What's New Splunk Enterprise & Splunk Cloud

.conf19

splunk>

Skip Bacon
VP | Splunk Enterprises

Sundeep Gupta
Director | Splunk Cloud

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.

Splunk, Splunk>, Turn Data Into Doing, 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. © 2019 Splunk Inc. All rights reserved.

Splunk Delivers a Complete Portfolio for Turning Data into Business Outcomes

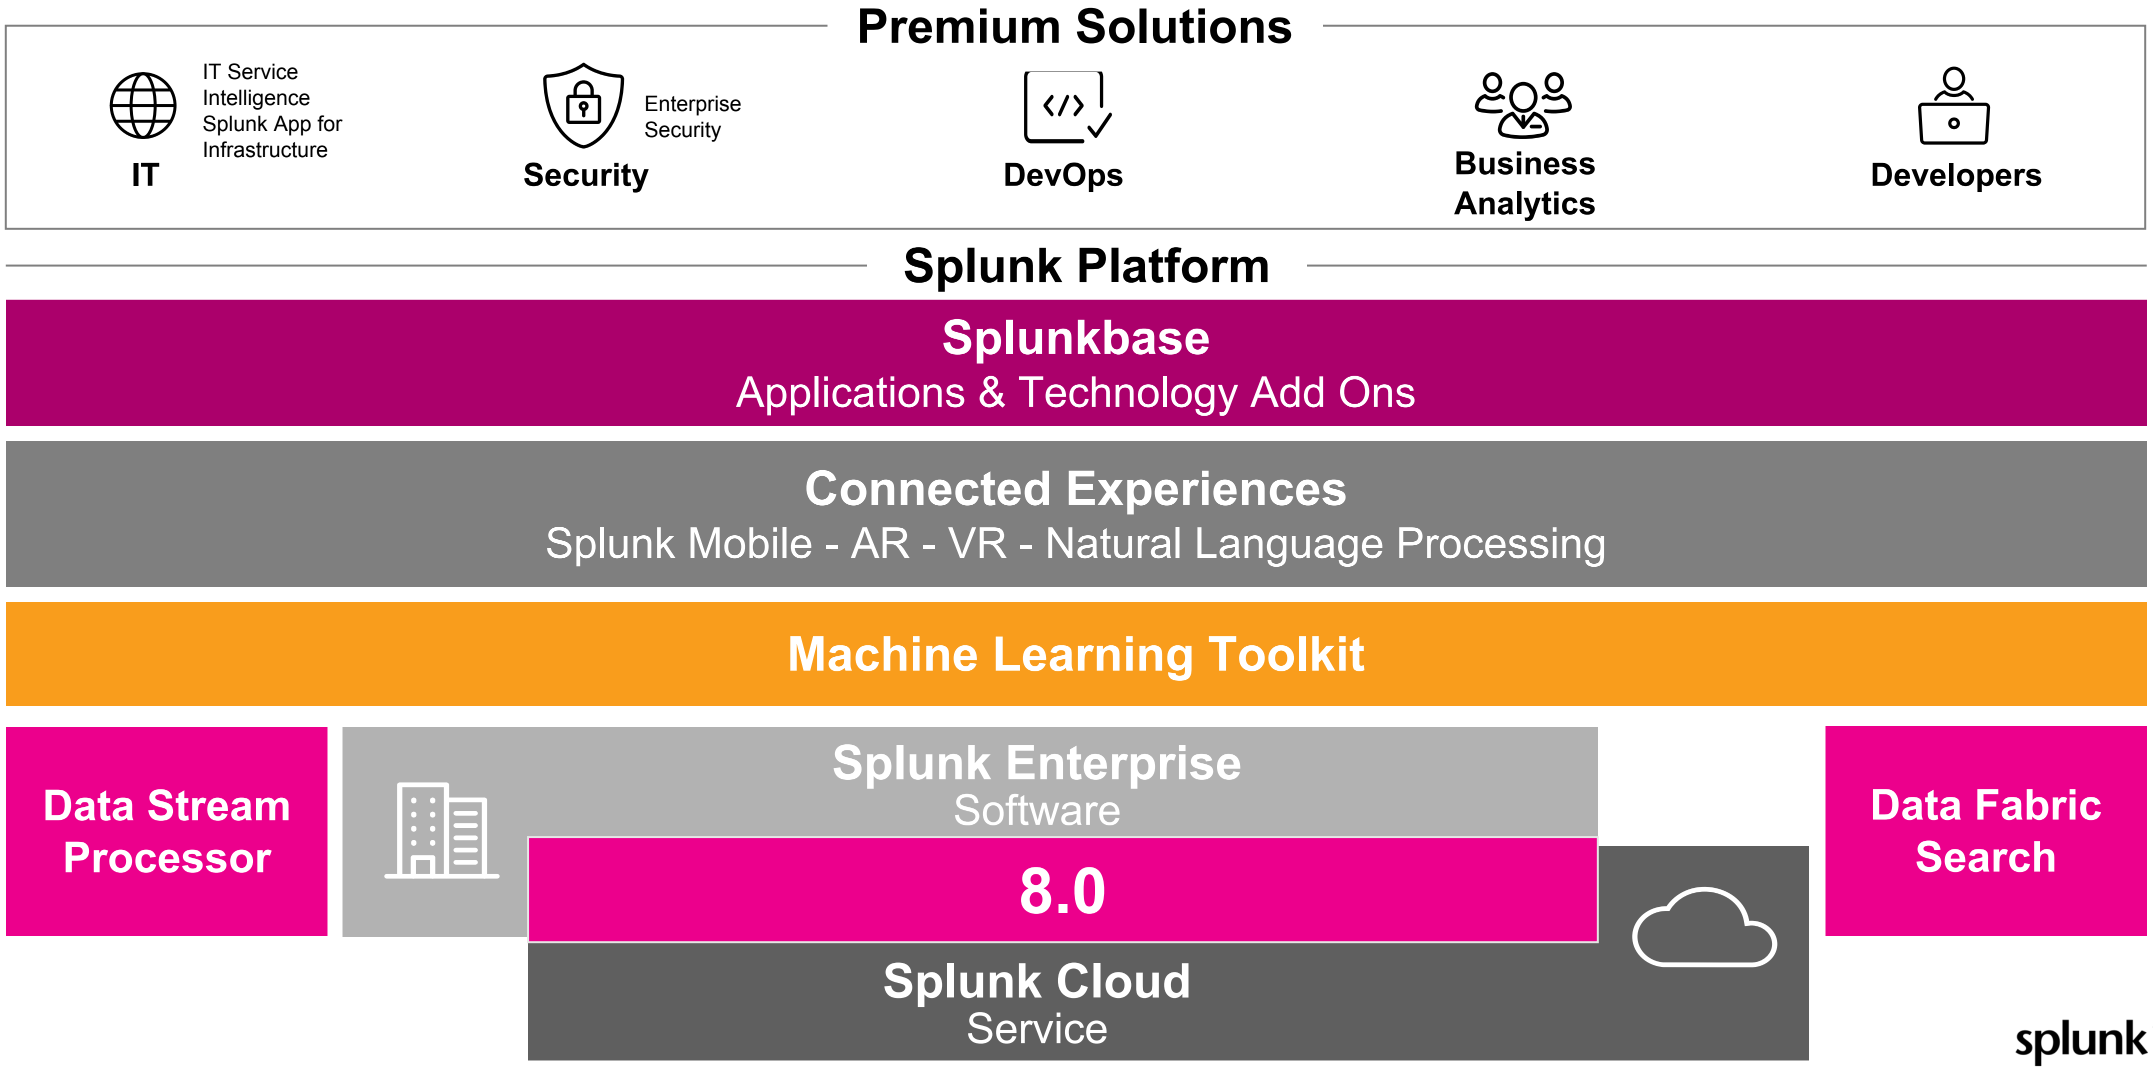

Premium Solutions

- IT: IT Service Intelligence, Splunk App for Infrastructure
- Security: Enterprise Security
- DevOps
- Business Analytics
- Developers

Splunk Platform

Splunkbase
Applications & Technology Add Ons

Connected Experiences
Splunk Mobile - AR - VR - Natural Language Processing

Machine Learning Toolkit

Data Stream Processor

Splunk Enterprise
Software

8.0

Splunk Cloud
Service

Data Fabric Search

What’s New: .conf19

Splunk Enterprise 8.0 and Splunk Cloud

Richer Insights for More Users

Expansive Data Access

Splunk Cloud for More Customers

Operability At Scale

What's New: .conf19

Splunk Enterprise 8.0 and Splunk Cloud

Richer Insights for More Users

- Splunk Dashboards
- Analytics Workspace
- Field Access Controls
- Python 3.7

Expansive Data Access

Splunk Cloud for More Customers

Operability At Scale

Richer

Splunk Dashboard s

Engaged Users

- **Beta Splunkbase app**
- **Pixel-perfect** dashboard layouts enable rich data storytelling
- **Images and background graphics** enhance presentation, expand business context
- Enables **common user experience** across Platform and Premium Solutions
- **Power visualizations** with SPL, as always!

