Sequel	4887
Mediocre Kingdoms	5922

SPL Optimizer

You Actually Don't Need to Know!

Under the covers, Splunk optimizes the search for you

Without OPTIMIZER

Scanned 232,630 events in 11.949 seconds

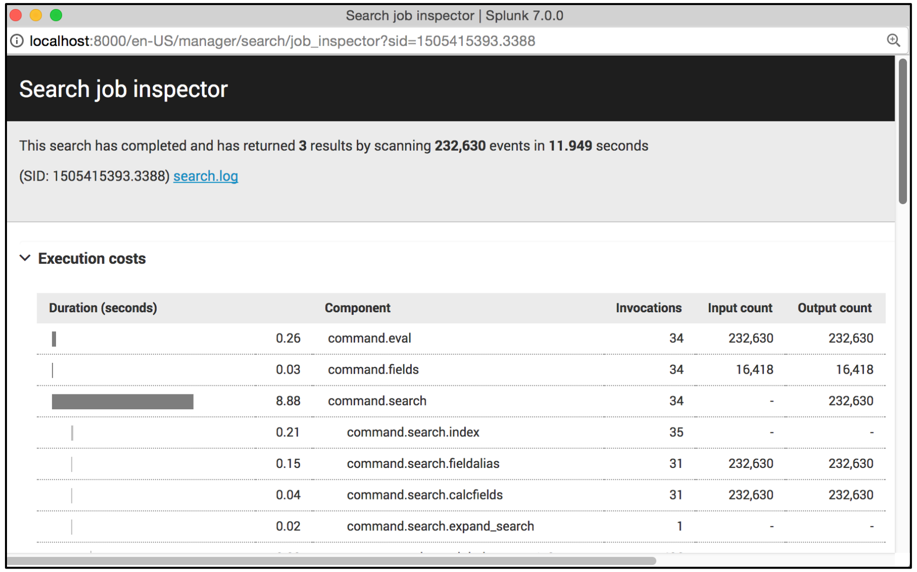

Search job inspector | Splunk 7.0.0

localhost:8000/en-US/manager/search/job_inspector?sid=1505415393.3388

Search job inspector

This search has completed and has returned **3** results by scanning **232,630** events in **11.949** seconds

(SID: 1505415393.3388) search.log

Execution costs

Duration (seconds)	Component	Invocations	Input count	Output count
0.26	command.eval	34	232,630	232,630
0.03	command.fields	34	16,418	16,418
8.88	command.search	34	-	232,630
0.21	command.search.index	35	-	-
0.15	command.search.fieldalias	31	232,630	232,630
0.04	command.search.calcfields	31	232,630	232,630
0.02	command.search.expand_search	1	-	-

```
buttercupgames
| eval new_price=price*1.15
| where new_price>20 AND categoryId="STRATEGY"
| stats count by product_name
```

With OPTIMIZER

Scanned 28,329 events in 2.184 seconds

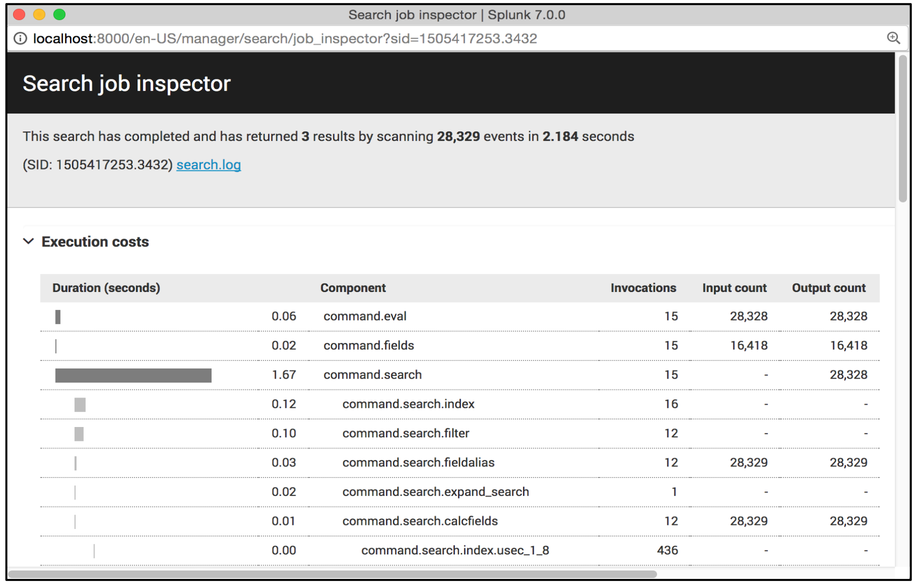

Search job inspector | Splunk 7.0.0

localhost:8000/en-US/manager/search/job_inspector?sid=1505417253.3432

Search job inspector

This search has completed and has returned **3** results by scanning **28,329** events in **2.184** seconds

(SID: 1505417253.3432) search.log

Execution costs

Duration (seconds)	Component	Invocations	Input count	Output count
0.06	command.eval	15	28,328	28,328
0.02	command.fields	15	16,418	16,418
1.67	command.search	15	-	28,328
0.12	command.search.index	16	-	-
0.10	command.search.filter	12	-	-
0.03	command.search.fieldalias	12	28,329	28,329
0.02	command.search.expand_search	1	-	-
0.01	command.search.calcfields	12	28,329	28,329
0.00	command.search.index.usec_1_8	436	-	-

predicate splitting

```
categoryId="STRATEGY" buttercupgames
| eval new_price=price*1.15
| where new_price>20
| stats count by product_name
```

splunk> .conf2017

6.6

Search Performance Optimizations

Predicate Splitting

Splits predicates and execute early
(Splits ANDs not Ors)

```
search 500 | eval x=a*b | where x=100 AND y="Test"
```

⬇

```
search 500 y="Test" | eval x=a*b | where x=100
```

limits.conf:

```
[search_optimization::predicate_split]
enabled = true
```

Projection Elimination

Removes commands that produce unneeded new fields

```
search 500 | eval x=a*b | lookup cust id OUTPUT name |
stats count by host
```

⬇

```
search 500 | stats count by host
```

limits.conf:

```
[search_optimization::projection_elimination]
enabled = true
cmds_black_list = <Commands List>
```

Automatically applies optimizations to optimize query execution speed

Faster Data Model Acceleration (DMA)

Big Improvements in Search Performance and Resource Usage

- Applicable to both platform and premium solutions
  - Acceleration searches run up to **3x** faster*.
  - Data Model summarization lag is up to **1/3** less*.
  - Less Resources used (CPU, MEM, IO)
  - Improved Responsiveness - acceleration data available more quickly for *tstats* querying
  - Enterprise Security models get a big boost in speed

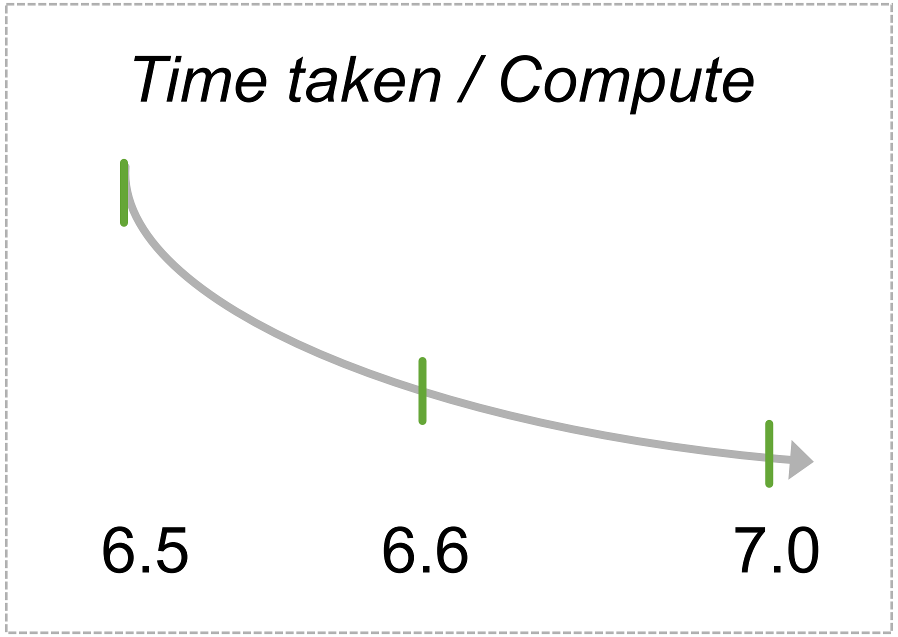

* From 6.5 → 7.0 (based on test lab data)