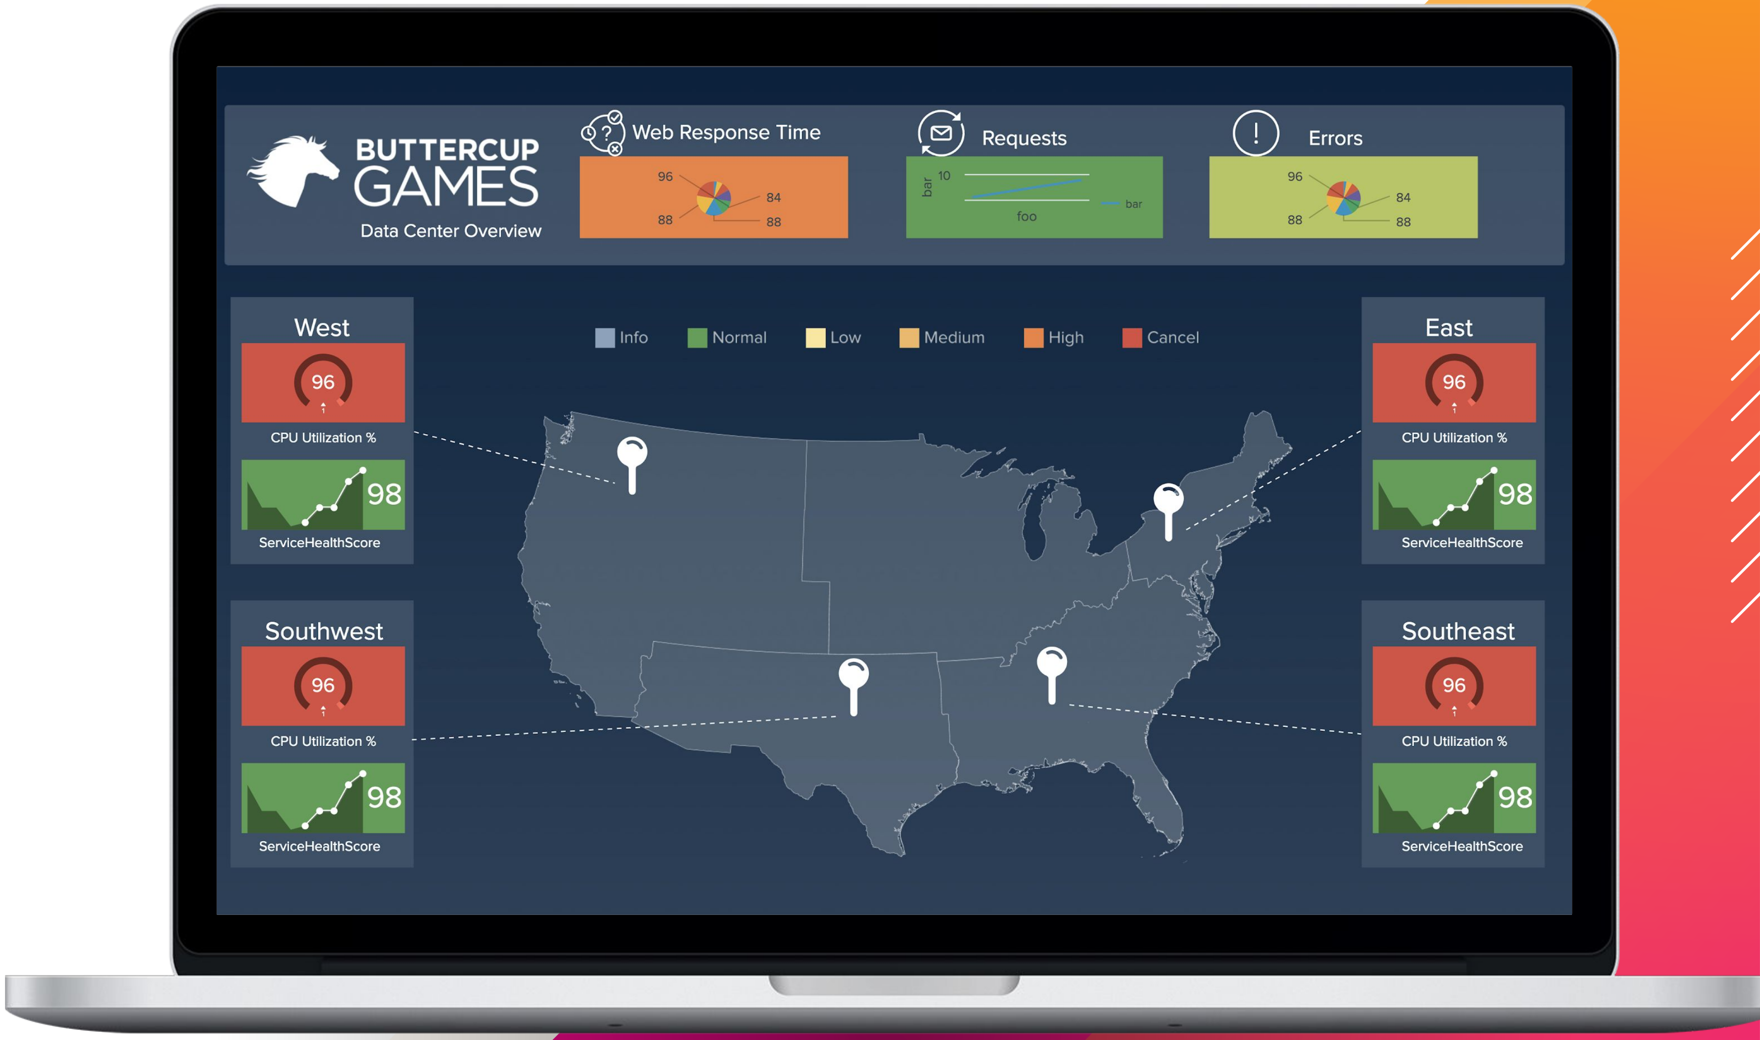

Splunk Enterprise

Splunk Cloud

Richer

Splunk Dashboards

Productive Developers

- **Editor UI** streamlines creation with familiar canvas metaphor
- Reduce **time-to-dashboard** including beautiful visualizations
- Operationalize dashboards with extended **take-action** capabilities
- Enable UI-based tokenization and filtering to **simplify entry points** to complex dashboards

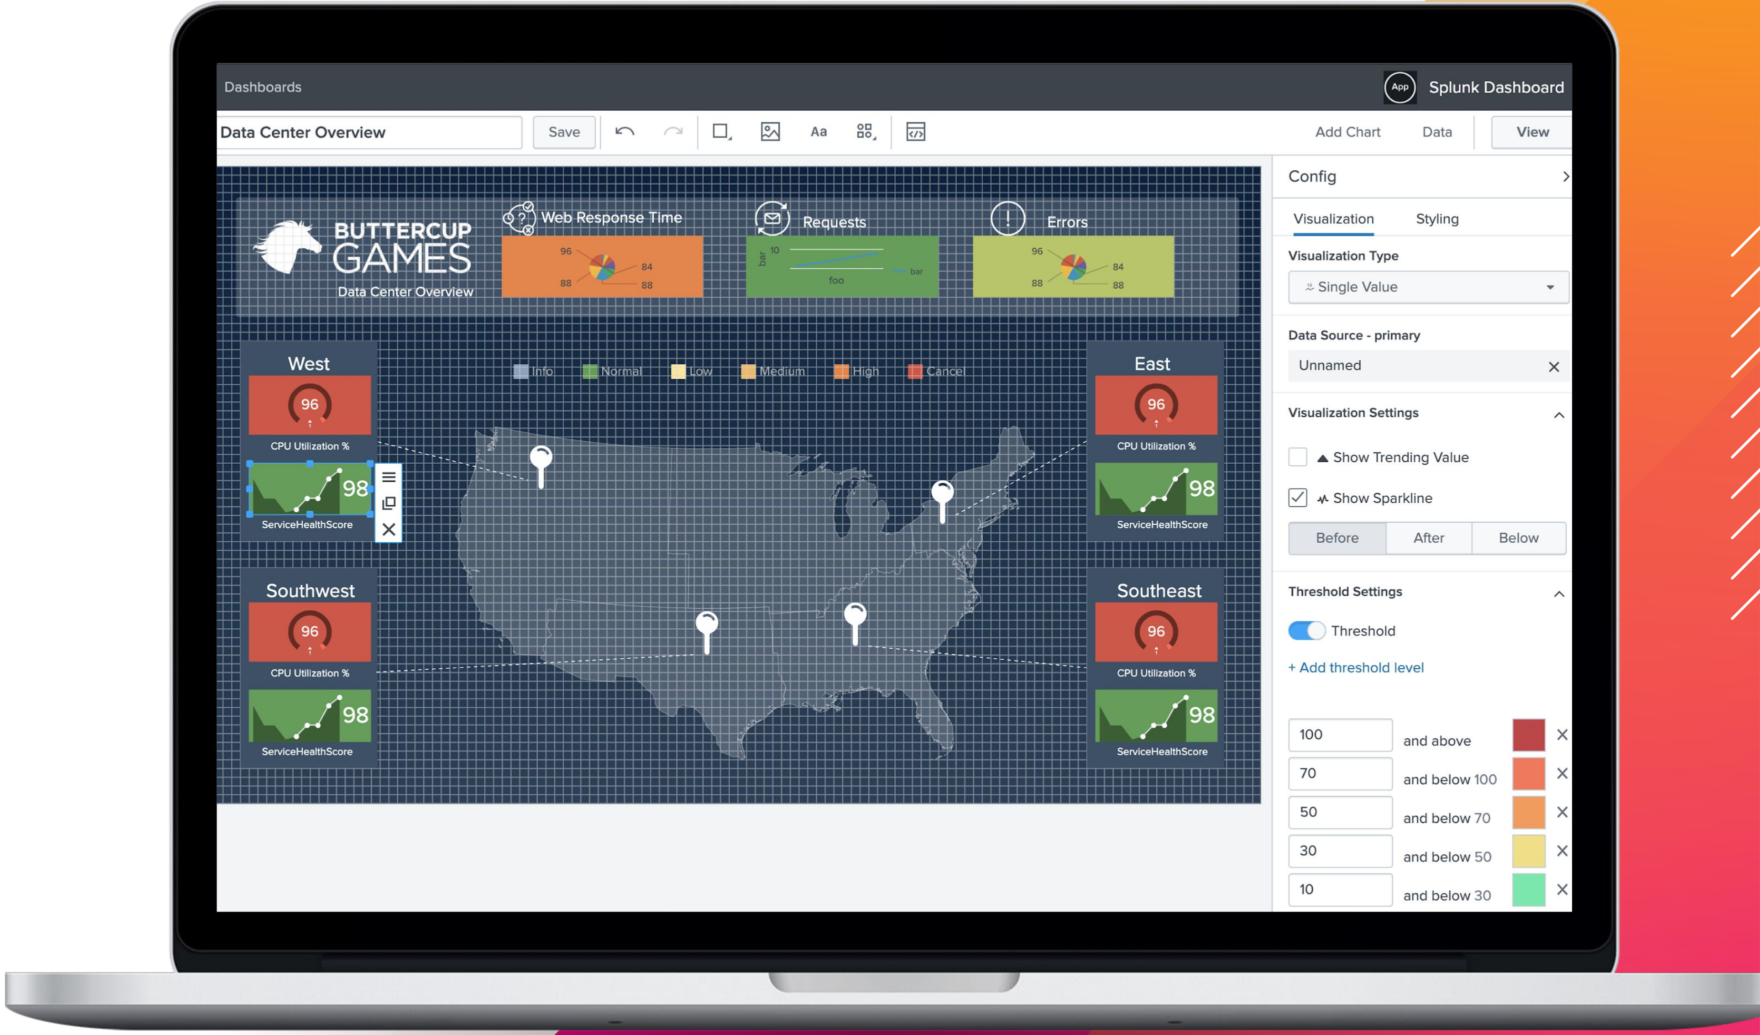

Config

Visualization | Styling

Visualization Type: Single Value

Data Source - primary: Unnamed ✕

Visualization Settings

☐ Show Trending Value

☑ Show Sparkline

Before | After | Below

Threshold Settings

● Threshold

+ Add threshold level

100	and above	🟥	✕
70	and below 100	🟥	✕
50	and below 70	🟧	✕
30	and below 50	🟨	✕
10	and below 30	🟩	✕

Splunk Enterprise

Splunk Cloud

Richer

Analytics Workspace

Self-Service Insights.
No SPL.

- **Next generation** of 7.x Metrics Workspace app
- **Visualize time-series data** in multiple ways: line, bar, categorical charts
- Include **aggregated dimensions and categories** for metrics and datasets
- Get to **desired data faster** with flexible time-range picker

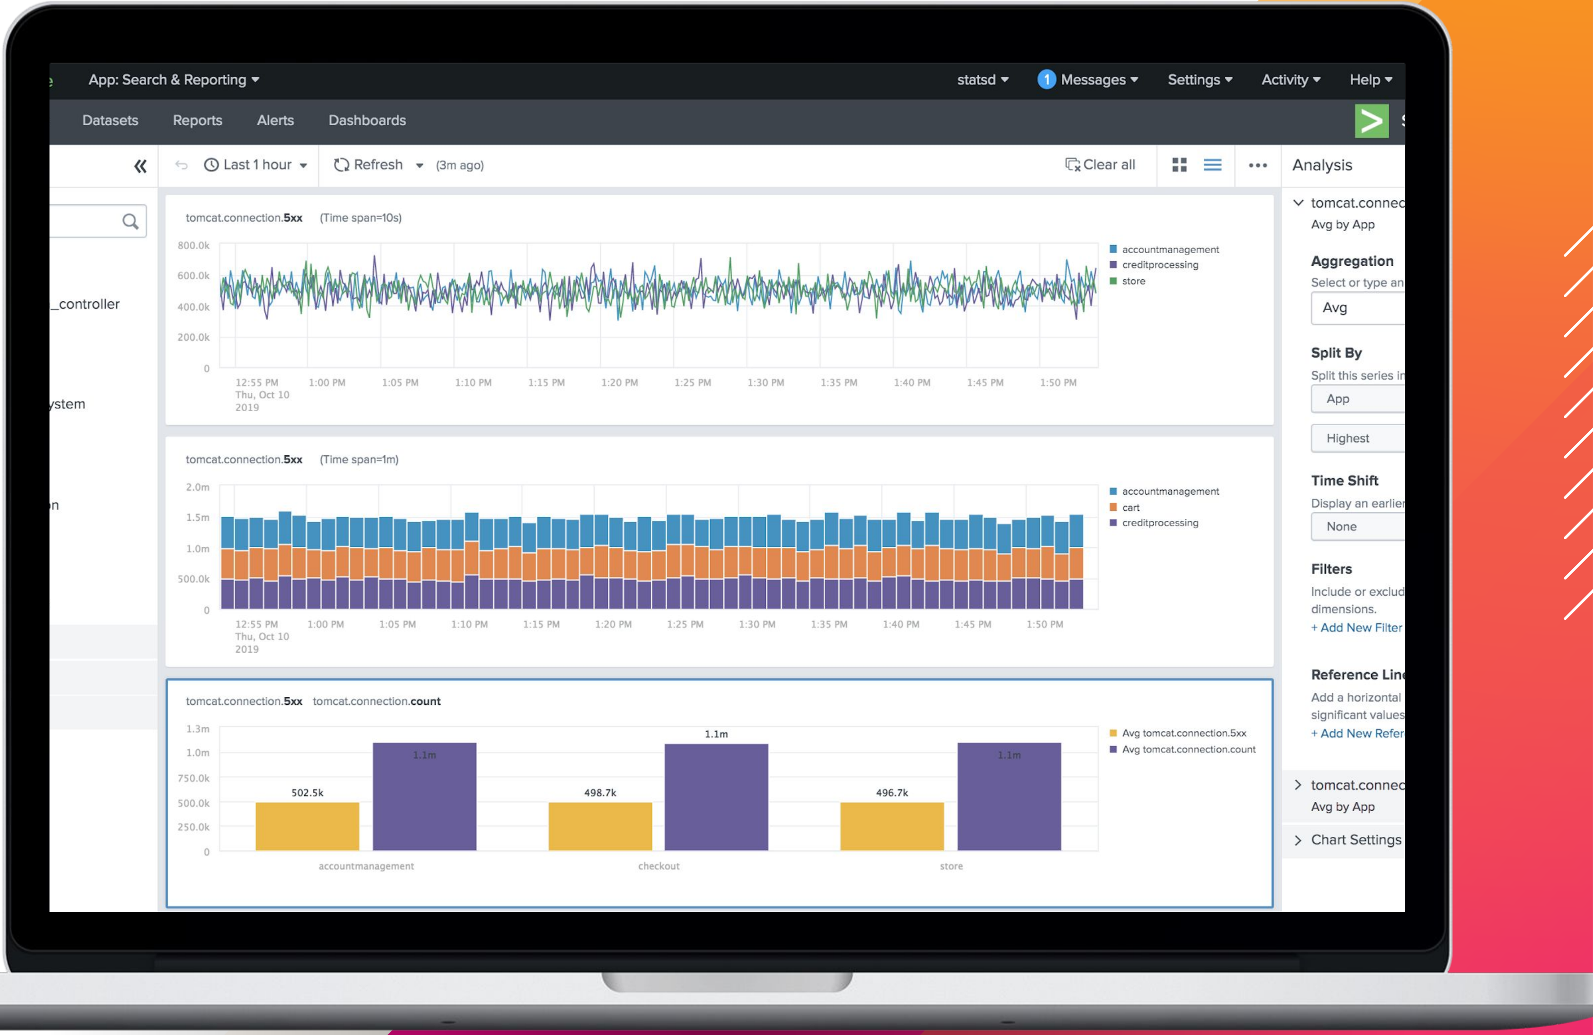

Splunk Enterprise

Splunk Cloud

splunk> .conf19

Richer

Analytics Workspace

Events + Metrics.
Peanut Butter + Jelly.

- Analyze **Alerts, Metrics, and Events (beta feature)** together
- Create better-performing **streaming metrics alerts** in a few quick steps directly from visual analysis
- **Visually compare variations** with statistical metrics
- Create **more types of dashboard panels and reports**

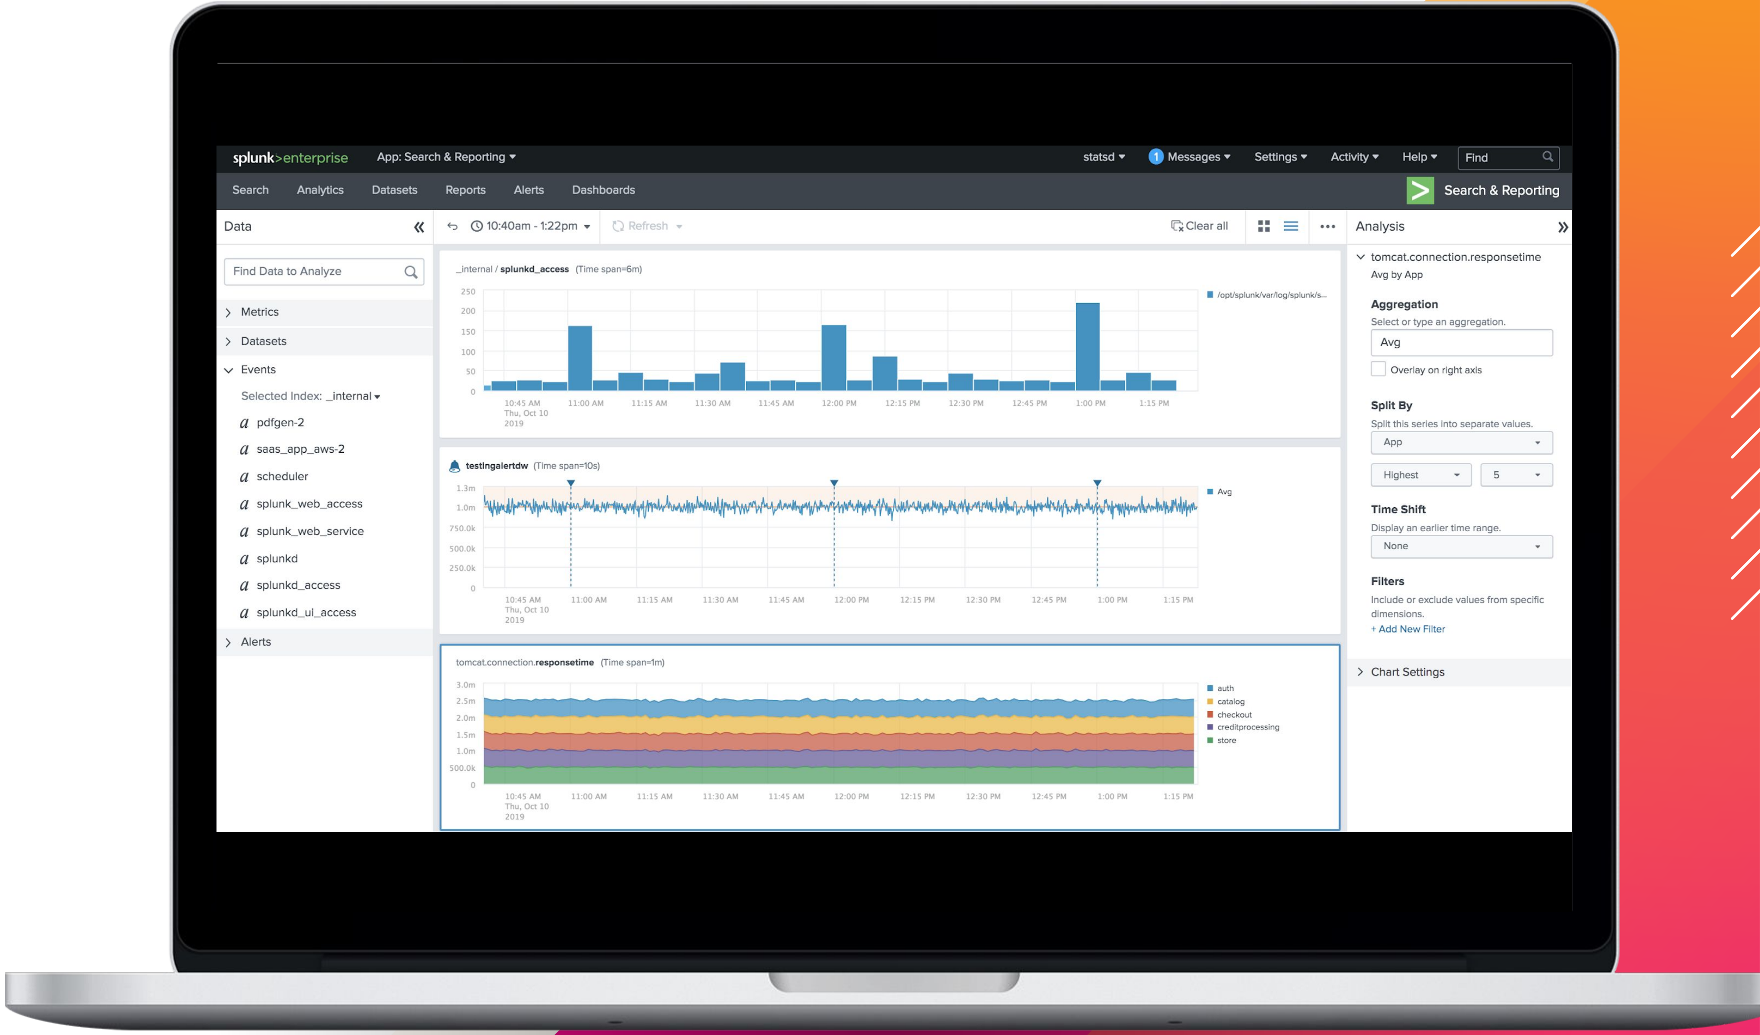

Splunk Enterprise

Splunk Cloud

splunk> .conf19

Indexed Field Access Controls

Share (Not Overshare) Data.