6.6

Enhanced Indexer Clustering Scalability

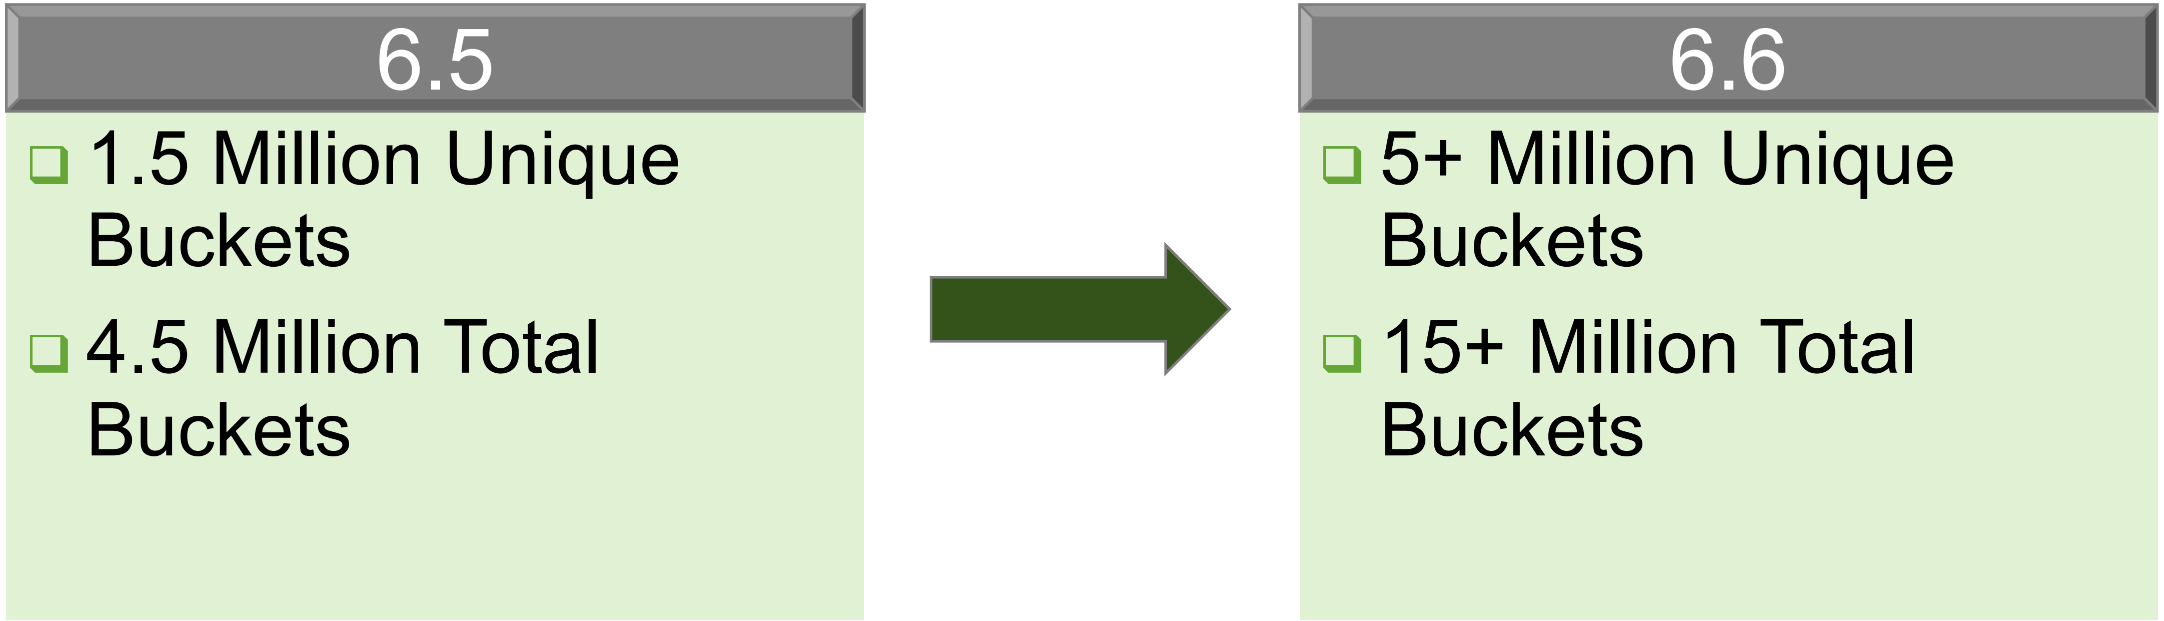

6.5

- ❑ 1.5 Million Unique Buckets
- ❑ 4.5 Million Total Buckets

6.6

- ❑ 5+ Million Unique Buckets
- ❑ 15+ Million Total Buckets

Delivers seamless scalability to 100s of Petabytes

Improved Manageability

Improved Resiliency and Developer Productivity

6.6

Cluster Resiliency

Improved resiliency to hardware and networking failures

Indexer Clustering

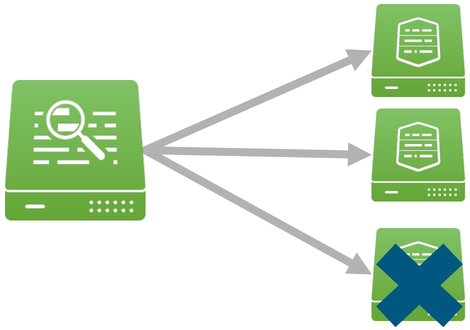

If an Indexer goes offline, search is not disrupted

Indexer recovery time is much faster

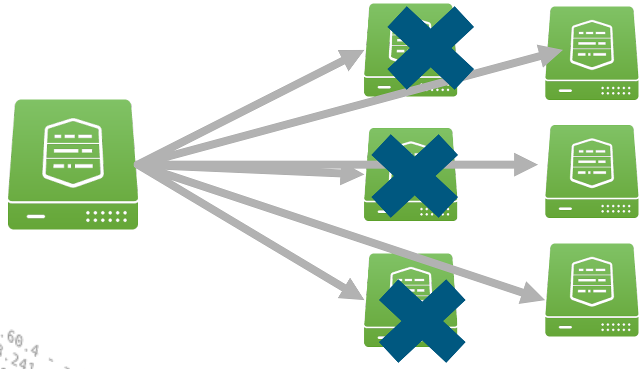

In a multi-site cluster, Forwarders can automatically failover to a secondary site

Search Head Clustering

Smarter knowledge object (config) replication across search heads

Intelligent captain selection based on configuration awareness

Improved replication speed across search heads and indexers for apps, lookup tables, etc

6.6

Cluster Management

Improved manageability as you grow and scale

Indexer Clustering

Manual Indexer Detention - Selectively control the flow of data into an Indexer

Deploy new apps (with reloadable configs) without cluster restart

Rollback to previous configuration if any operational issues encountered

Search Head Clustering

New User Interface

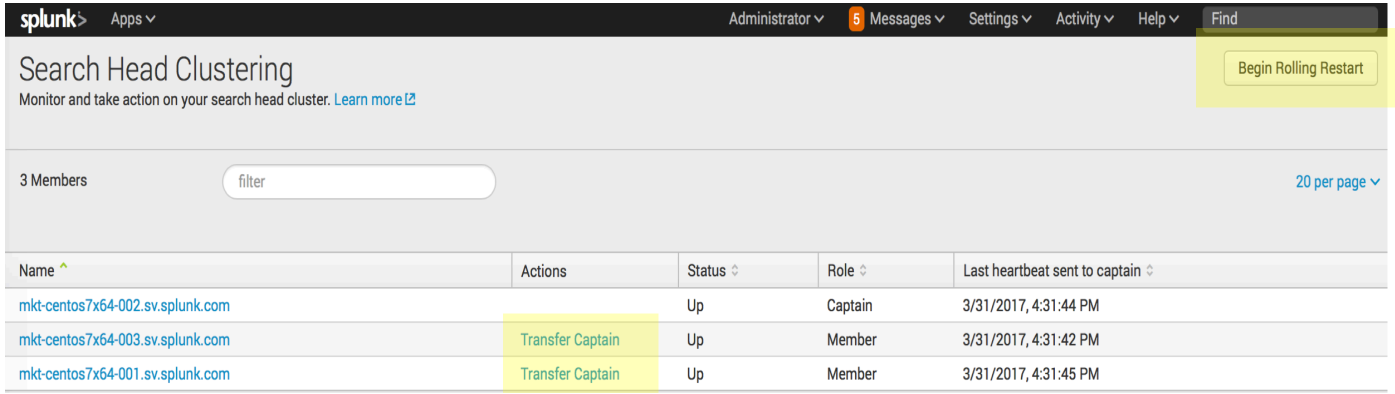

Search Head Clustering

Monitor and take action on your search head cluster. Learn more

3 Members

Name	Actions	Status	Role	Last heartbeat sent to captain
mkt-centos7x64-002.sv.splunk.com		Up	Captain	3/31/2017, 4:31:44 PM
mkt-centos7x64-003.sv.splunk.com	Transfer Captain	Up	Member	3/31/2017, 4:31:42 PM
mkt-centos7x64-001.sv.splunk.com	Transfer Captain	Up	Member	3/31/2017, 4:31:45 PM

Transfer Search Head Captain

Begin Rolling Restart

6.6

Knowledge Object Management

Automatically detect orphaned Knowledge Objects when an employee leaves your company

Manually Re-assign to a different user

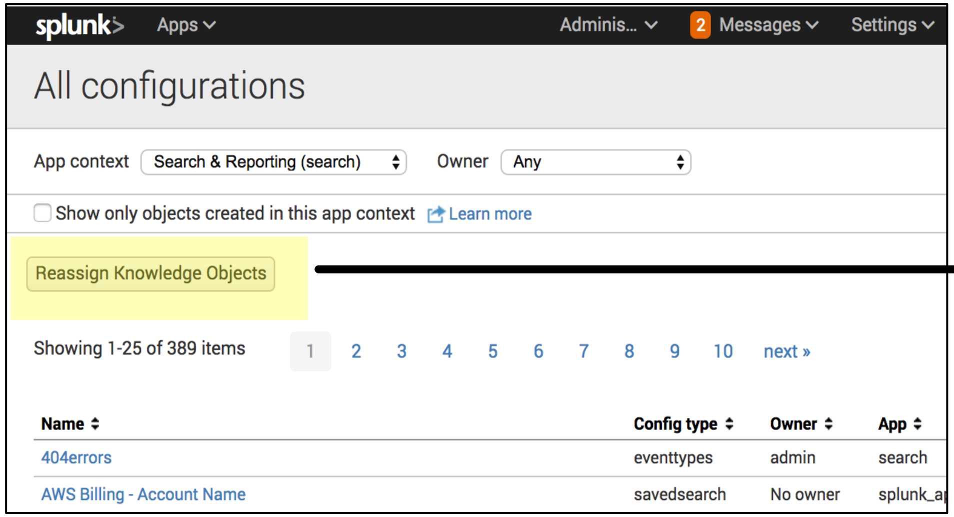

All configurations

App context: Search & Reporting (search) Owner: Any

□ Show only objects created in this app context Learn more

Reassign Knowledge Objects

Showing 1-25 of 389 items 1 2 3 4 5 6 7 8 9 10 next »

Name	Config type	Owner	App
404errors	eventtypes	admin	search
AWS Billing - Account Name	savedsearch	No owner	splunk_ap

Settings → All Configurations

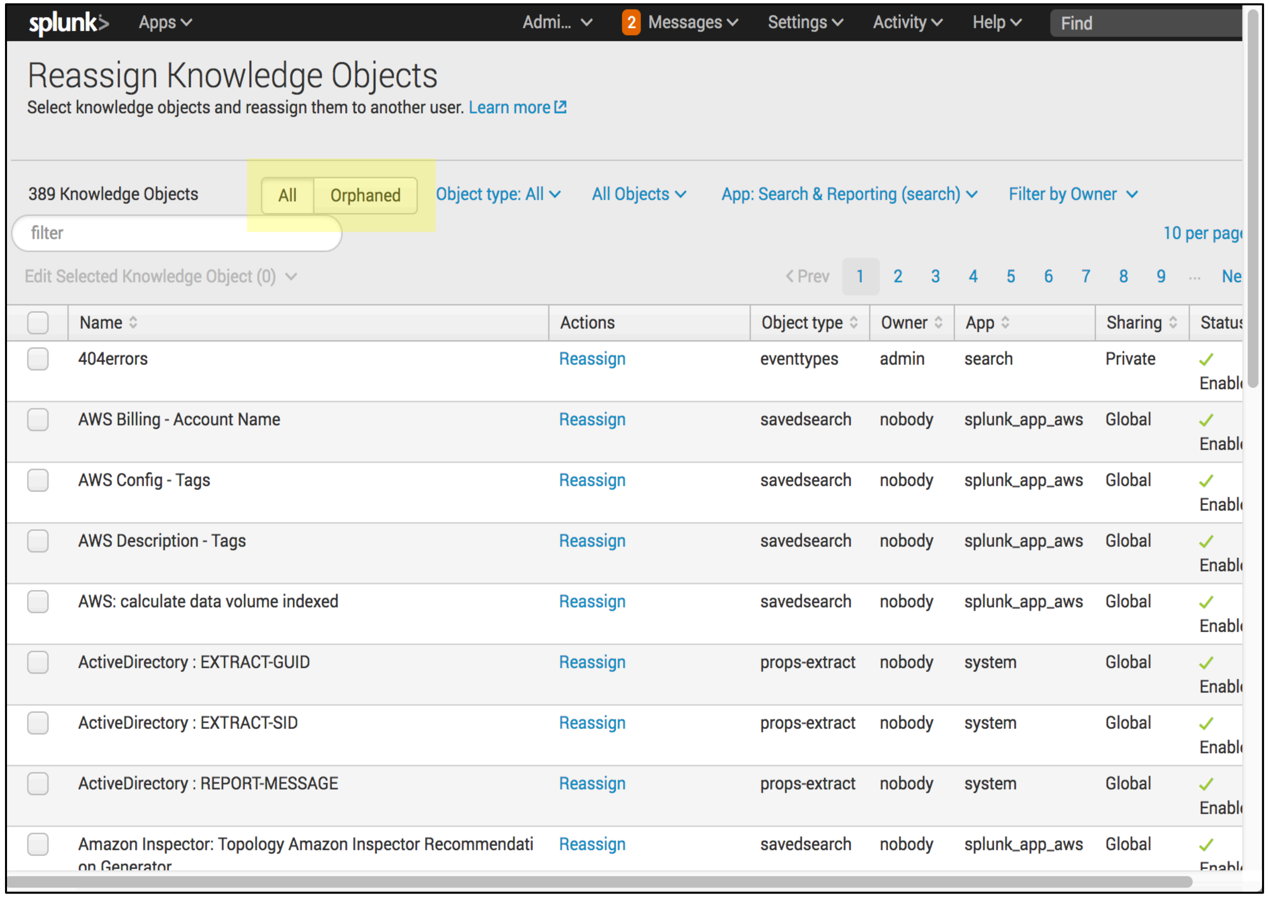

Reassign Knowledge Objects

Select knowledge objects and reassign them to another user. Learn more

389 Knowledge Objects All Orphaned Object type: All All Objects App: Search & Reporting (search) Filter by Owner

filter 10 per pag

Edit Selected Knowledge Object (0) ‹ Prev 1 2 3 4 5 6 7 8 9 ... Ne

□	Name	Actions	Object type	Owner	App	Sharing	Status
□	404errors	Reassign	eventtypes	admin	search	Private	✓ Enabl
□	AWS Billing - Account Name	Reassign	savedsearch	nobody	splunk_app_aws	Global	✓ Enabl
□	AWS Config - Tags	Reassign	savedsearch	nobody	splunk_app_aws	Global	✓ Enabl
□	AWS Description - Tags	Reassign	savedsearch	nobody	splunk_app_aws	Global	✓ Enabl
□	AWS: calculate data volume indexed	Reassign	savedsearch	nobody	splunk_app_aws	Global	✓ Enabl
□	ActiveDirectory : EXTRACT-GUID	Reassign	props-extract	nobody	system	Global	✓ Enabl
□	ActiveDirectory : EXTRACT-SID	Reassign	props-extract	nobody	system	Global	✓ Enabl
□	ActiveDirectory : REPORT-MESSAGE	Reassign	props-extract	nobody	system	Global	✓ Enabl
□	Amazon Inspector: Topology Amazon Inspector Recommendation Generator	Reassign	savedsearch	nobody	splunk_app_aws	Global	✓ Enabl

7.0

Data Ingestion Management

New Regex Profiling views for Admins to monitor incoming data

- Found in the Monitoring Console under “Indexing→ Indexing Performance”
- Find CPU time spent on Regex Extraction
- Split-by source, sourcetype, index and host

New Monitoring Console Views

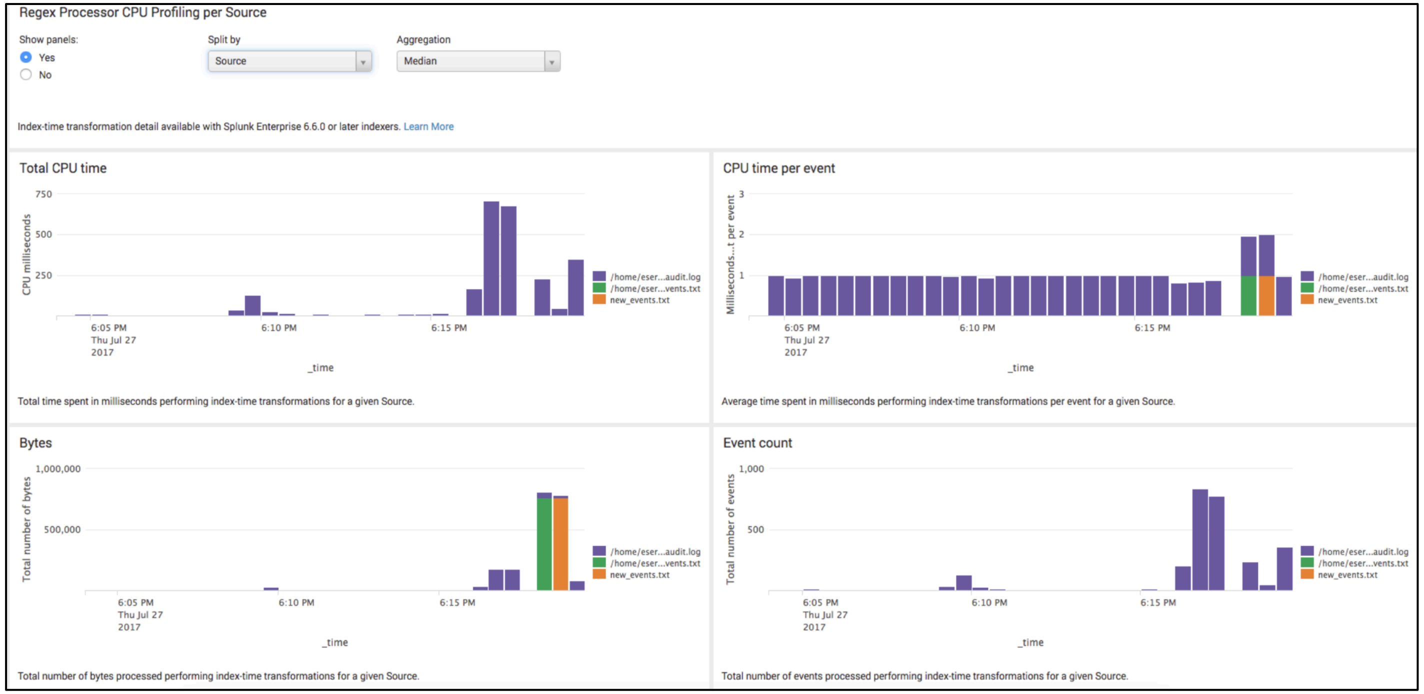

Regex Processor CPU Profiling per Source

Show panels:
● Yes
○ No

Split by: Source

Aggregation: Median

Index-time transformation detail available with Splunk Enterprise 6.6.0 or later indexers. Learn More

Total CPU time

Total time spent in milliseconds performing index-time transformations for a given Source.

CPU time per event

Average time spent in milliseconds performing index-time transformations per event for a given Source.

Bytes

Total number of bytes processed performing index-time transformations for a given Source.

Event count

Total number of events processed performing index-time transformations for a given Source.

6.6

Data Ingestion Management

New Data Quality views for Admins to check the accuracy of incoming data

- Found in the Monitoring Console under Indexing→Inputs
- Surfaces index-time warnings and errors (splunkd.log)
- Find Timestamp, Linebreaking, and Aggregation issues with incoming data

New Monitoring Console Views

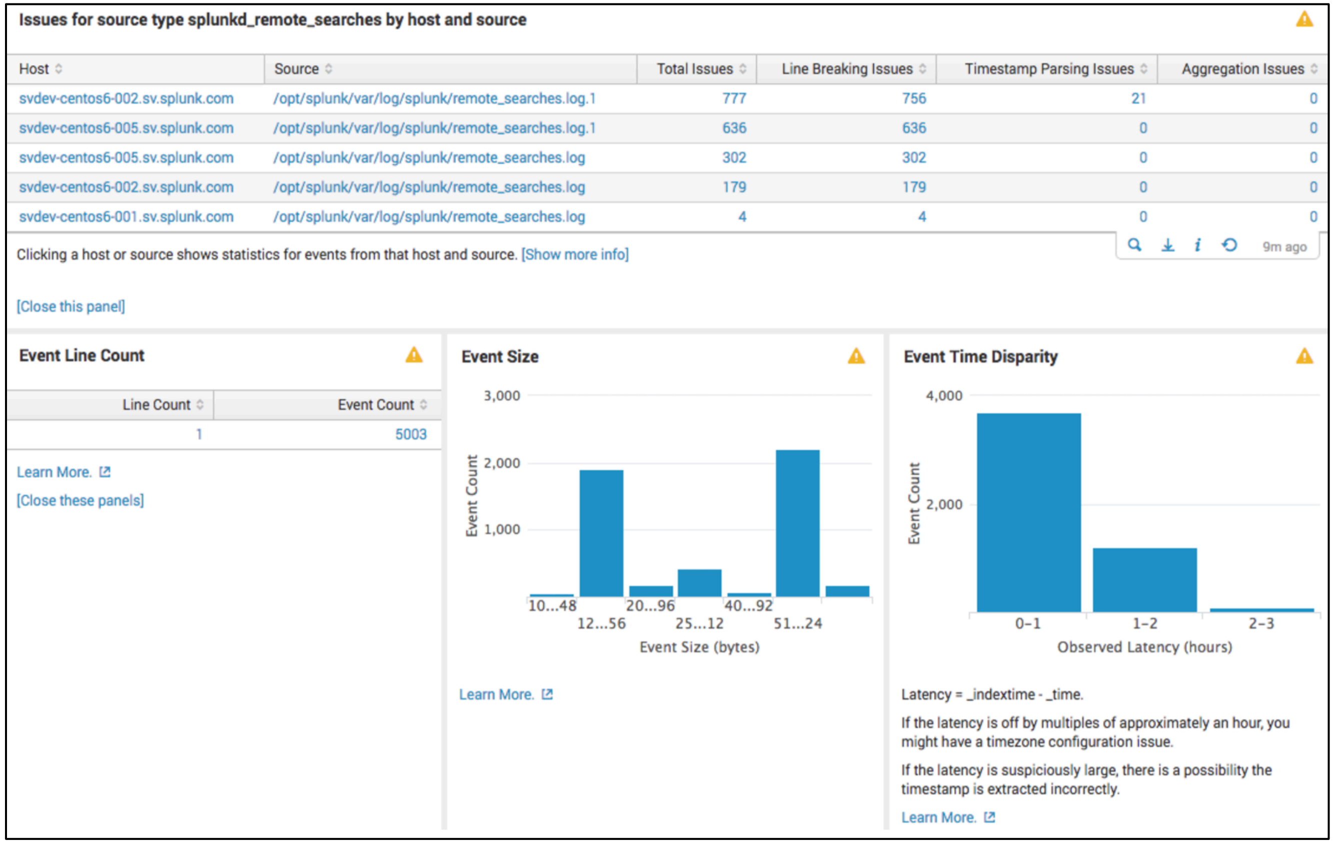

Issues for source type splunkd_remote_searches by host and source

Host	Source	Total Issues	Line Breaking Issues	Timestamp Parsing Issues	Aggregation Issues
svdev-centos6-002.sv.splunk.com	/opt/splunk/var/log/splunk/remote_searches.log.1	777	756	21	0
svdev-centos6-005.sv.splunk.com	/opt/splunk/var/log/splunk/remote_searches.log.1	636	636	0	0
svdev-centos6-005.sv.splunk.com	/opt/splunk/var/log/splunk/remote_searches.log	302	302	0	0
svdev-centos6-002.sv.splunk.com	/opt/splunk/var/log/splunk/remote_searches.log	179	179	0	0
svdev-centos6-001.sv.splunk.com	/opt/splunk/var/log/splunk/remote_searches.log	4	4	0	0

Clicking a host or source shows statistics for events from that host and source. [Show more info]

[Close this panel]

Event Line Count

Line Count	Event Count
1	5003

Learn More.

[Close these panels]

Event Size

Learn More.

Event Time Disparity

Latency = _indextime - _time.

If the latency is off by multiples of approximately an hour, you might have a timezone configuration issue.

If the latency is suspiciously large, there is a possibility the timestamp is extracted incorrectly.