- Add field access controls to roles using the **dynamic Search Filter generator**
- **Limit event access within an index** using source type, source, or any other indexed field::value combination
- Enables **secure separation** across **diverse users** on same indexes
- **Reduce indexes and data duplication** for smaller, more manageable deployments

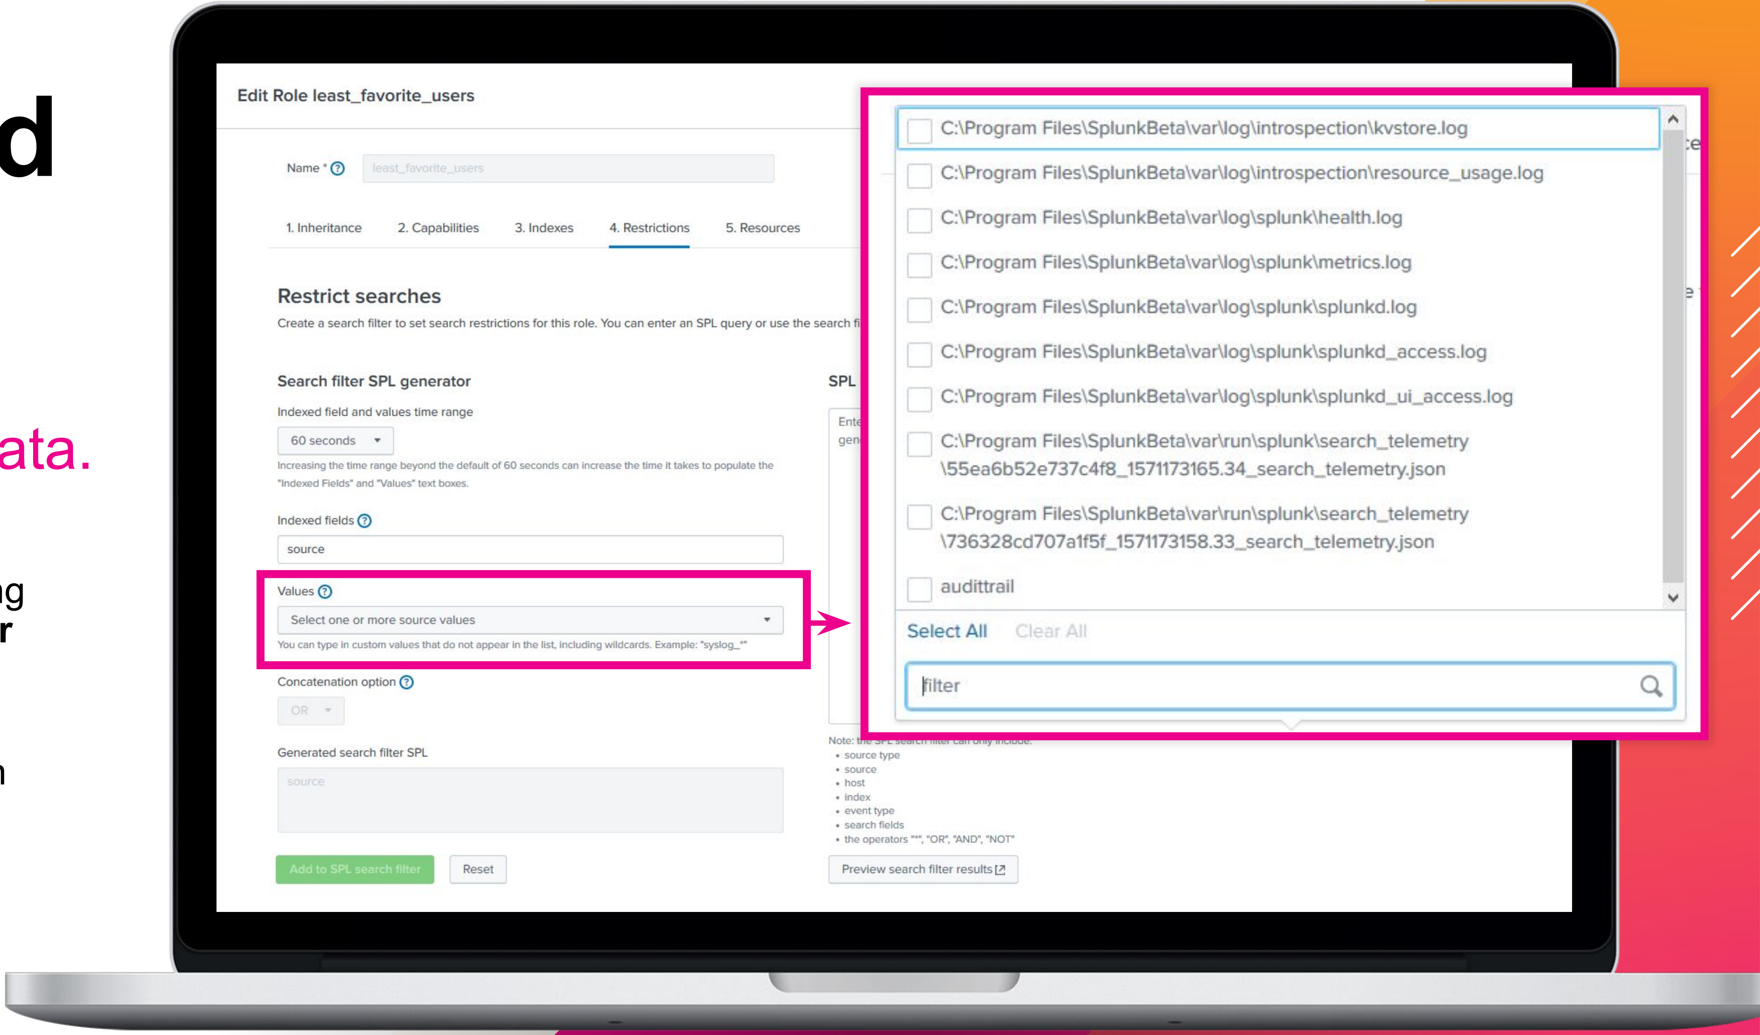

Edit Role least_favorite_users

Name * least_favorite_users

1. Inheritance 2. Capabilities 3. Indexes 4. Restrictions 5. Resources

Restrict searches

Create a search filter to set search restrictions for this role. You can enter an SPL query or use the search fi

Search filter SPL generator

Indexed field and values time range: 60 seconds

Increasing the time range beyond the default of 60 seconds can increase the time it takes to populate the "Indexed Fields" and "Values" text boxes.

Indexed fields: source

Values: Select one or more source values

You can type in custom values that do not appear in the list, including wildcards. Example: "syslog_*"

- ☐ C:\Program Files\SplunkBeta\var\log\introspection\kvstore.log
- ☐ C:\Program Files\SplunkBeta\var\log\introspection\resource_usage.log
- ☐ C:\Program Files\SplunkBeta\var\log\splunk\health.log
- ☐ C:\Program Files\SplunkBeta\var\log\splunk\metrics.log
- ☐ C:\Program Files\SplunkBeta\var\log\splunk\splunkd.log
- ☐ C:\Program Files\SplunkBeta\var\log\splunk\splunkd_access.log
- ☐ C:\Program Files\SplunkBeta\var\log\splunk\splunkd_ui_access.log
- ☐ C:\Program Files\SplunkBeta\var\run\splunk\search_telemetry\55ea6b52e737c4f8_1571173165.34_search_telemetry.json
- ☐ C:\Program Files\SplunkBeta\var\run\splunk\search_telemetry\736328cd707a1f5f_1571173158.33_search_telemetry.json
- ☐ audittrail

Select All Clear All

filter

Concatenation option: OR

Generated search filter SPL: source

SPL

Ente gen

Note: the SPL search filter can only include:

- source type
- source
- host
- index
- event type
- search fields
- the operators "*", "OR", "AND", "NOT"

Splunk Enterprise

Splunk Cloud

Python 3 Support

Expanded Customizability

- Enterprise 8.0 runs **natively on Python 3.7**, also supports 2.7
- Premium Solutions and Splunk-supported Apps being **upgraded for 3.7 compatibility**
- **Use Upgrade Readiness App** to assess 7.x dependencies. Use **Six** and **Future** libraries to make apps dual-version compatible
- **Python 2 supported** for lifecycle of Enterprise 6.6/7.x; removed from post-8.0 release