Learn More.

splunk> .conf2017

Cloud App Management – Certified Apps

Install and Update Splunk certified apps with the click of a button

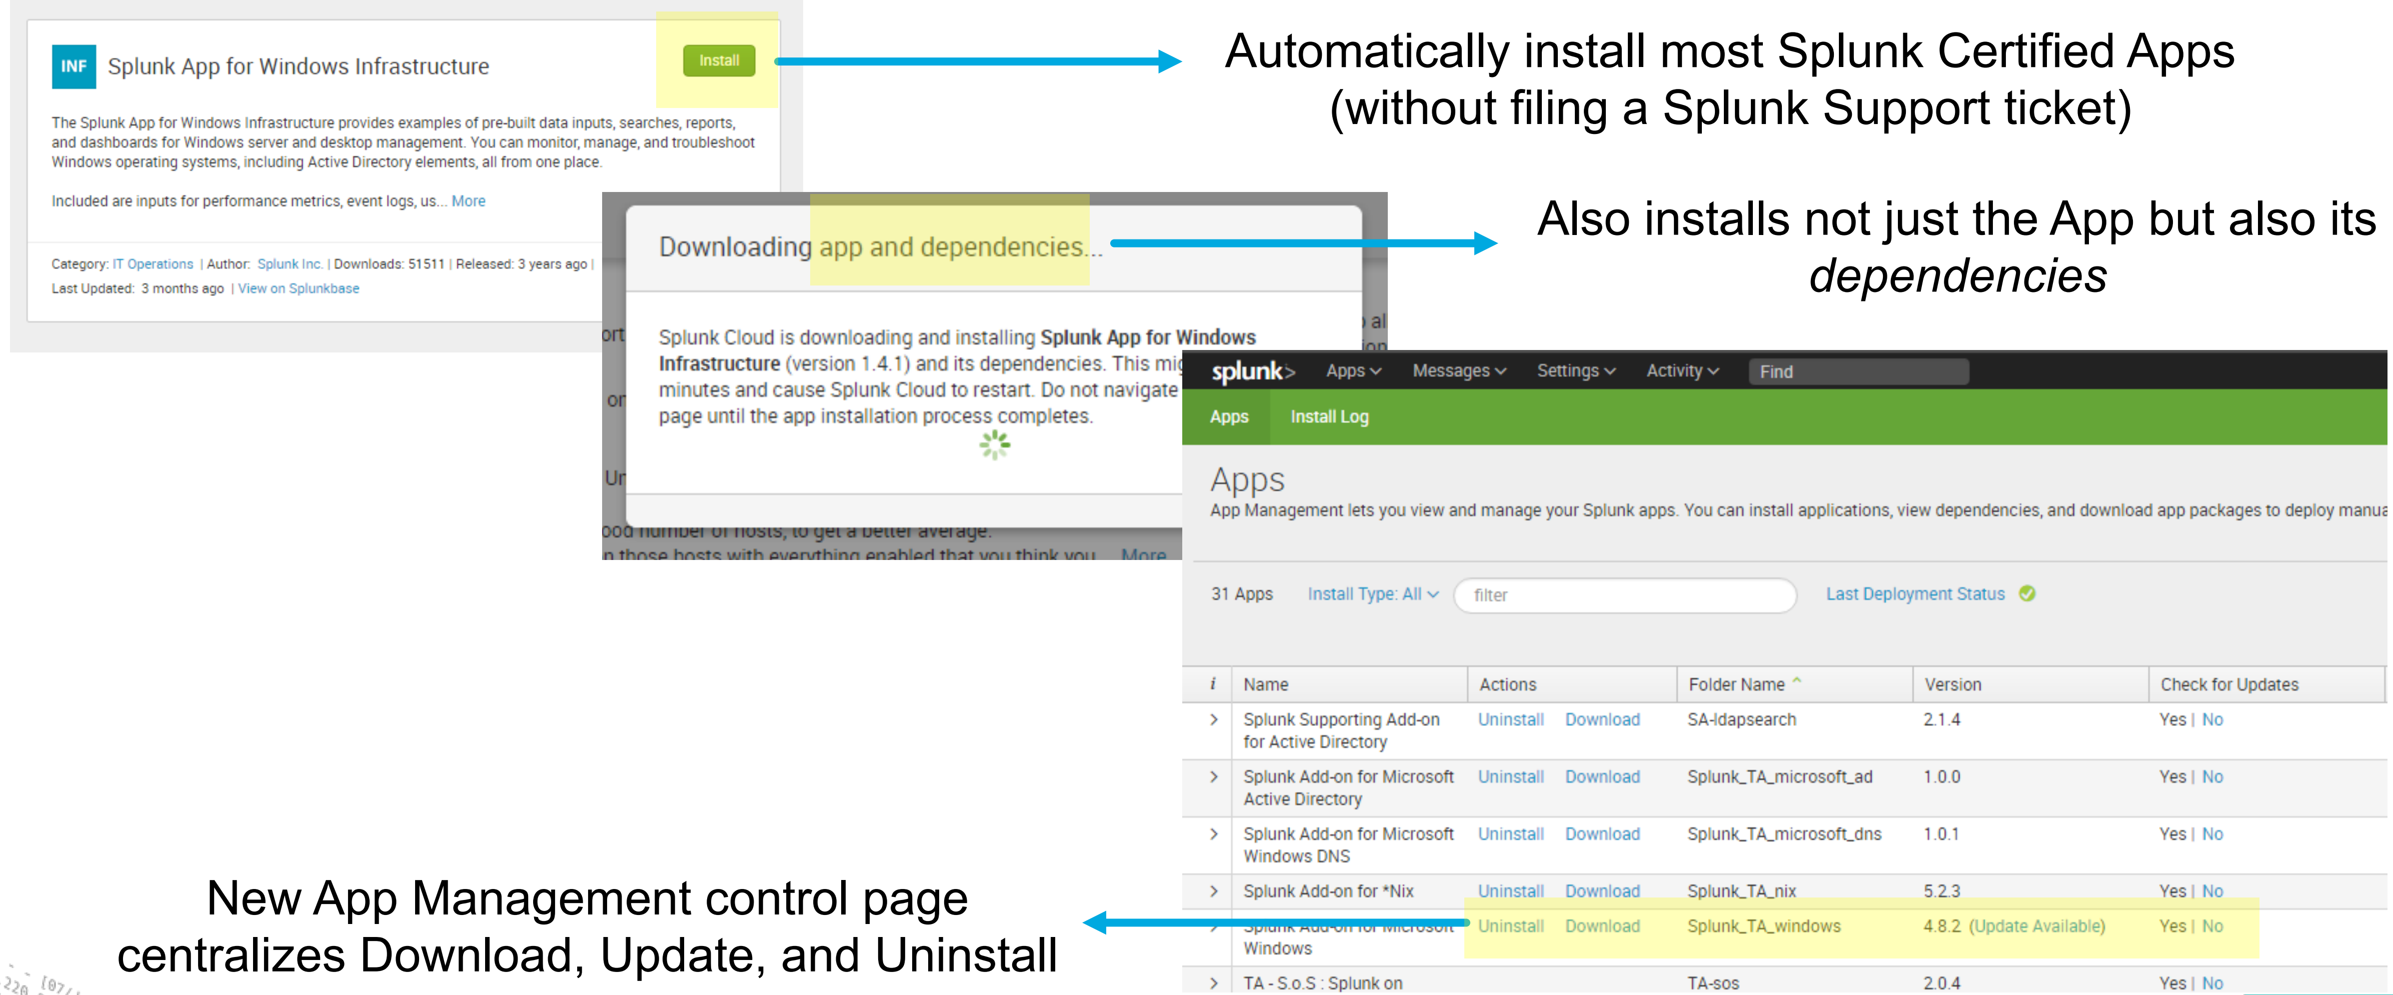

INF Splunk App for Windows Infrastructure

The Splunk App for Windows Infrastructure provides examples of pre-built data inputs, searches, reports, and dashboards for Windows server and desktop management. You can monitor, manage, and troubleshoot Windows operating systems, including Active Directory elements, all from one place.

Included are inputs for performance metrics, event logs, us... More

Category: IT Operations | Author: Splunk Inc. | Downloads: 51511 | Released: 3 years ago | Last Updated: 3 months ago | View on Splunkbase

Downloading app and dependencies...

Splunk Cloud is downloading and installing **Splunk App for Windows Infrastructure** (version 1.4.1) and its dependencies. This mi… minutes and cause Splunk Cloud to restart. Do not navigate … page until the app installation process completes.

Apps

App Management lets you view and manage your Splunk apps. You can install applications, view dependencies, and download app packages to deploy manua…

31 Apps | Install Type: All | filter | Last Deployment Status ✓

i	Name	Actions	Folder Name	Version	Check for Updates
>	Splunk Supporting Add-on for Active Directory	Uninstall Download	SA-ldapsearch	2.1.4	Yes | No
>	Splunk Add-on for Microsoft Active Directory	Uninstall Download	Splunk_TA_microsoft_ad	1.0.0	Yes | No
>	Splunk Add-on for Microsoft Windows DNS	Uninstall Download	Splunk_TA_microsoft_dns	1.0.1	Yes | No
>	Splunk Add-on for *Nix	Uninstall Download	Splunk_TA_nix	5.2.3	Yes | No
>	Splunk Add-on for Microsoft Windows	Uninstall Download	Splunk_TA_windows	4.8.2 (Update Available)	Yes | No
>	TA - S.o.S : Splunk on		TA-sos	2.0.4	Yes | No

Automatically install most Splunk Certified Apps (without filing a Splunk Support ticket)

Also installs not just the App but also its *dependencies*

New App Management control page centralizes Download, Update, and Uninstall

Cloud*

Cloud Application Management – Private Apps

Quickly install and uninstall ***private apps***

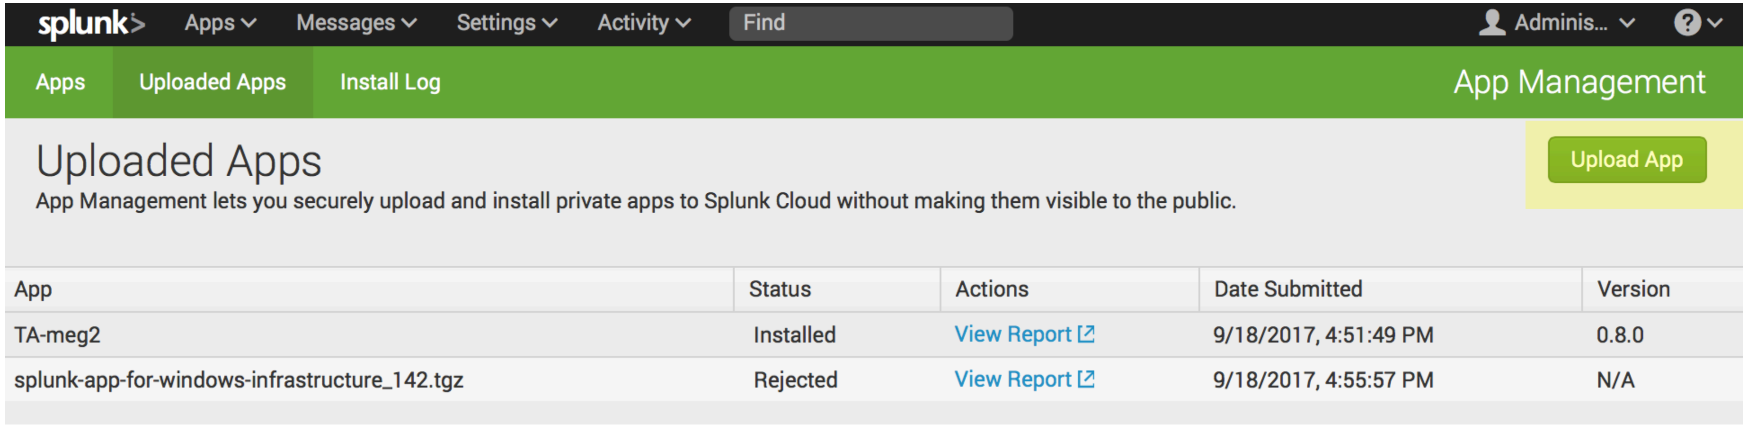

Uploaded Apps

App Management lets you securely upload and install private apps to Splunk Cloud without making them visible to the public.

App	Status	Actions	Date Submitted	Version
TA-meg2	Installed	View Report	9/18/2017, 4:51:49 PM	0.8.0
splunk-app-for-windows-infrastructure_142.tgz	Rejected	View Report	9/18/2017, 4:55:57 PM	N/A

- New self-service Upload App interface for the cloud
  - No support ticket needed!
- Automatically runs AppInspect and generates a detailed AppInspect vetting report
- If approved, Install/Update private apps with a single click

*7.0 deployment dates in the cloud will vary across customers (planned initial delivery starting later this year)

Application Development

Streamlined App Development Lifecycle

Develop

Add-on Builder App

AppInspect Tool

Package

New Packaging Toolkit (1.0)

Certify

Faster App Certification Process

Deploy

More Self-service App Installation & Management

Get your apps built and deployed faster than ever!

Monitoring & Investigation

Increased productivity for a wider range of users

Performance, Scale & Manageability

Faster, more scalable and easier to manage than ever before

Machine Learning & Intelligence

More powerful and extensible machine learning capabilities

Machine Learning Overview

Learning by Example

Start

Use Case / Problem
(hard for human to figure out)

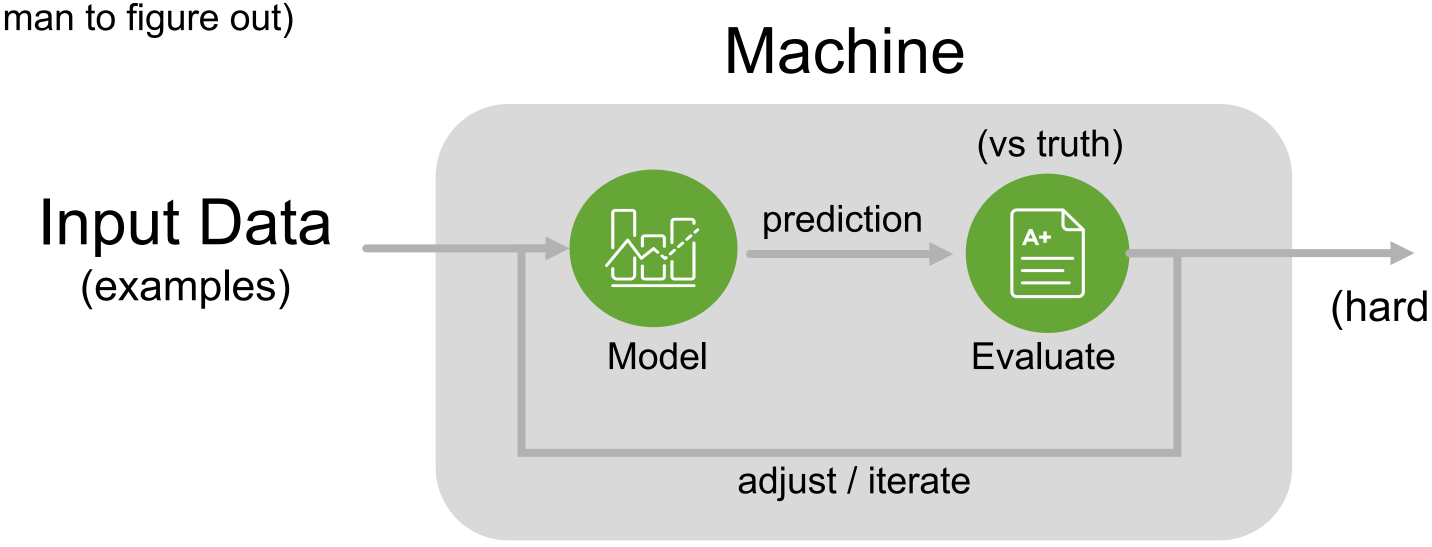

Intelligence
(hard for human to figure out)