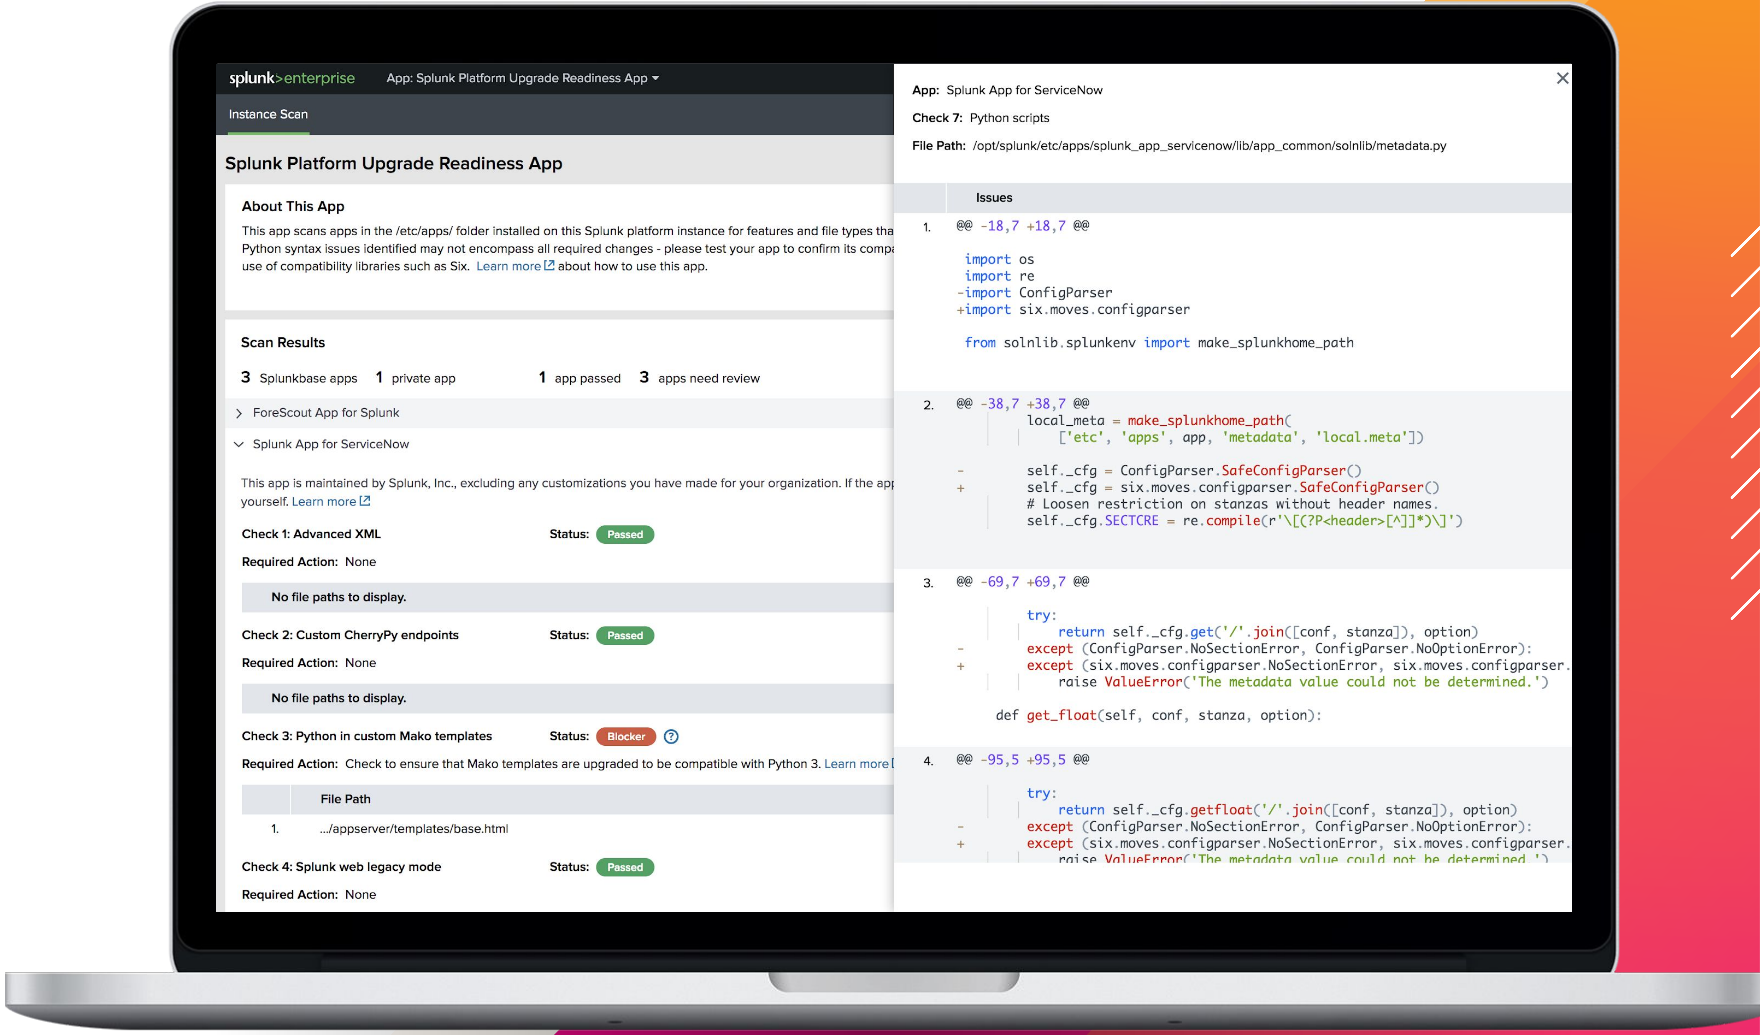

Splunk Platform Upgrade Readiness App

About This App

This app scans apps in the /etc/apps/ folder installed on this Splunk platform instance for features and file types tha
Python syntax issues identified may not encompass all required changes - please test your app to confirm its compa
use of compatibility libraries such as Six. Learn more about how to use this app.

Scan Results

3 Splunkbase apps 1 private app 1 app passed 3 apps need review

› ForeScout App for Splunk

⌄ Splunk App for ServiceNow

This app is maintained by Splunk, Inc., excluding any customizations you have made for your organization. If the app
yourself. Learn more

Check 1: Advanced XML **Status:** Passed

Required Action: None

No file paths to display.

Check 2: Custom CherryPy endpoints **Status:** Passed

Required Action: None

No file paths to display.

Check 3: Python in custom Mako templates **Status:** Blocker ?

Required Action: Check to ensure that Mako templates are upgraded to be compatible with Python 3. Learn more

	File Path
1.	.../appserver/templates/base.html

Check 4: Splunk web legacy mode **Status:** Passed

Required Action: None

App: Splunk App for ServiceNow

Check 7: Python scripts

File Path: /opt/splunk/etc/apps/splunk_app_servicenow/lib/app_common/solnlib/metadata.py

Issues

1.
```
@@ -18,7 +18,7 @@

 import os
 import re
-import ConfigParser
+import six.moves.configparser

 from solnlib.splunkenv import make_splunkhome_path
```

2.
```
@@ -38,7 +38,7 @@
        local_meta = make_splunkhome_path(
            ['etc', 'apps', app, 'metadata', 'local.meta'])

-       self._cfg = ConfigParser.SafeConfigParser()
+       self._cfg = six.moves.configparser.SafeConfigParser()
        # Loosen restriction on stanzas without header names.
        self._cfg.SECTCRE = re.compile(r'\[(?P<header>[^]]*)\]')
```

3.
```
@@ -69,7 +69,7 @@

        try:
            return self._cfg.get('/'.join([conf, stanza]), option)
-       except (ConfigParser.NoSectionError, ConfigParser.NoOptionError):
+       except (six.moves.configparser.NoSectionError, six.moves.configparser.
            raise ValueError('The metadata value could not be determined.')

    def get_float(self, conf, stanza, option):
```

4.
```
@@ -95,5 +95,5 @@

        try:
            return self._cfg.getfloat('/'.join([conf, stanza]), option)
-       except (ConfigParser.NoSectionError, ConfigParser.NoOptionError):
+       except (six.moves.configparser.NoSectionError, six.moves.configparser.
            raise ValueError('The metadata value could not be determined.')
```

Splunk Enterprise

Splunk Cloud

What’s New: .conf19

Splunk Enterprise 8.0 and Splunk Cloud

Richer Insights for More Users

- Splunk Dashboards
- Analytics Workspace
- Field Access Controls
- Python 3.7

Expansive Data Access

- **Search Performance**
- **Data Model Summaries**
- **Metrics Optimization**

Splunk Cloud for More Customers

Operability At Scale

SearchEvaluator

Faster. More efficient. Cake. Eat.

- **SearchEvaluator filters results** to return specified event types
- **Much faster new algorithm**, adopted from Stanford University research
- **~30% faster** processing than 7.3
- **Reduces memory usage** for searches across multiple indexes

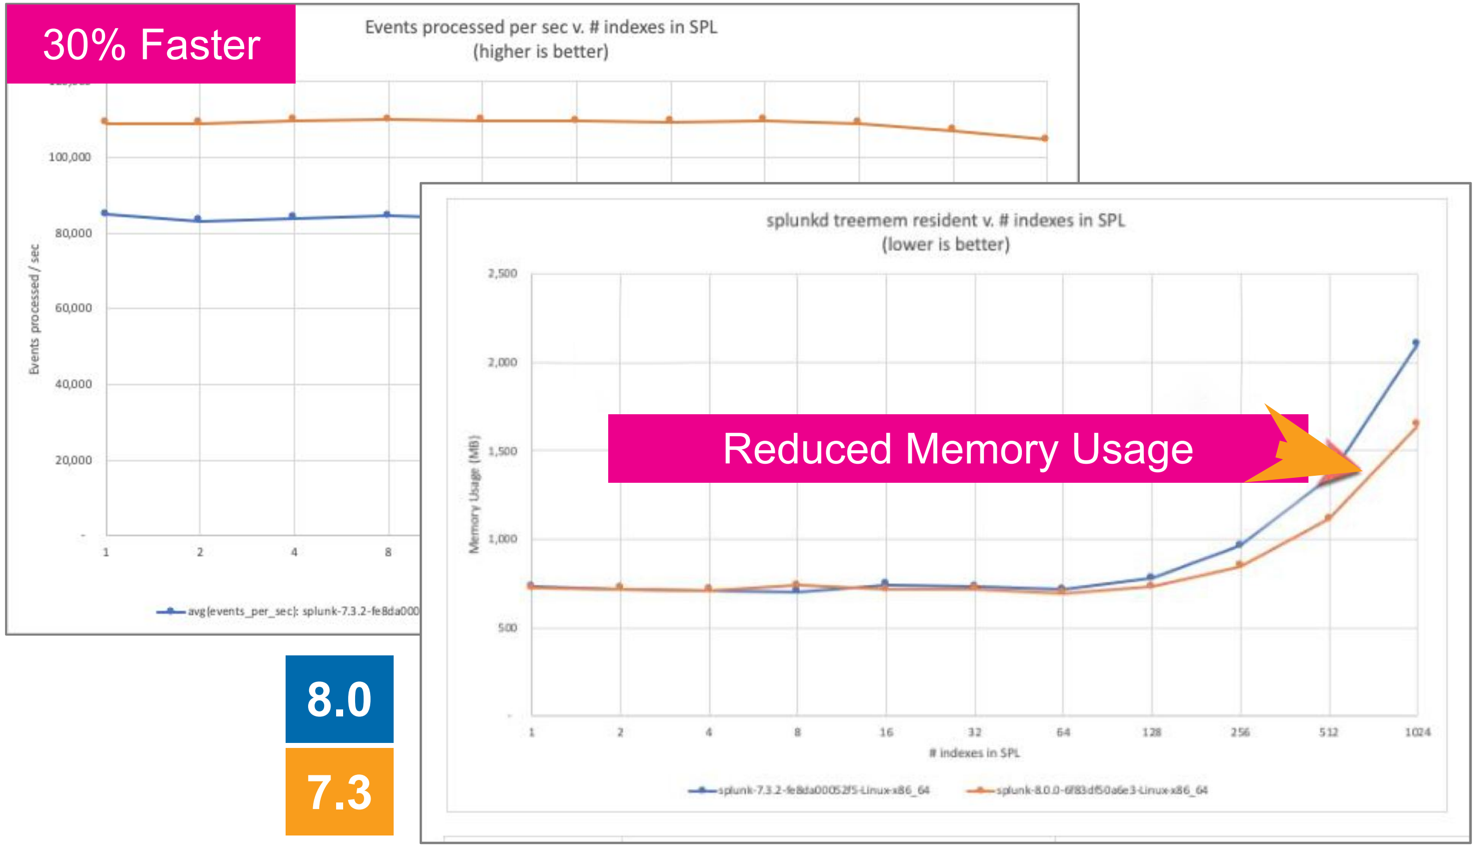

Splunk Enterprise

Splunk Cloud

Cited improvements from Splunk testing. Actual results may vary.