RECURSION pharmaceuticals

TELUS
Network Optimization
Detect & Prevent Equipment Failure

NTT docomo
Security / Fraud Prevention

Telco
Prevent Cell Tower Failure
Optimize Repair Operations

Zillow®
Prioritize Website Issues
and Predict Root Cause

Entertainment Company
Predict Gaming Outages
Fraud Prevention

CONCANON™ INSIGHT ON DEMAND
Machine Learning Consulting Services

SCIANTA ANALYTICS DEEP INSIGHT™
Analytics App built on ML Toolkit

Machine Learning Process

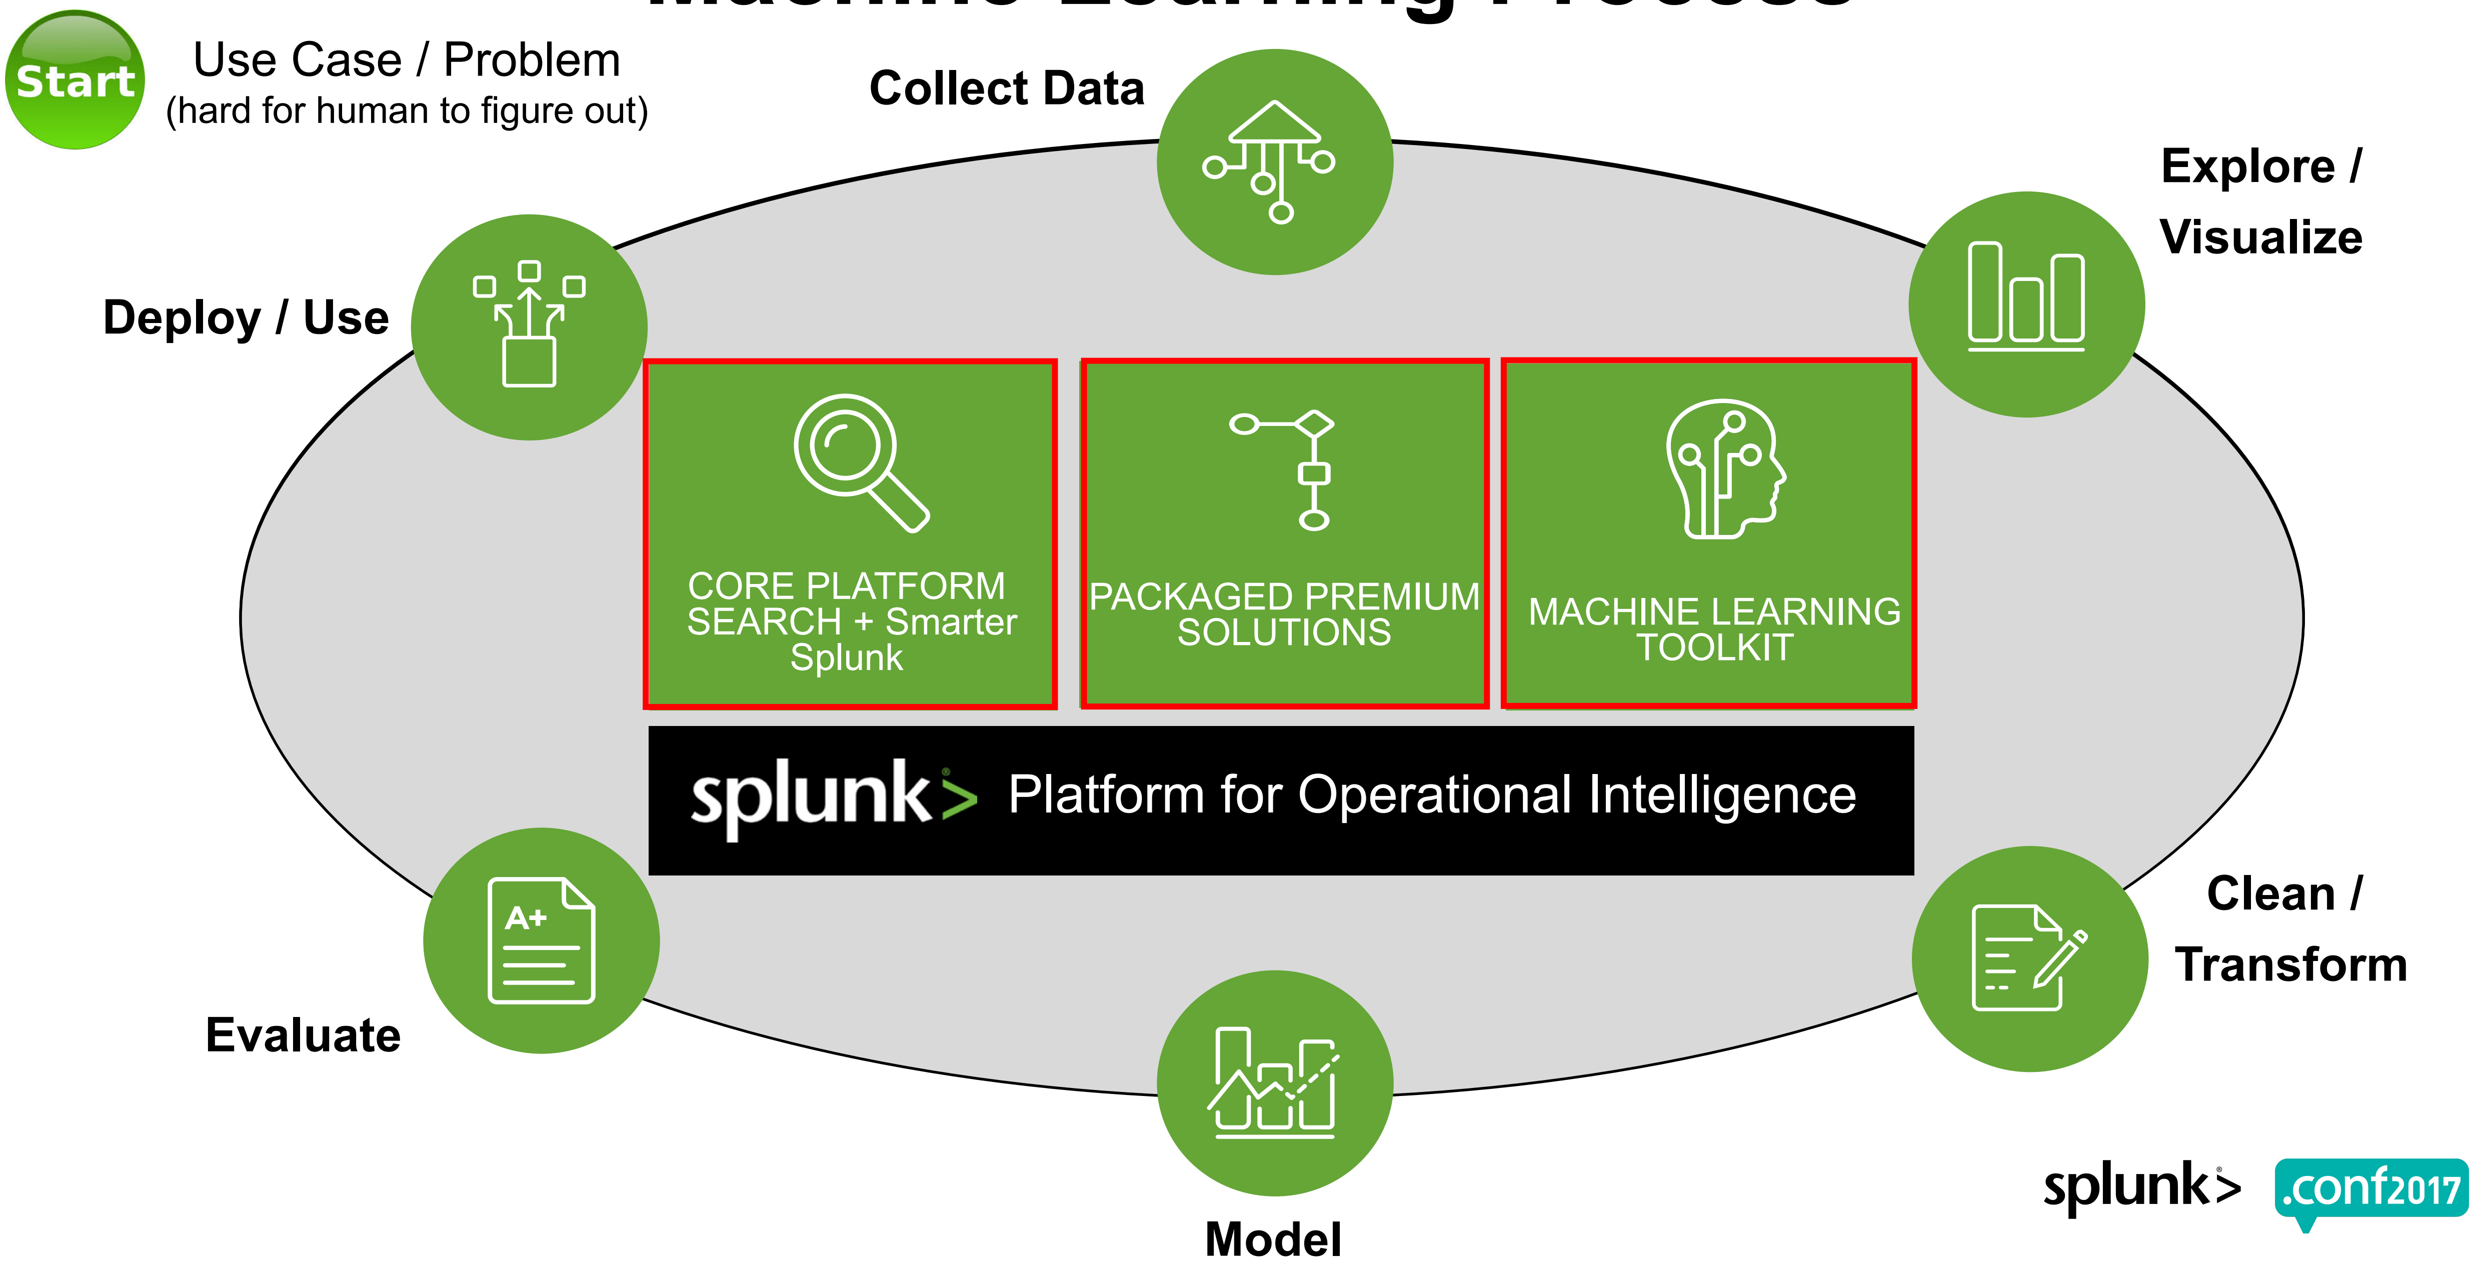

Start

Use Case / Problem
(hard for human to figure out)

Collect Data
Explore / Visualize
Clean / Transform
Model
Evaluate
Deploy / Use
CORE PLATFORM SEARCH + Smarter Splunk
PACKAGED PREMIUM SOLUTIONS
MACHINE LEARNING TOOLKIT
splunk> Platform for Operational Intelligence

Machine Learning Toolkit

DEMO – Machine Learning Toolkit

Send me an email alert when Server Power consumption is way outside of predicted values

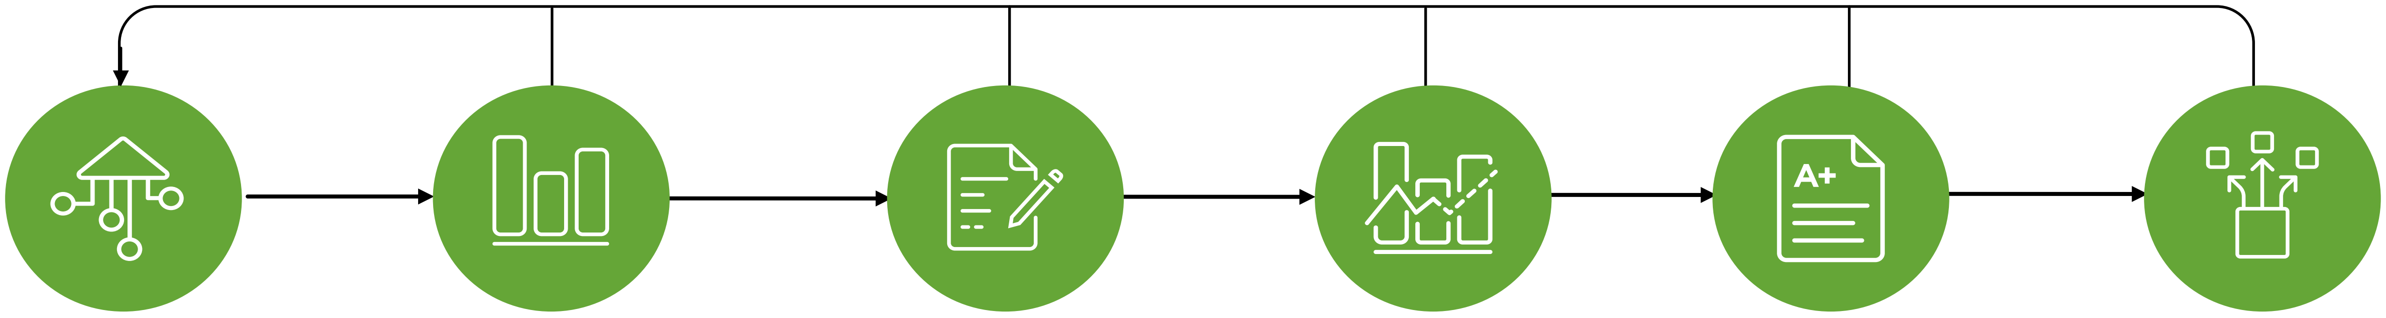

Collect Data

Input server_pow.csv

Explore / Visualize

Checkout the data

Clean / Transform

Normalize the data (apply pre-processing step)

Model

Predict power consumption
Detect numeric outliers

Evaluate

Deploy / Use

Setup thresholds, alerts, and actions

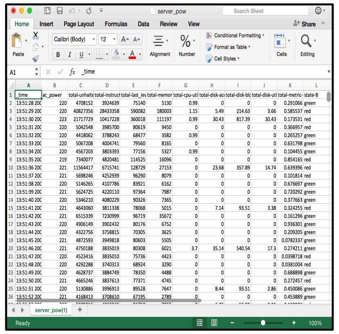

_time

	A	B	C	D	E	F	G	H	I	J	K	L
1	_time	ac_power	total-unhalte	total-instruct	total-last_lev	total-memor	total-cpu-uti	total-disk-ac	total-disk-blc	total-disk-uti	total-metric-	state-B
2	13:51:28 20(	220	4708152	3924639	75140	5130	0.99	0	0	0	0.291066	green
3	13:51:29 20(	220	40827356	28433358	590082	180003	1.15	5.49	214.63	3.66	0.585537	red
4	13:51:30 20(	223	21717729	10417228	360018	111197	0.99	30.43	817.39	30.43	0.173531	red
5	13:51:31 20(	220	5042548	3985700	80619	9450	0	0	0	0	0.366957	red
6	13:51:32 20(	220	4418062	3788243	68477	3582	0.99	0	0	0	0.265257	green
7	13:51:33 20(	220	5067208	4004741	79560	8165	0	0	0	0	0.631798	green
8	13:51:34 20(	220	4567203	3803393	77156	5327	0.99	0	0	0	0.104455	green
9	13:51:35 20(	219	7340077	4820481	114525	16096	0	0	0	0	0.854165	red
10	13:51:36 20(	221	11564417	6715741	128729	27153	0	23.68	357.89	14.74	0.639396	red
11	13:51:37 20(	221	5698246	4252939	96290	8079	0	0	0	0	0.101814	red
12	13:51:38 20(	220	5146265	4107786	83921	6162	0	0	0	0	0.676697	green
13	13:51:39 20(	221	5624725	4220110	97364	7987	0	0	0	0	0.720292	green
14	13:51:40 20(	220	5346210	4080229	90326	7365	0	0	0	0	0.377663	green
15	13:51:41 20(	221	4643060	3811338	78068	5015	0	7.14	93.51	3.38	0.324255	red
16	13:51:42 20(	221	6515339	7230999	96719	35672	0	0	0	0	0.161296	green
17	13:51:43 20(	220	4906149	3902432	80176	6752	0	0	0	0	0.936301	green
18	13:51:44 20(	220	4322756	3758815	70305	3625	0	0	0	0	0.209205	green
19	13:51:45 20(	221	4872593	3949818	80603	5505	0	0	0	0	0.0782337	green
20	13:51:46 20(	221	4750188	3835019	80308	6021	3.7	35.14	540.54	17.3	0.274211	green
21	13:51:47 20(	220	4523416	3835010	75736	4423	0	0	0	0	0.0398718	red
22	13:51:48 20(	220	4292288	3740313	68924	3290	0	0	0	0	0.0381004	red
23	13:51:49 20(	220	4628737	3884749	78350	4488	0	0	0	0	0.688898	green
24	13:51:50 20(	221	4665246	3837613	77371	4745	0	0	0	0	0.272457	red
25	13:51:51 20(	220	5130886	3996913	89528	7647	0	8.44	93.51	2.86	0.450086	green
26	13:51:52 20(	221	4168413	3708610	67195	2789	0	0	0	0	0.453889	green

server_pow(1)

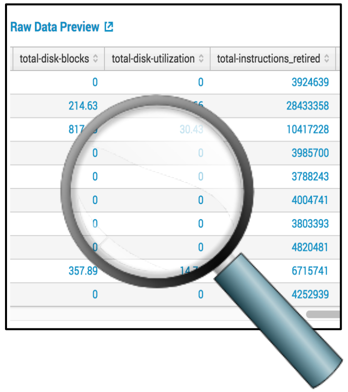

Raw Data Preview

total-disk-blocks	total-disk-utilization	total-instructions_retired
0	0	3924639
214.63	[illegible]	28433358
817.39	30.43	10417228
0	0	3985700
0	0	3788243
0	0	4004741
0	0	3803393
0	0	4820481
357.89	14.7	6715741
0	0	4252939