stats Command

Lean, mean analytics machine

- Rewrote one of **most-used search commands**
- Calculates **aggregate statistics**, such as average, count, and sum, over the search results set
- Also used in the reduce phase of commands like **mstats** and **tstats**
- **Aggregate performance improved 2x** over Splunk Enterprise 7.3
- **Processing time decreased up to 75%** for high-cardinality operations

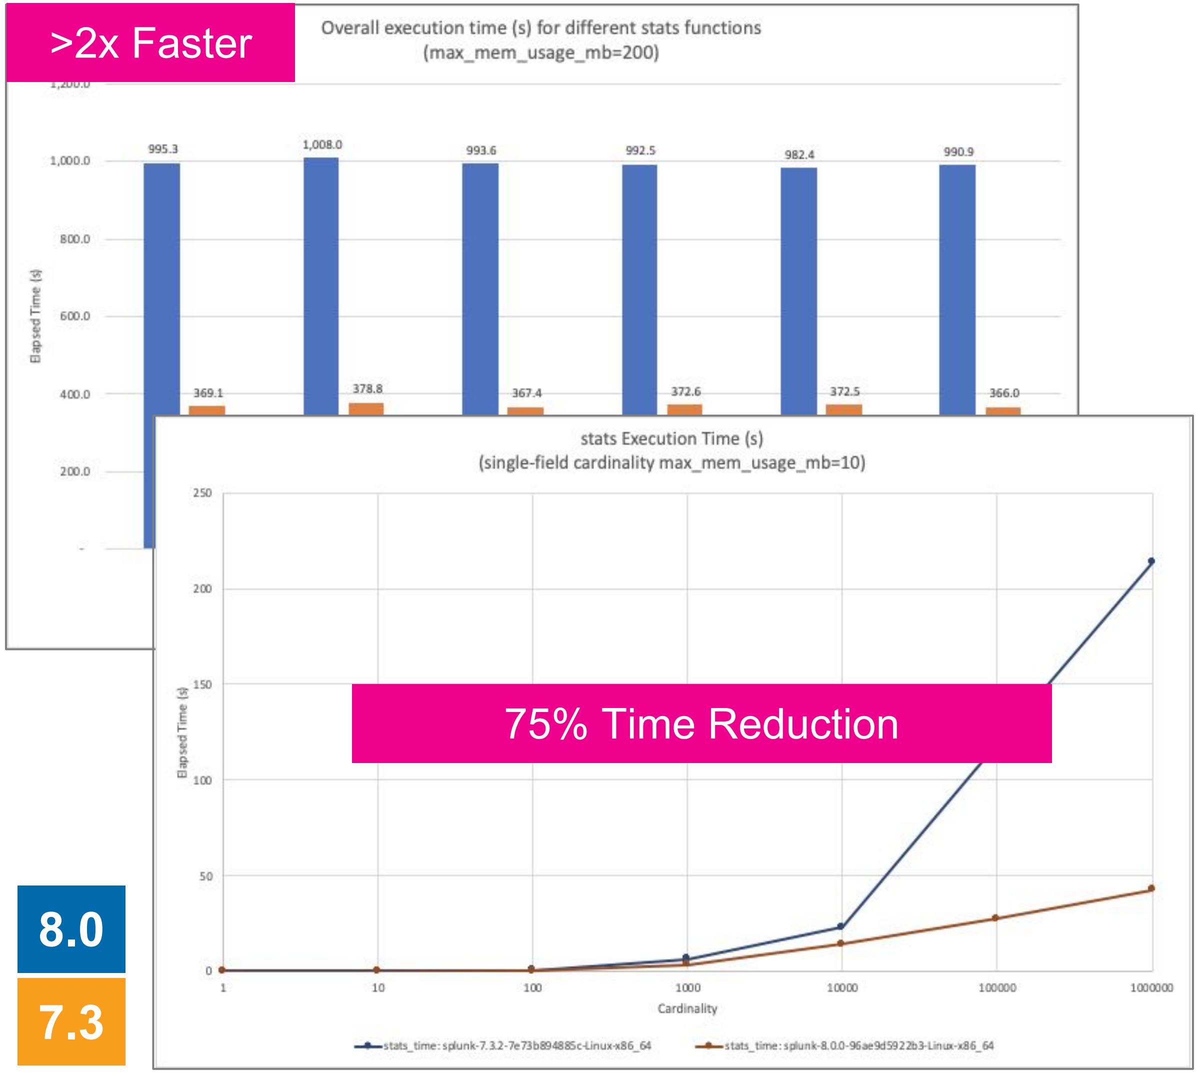
>2x Faster

75% Time Reduction

Splunk Enterprise

Splunk Cloud

Cited improvements from Splunk testing. Actual results may vary.

Faster Searches. Automatically.

tstats Optimizer

- **Automatically converts** “| from datamodel **stats**” to “| **tstats**” for supported aggregations

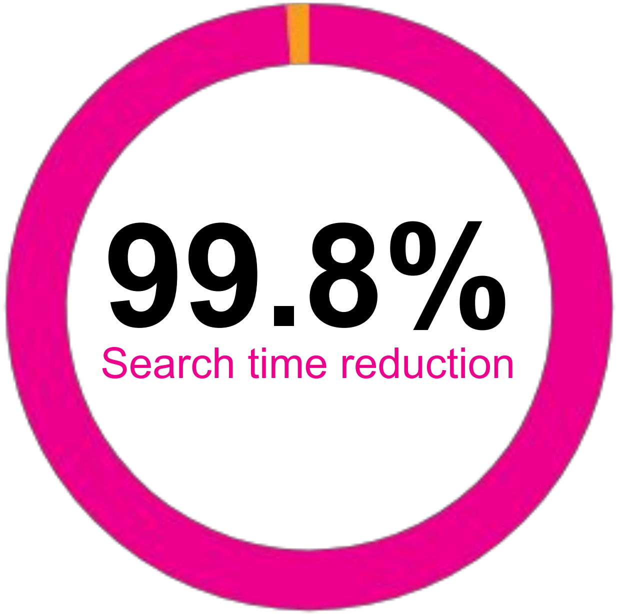

Shared lookup

- Automatically reuses **lookups** within search process for **faster processing** and **lower memory consumption**.

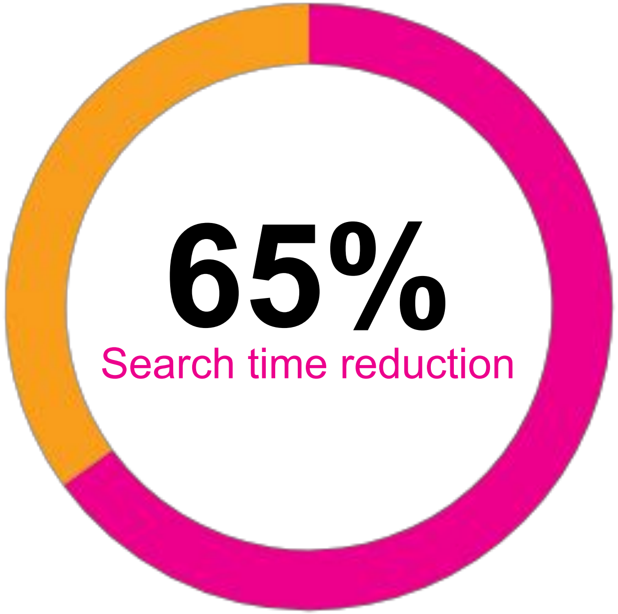

Splunk Enterprise

Splunk Cloud

Cited improvements from Splunk testing. Actual results may vary.

DMA Summary Sharing

Fast. Consistently and Efficiently.

- Data Model Acceleration (DMA) Summaries on indexers can be shared **across multiple search heads**
- Helps to ensure **consistent search results**
- Maintaining fewer summaries **reduces indexer cluster resource utilization**
- Can be **configured per data model**

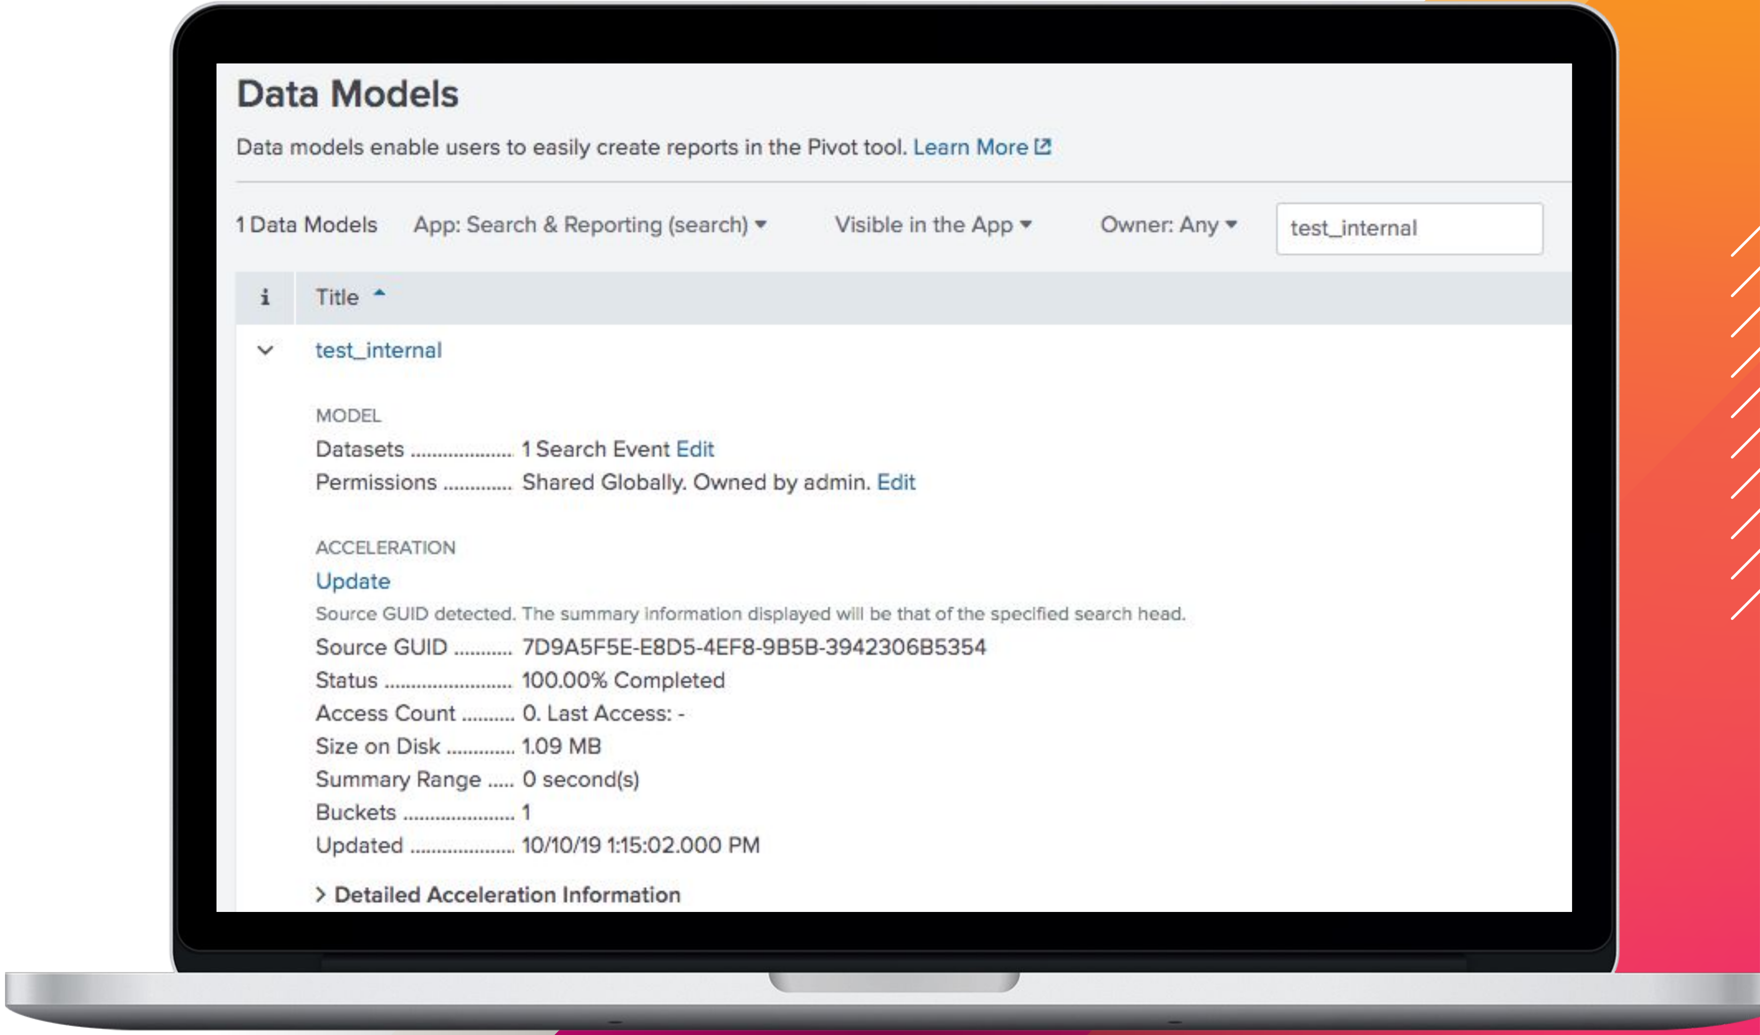

Data Models

Data models enable users to easily create reports in the Pivot tool. Learn More

1 Data Models   App: Search & Reporting (search) ▾   Visible in the App ▾   Owner: Any ▾   test_internal

i	Title ▴
⌄	test_internal

MODEL

Datasets .... 1 Search Event Edit

Permissions .... Shared Globally. Owned by admin. Edit

ACCELERATION

Update

Source GUID detected. The summary information displayed will be that of the specified search head.

Source GUID 7D9A5F5E-E8D5-4EF8-9B5B-3942306B5354

Status .... 100.00% Completed

Access Count .... 0. Last Access: -

Size on Disk .... 1.09 MB

Summary Range ..... 0 second(s)

Buckets 1

Updated .... 10/10/19 1:15:02.000 PM

› Detailed Acceleration Information

Splunk Enterprise

Splunk Cloud

Expansive Data: DMA Summaries

DMA Monitoring

Keep the wheels on

Deep monitoring helps **ensure robust operation**

Provides **detailed status** info...

- Access Count and Last Access date/time
- Size on Disk
- Summary Range
- Last Update time/date

… **Operating statistics**

- Last Run Statistics
- SSID, Start and Run time

… and **configuration details**

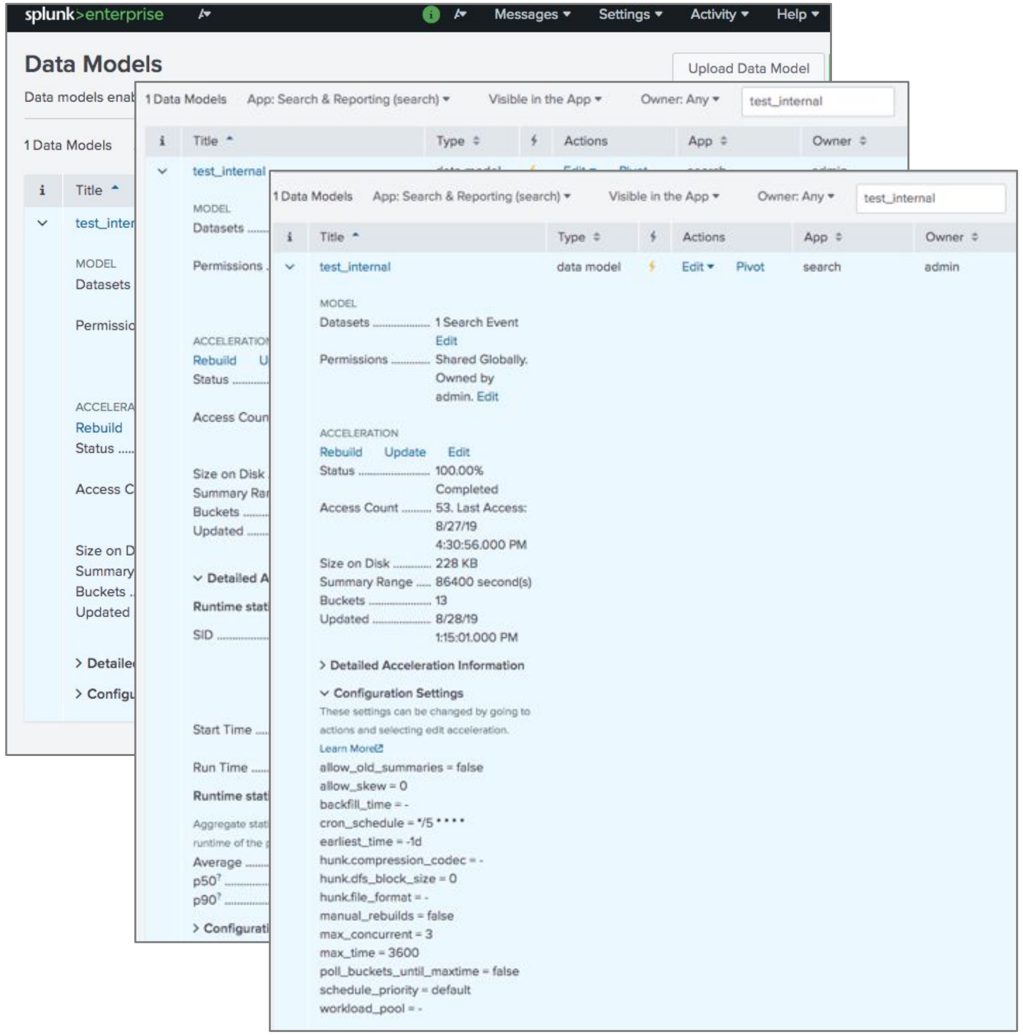

i	Title	Type	⚡	Actions	App	Owner
⌄	test_internal	data model	⚡	Edit ▾ Pivot	search	admin

MODEL
Datasets .... 1 Search Event
Edit
Permissions .... Shared Globally. Owned by admin. Edit

ACCELERATION
Rebuild Update Edit
Status .... 100.00% Completed
Access Count .... 53. Last Access: 8/27/19 4:30:56.000 PM
Size on Disk .... 228 KB
Summary Range ..... 86400 second(s)
Buckets 13
Updated .... 8/28/19 1:15:01.000 PM

› Detailed Acceleration Information

⌄ Configuration Settings

These settings can be changed by going to actions and selecting edit acceleration.
Learn More

```
allow_old_summaries = false
allow_skew = 0
backfill_time = -
cron_schedule = */5 * * * *
earliest_time = -1d
hunk.compression_codec = -
hunk.dfs_block_size = 0
hunk.file_format = -
manual_rebuilds = false
max_concurrent = 3
max_time = 3600
poll_buckets_until_maxtime = false
schedule_priority = default
workload_pool = -
```

Splunk Enterprise

Splunk Cloud

Expansive Data: Metrics

Metrics Optimizations

More Metrics.
Fewer Resources.