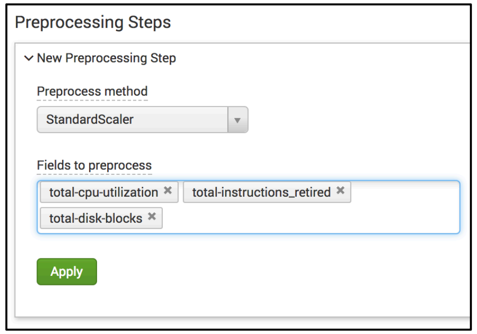

Preprocessing Steps

New Preprocessing Step

Preprocess method: StandardScaler

Fields to preprocess: total-cpu-utilization, total-instructions_retired, total-disk-blocks

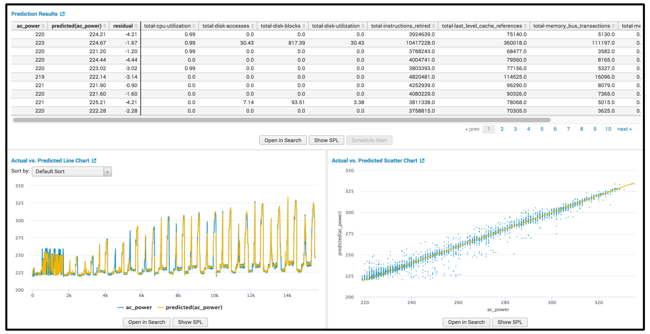

Prediction Results

ac_power	predicted(ac_power)	residual	total-cpu-utilization	total-disk-accesses	total-disk-blocks	total-disk-utilization	total-instructions_retired	total-last_level_cache_references	total-memory_bus_transactions	total-me
220	224.21	-4.21	0.99	0.0	0.0	0.0	3924639.0	75140.0	5130.0	0.
223	224.67	-1.67	0.99	30.43	817.39	30.43	10417228.0	360018.0	111197.0	0.
220	221.20	-1.20	0.99	0.0	0.0	0.0	3788243.0	68477.0	3582.0	0.
220	224.44	-4.44	0.0	0.0	0.0	0.0	4004741.0	79560.0	8165.0	0.
220	223.02	-3.02	0.99	0.0	0.0	0.0	3803393.0	77156.0	5327.0	0.
219	222.14	-3.14	0.0	0.0	0.0	0.0	4820481.0	114525.0	16096.0	0.
221	221.90	-0.90	0.0	0.0	0.0	0.0	4252939.0	96290.0	8079.0	0.
220	221.60	-1.60	0.0	0.0	0.0	0.0	4080229.0	90326.0	7365.0	0.
221	225.21	-4.21	0.0	7.14	93.51	3.38	3811338.0	78068.0	5015.0	0.
220	222.28	-2.28	0.0	0.0	0.0	0.0	3758815.0	70305.0	3625.0	0.

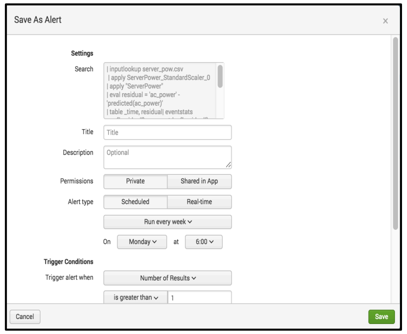

Save As Alert

Settings

Search:

```
| inputlookup server_pow.csv
 | apply ServerPower_StandardScaler_0
| apply "ServerPower"
| eval residual = 'ac_power' -
'predicted(ac_power)'
| table _time, residual| eventstats
```

Title: Title

Description: Optional

Permissions: Private | Shared in App

Alert type: Scheduled | Real-time

Run every week

On Monday at 6:00

Trigger Conditions

Trigger alert when: Number of Results

is greater than 1

Machine Learning Toolkit – What's New

Enhanced Detect Numeric Outliers assistant

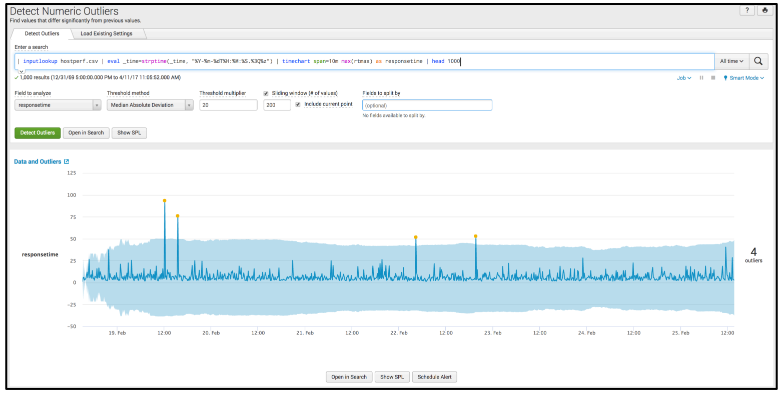

Detect Numeric Outliers

Find values that differ significantly from previous values.

Detect Outliers | Load Existing Settings

Enter a search

```
| inputlookup hostperf.csv | eval _time=strptime(_time, "%Y-%m-%dT%H:%M:%S.%3Q%z") | timechart span=10m max(rtmax) as responsetime | head 1000
```

✓ 1,000 results (12/31/69 5:00:00.000 PM to 4/11/17 11:05:52.000 AM)

Field to analyze: responsetime

Threshold method: Median Absolute Deviation

Threshold multiplier: 20

☑ Sliding window (# of values): 200

☑ Include current point

Fields to split by: (optional)

No fields available to split by.

Segment outlier detection by field

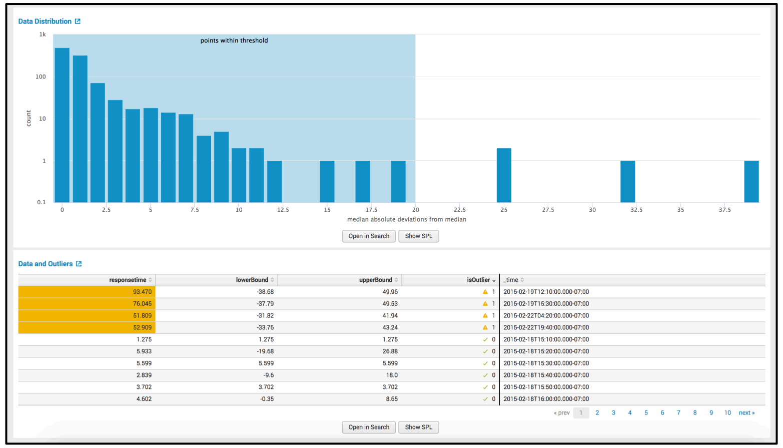

responsetime	lowerBound	upperBound	isOutlier	_time
93.470	-38.68	49.96	⚠ 1	2015-02-19T12:10:00.000-07:00
76.045	-37.79	49.53	⚠ 1	2015-02-19T15:30:00.000-07:00
51.809	-31.82	41.94	⚠ 1	2015-02-22T04:20:00.000-07:00
52.909	-33.76	43.24	⚠ 1	2015-02-22T19:40:00.000-07:00
1.275	1.275	1.275	✓ 0	2015-02-18T15:10:00.000-07:00
5.933	-19.68	26.88	✓ 0	2015-02-18T15:20:00.000-07:00
5.599	5.599	5.599	✓ 0	2015-02-18T15:30:00.000-07:00
2.839	-9.6	18.0	✓ 0	2015-02-18T15:40:00.000-07:00
3.702	3.702	3.702	✓ 0	2015-02-18T15:50:00.000-07:00
4.602	-0.35	8.65	✓ 0	2015-02-18T16:00:00.000-07:00

New visualizations, including data distribution histogram

Expanded pre-processing capabilities

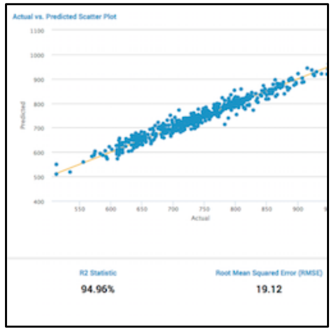

R2 Statistic: 94.96%

Root Mean Squared Error (RMSE): 19.12

Predict Numeric Fields

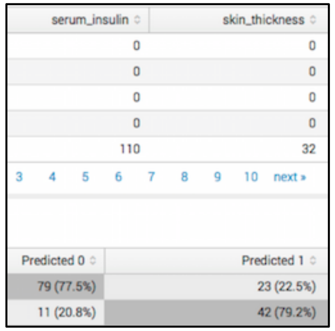

serum_insulin	skin_thickness
0	0
0	0
0	0
0	0
110	32

Predicted 0	Predicted 1
79 (77.5%)	23 (22.5%)
11 (20.8%)	42 (79.2%)

Predict Categorical Fields

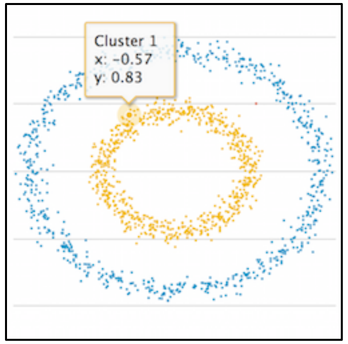

Cluster Numeric Events

ML-SPL API

Makes it easier to import open source and proprietary algorithms

splunk> .conf2017

Machine Learning Toolkit – What's New

Model Management (MLTK 3.0)

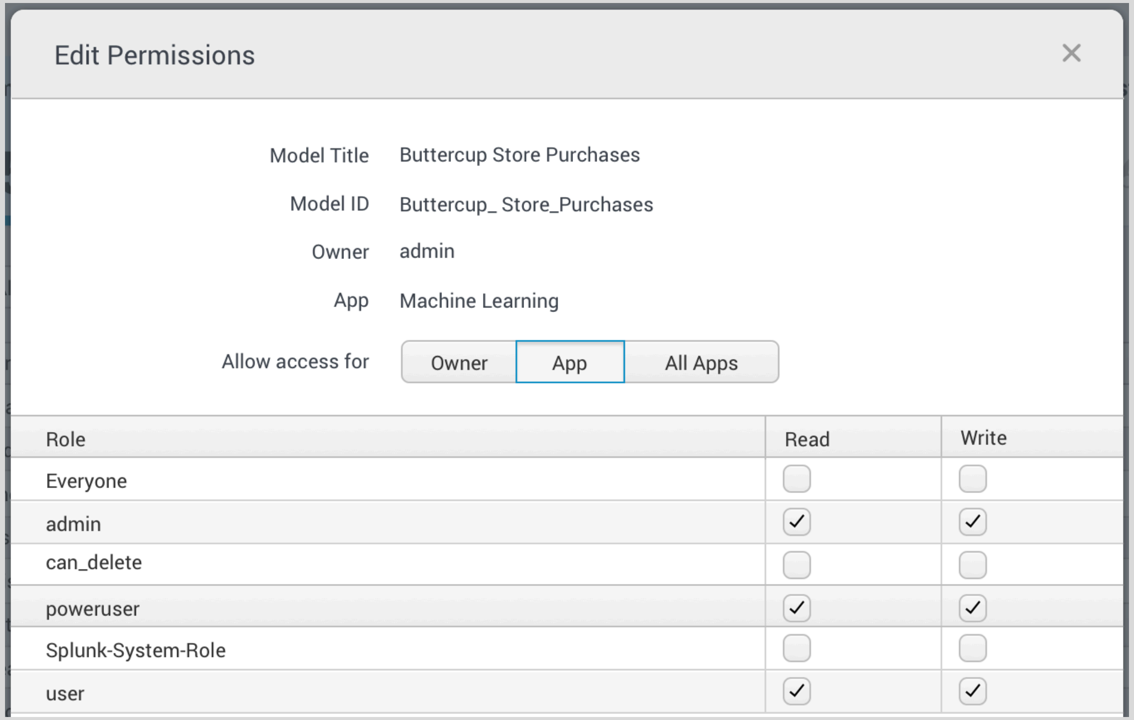

Edit Permissions

Model Title: Buttercup Store Purchases

Model ID: Buttercup_ Store_Purchases

Owner: admin

App: Machine Learning

Allow access for: Owner [✓ App] All Apps

Role	Read	Write
Everyone	☐	☐
admin	☑	☑
can_delete	☐	☐
poweruser	☑	☑
Splunk-System-Role	☐	☐
user	☑	☑

Assign permissions to models to control access

Spark Support (private beta open now)

- Use your existing Spark cluster with MLTK
  - Distributed fit on massive datasets
  - Apply MLlib models for supported algos
- Go to Splunk Innovation Labs (NDA only)
- Contact mlprogram@splunk.com
  - What is your use case (e.g., predicting server downtime)?
  - How are you using Spark today?

Leverage your Spark Infrastructure

Monitoring & Investigation

Increased productivity for a wider range of users

Native Metrics Ingestion & Event Integration
Event Annotation
Chart Enhancements
Report Actions
Trellis Layout
Dashboard Drilldown editor
Dashboard Search Bar
Search Productivity Enhancements
Table Datasets Explorer
AWS Kinesis Firehose Support (Private Beta)

Performance, Scale & Manageability

Faster, more scalable and easier to manage than ever before

New High Speed Metrics Engine
(20x faster)
Improved SPL Optimizer
Faster Data Model Acceleration
Improved Indexer Scalability (5M buckets)
Improved SHC/IC Resiliency and Management
Orphaned knowledge object detection
New Data Ingestion Management views
Cloud Self-Service App Installation (private apps too)
New App Packaging Toolkit

Machine Learning & Intelligence

More powerful and extensible machine learning capabilities

Fully Integrated ML across portfolio
Enhanced Detect Numeric Outliers assistant
New visualizations
Expanded pre-processing capabilities
Extended ML-SPL API
Additional Algorithms
Model Management (3.0)
Spark Support (Private Beta)

….plus lots more we didn’t get a chance to cover

Splunk Enterprise 7.0 Available Now

Plus check out the Overview Apps for What's New

Splunk 6.6 Overview App

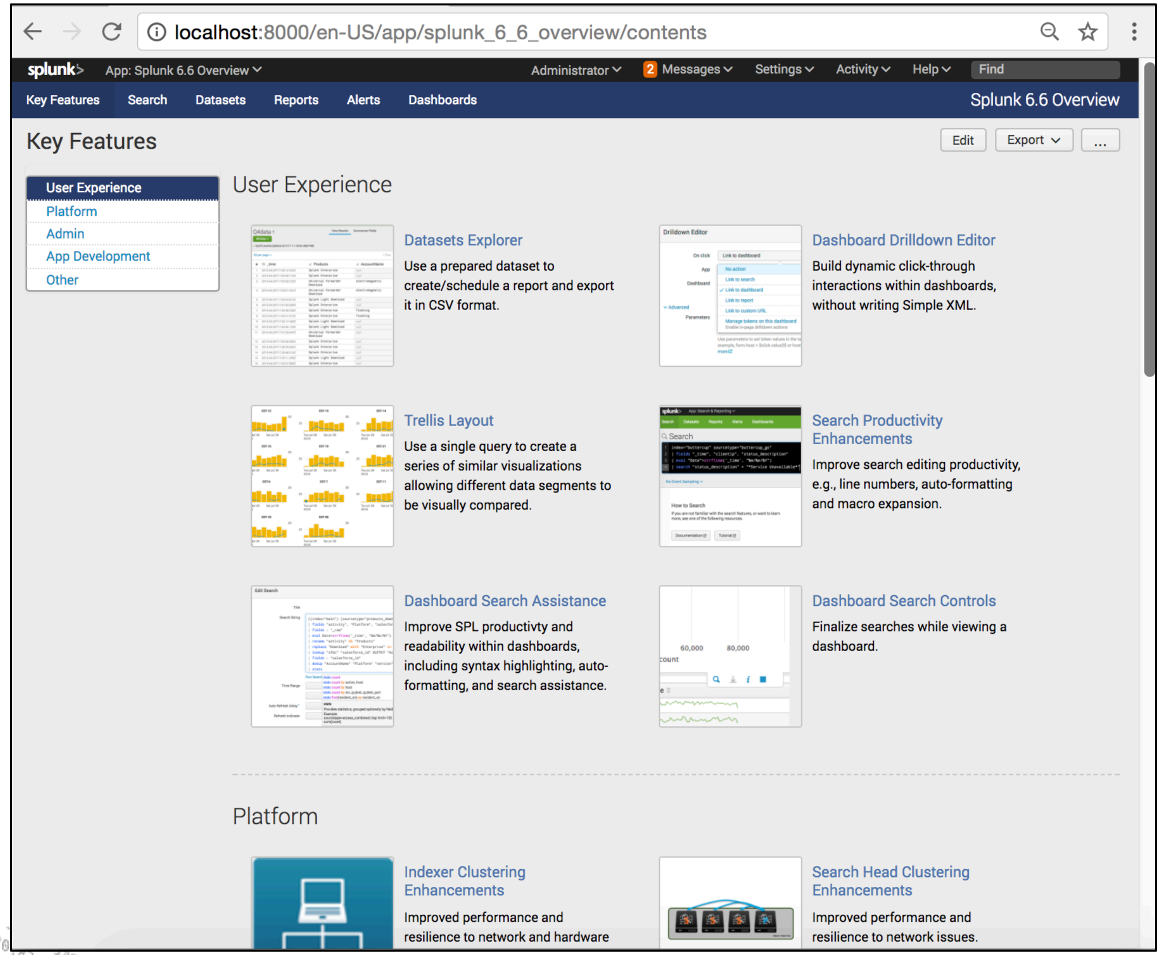

localhost:8000/en-US/app/splunk_6_6_overview/contents

Key Features

User Experience

Datasets Explorer
Use a prepared dataset to create/schedule a report and export it in CSV format.

Dashboard Drilldown Editor
Build dynamic click-through interactions within dashboards, without writing Simple XML.

Trellis Layout
Use a single query to create a series of similar visualizations allowing different data segments to be visually compared.

Search Productivity Enhancements
Improve search editing productivity, e.g., line numbers, auto-formatting and macro expansion.

Dashboard Search Assistance
Improve SPL productivty and readability within dashboards, including syntax highlighting, auto-formatting, and search assistance.

Dashboard Search Controls
Finalize searches while viewing a dashboard.

Platform

Indexer Clustering Enhancements
Improved performance and resilience to network and hardware

Search Head Clustering Enhancements
Improved performance and resilience to network issues.

Splunk 7.0 Overview App

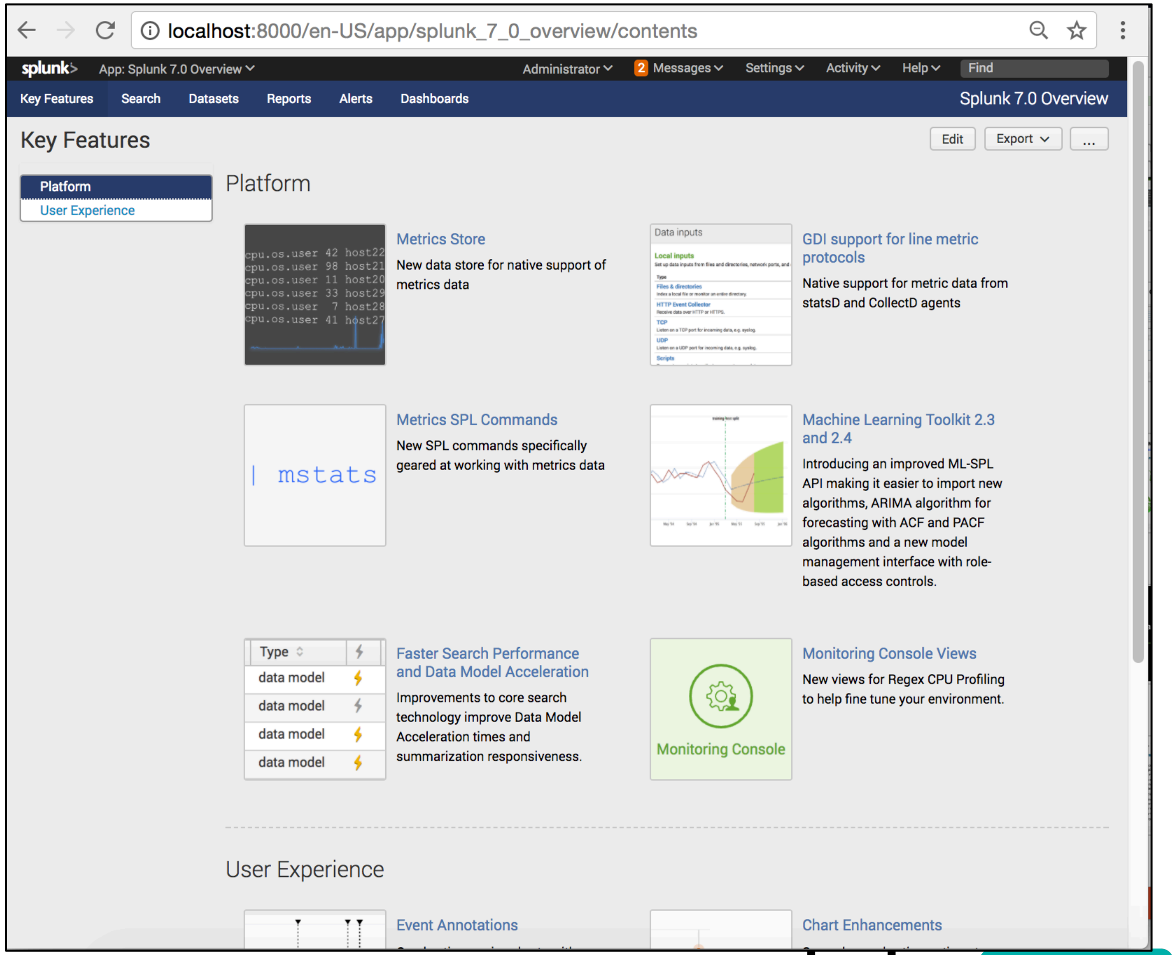

localhost:8000/en-US/app/splunk_7_0_overview/contents

Key Features

Platform

Metrics Store
New data store for native support of metrics data

GDI support for line metric protocols
Native support for metric data from statsD and CollectD agents

Metrics SPL Commands
New SPL commands specifically geared at working with metrics data

Machine Learning Toolkit 2.3 and 2.4
Introducing an improved ML-SPL API making it easier to import new algorithms, ARIMA algorithm for forecasting with ACF and PACF algorithms and a new model management interface with role-based access controls.

Faster Search Performance and Data Model Acceleration
Improvements to core search technology improve Data Model Acceleration times and summarization responsiveness.

Monitoring Console Views
New views for Regex CPU Profiling to help fine tune your environment.

User Experience

Event Annotations

Chart Enhancements

INNOVATION

Visit Splunk Labs

Room: 301-303
Tues 10:30am – 2:30pm, 4:30pm – 6:30pm
Wed 10:30am – 3:00pm

splunk> .conf2017

Thank You

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

splunk> .conf2017