- **Multiple measures per metric event** yield faster processing and reduced storage
- **Histograms** provide **rich insight** into complex distributions
- New **all_nums** command in Logs-to-Metrics **automatically converts numbers** to metrics
- **Multiple aggregation functions** enable broader use of Rollups (introduced in 7.3)

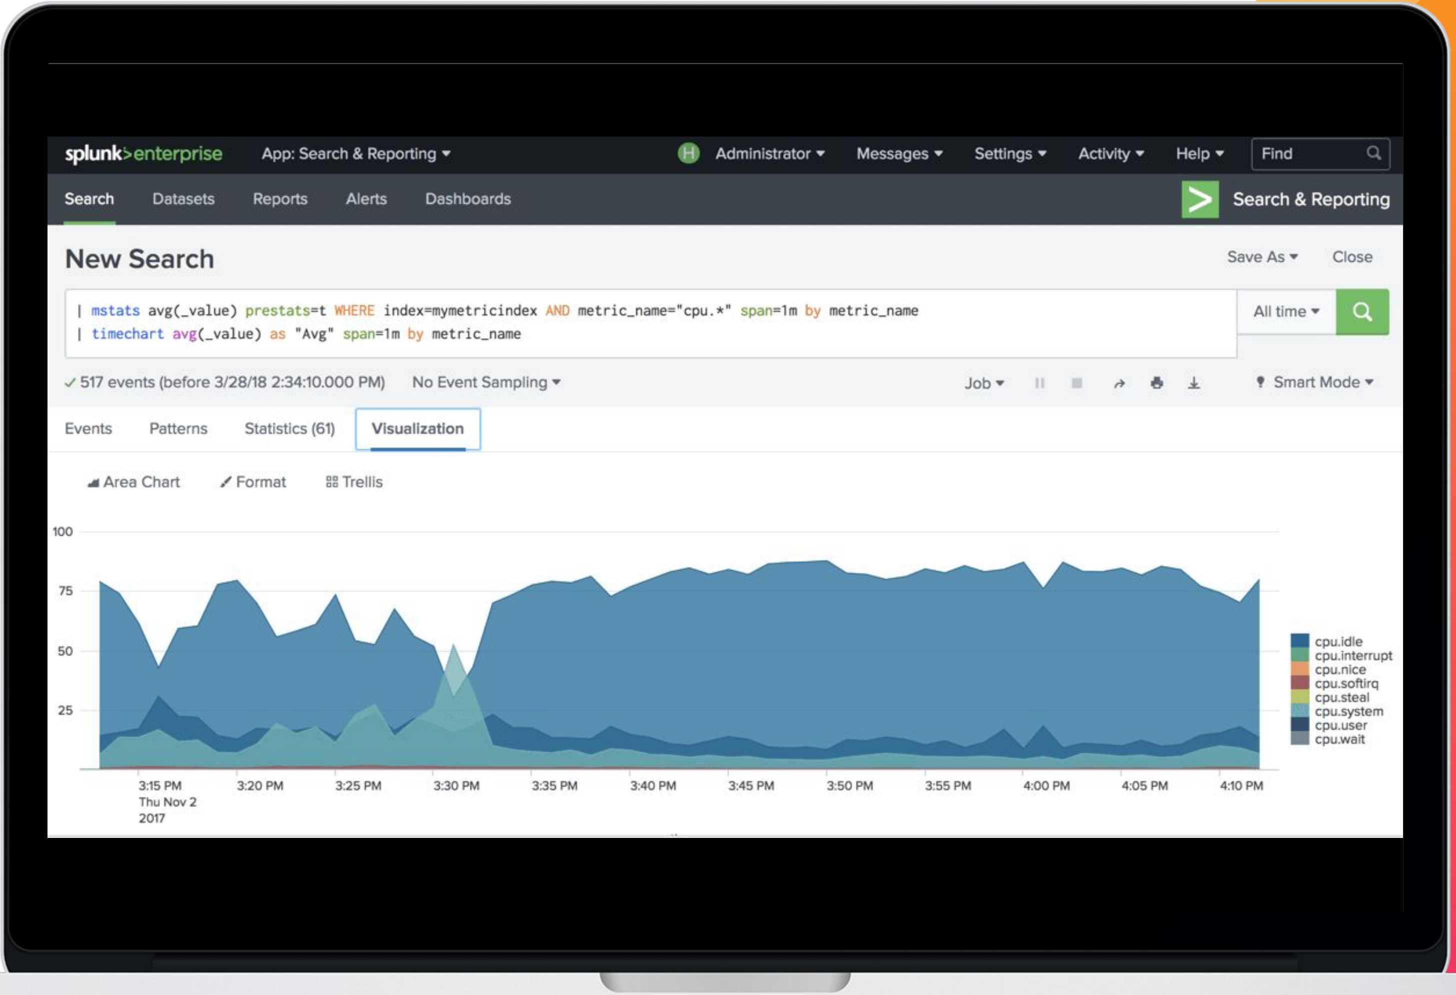

```
| mstats avg(_value) prestats=t WHERE index=mymetricindex AND metric_name="cpu.*" span=1m by metric_name
| timechart avg(_value) as "Avg" span=1m by metric_name
```

✓ 517 events (before 3/28/18 2:34:10.000 PM) No Event Sampling ▾

Splunk Enterprise

Splunk Cloud

What’s New: .conf19

Splunk Enterprise 8.0 and Splunk Cloud

Richer Insights for More Users

- Splunk Dashboards
- Analytics Workspace
- Field Access Controls
- Python 3.7

Expansive Data Access

- Search Performance
- Data Model Summaries
- Metrics Optimization

Splunk Cloud for More Customers

- **Global Growth**
- **FedRAMP**

Operability At Scale

Splunk Cloud: By The Numbers

10 Global Regions

500 TB/Day Ingest

80% YoY Revenue Growth*

5 Compliance Standards

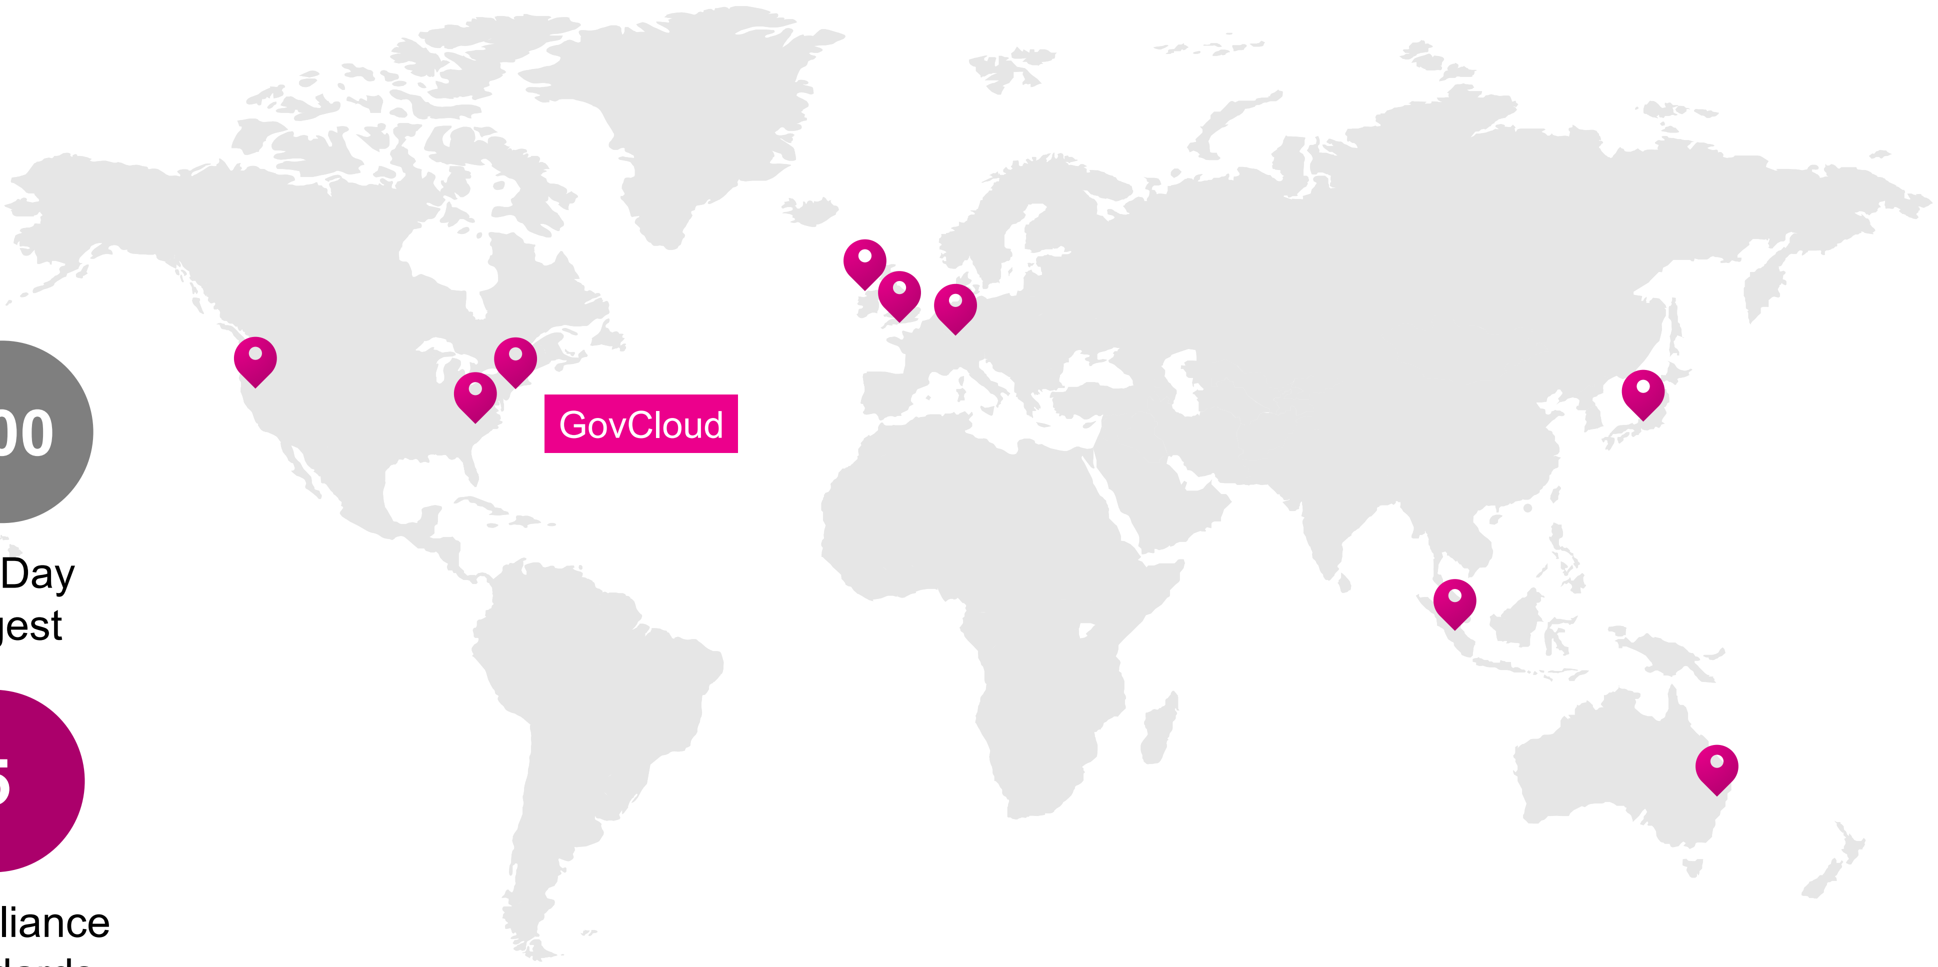

Splunk Cloud FedRAMP

Empowering agencies to drive decisions and actions at mission speeds

- FedRAMP Authorized by GSA for Moderate Impact Level SaaS
- Accelerates agency ATOs (Authority To Operate)
- Proactive risk management from the start
- Quick, real-time risk status of risk profiles
- Monitoring of any environment – cloud, on-premises or hybrid

What’s New: .conf19

Splunk Enterprise 8.0 and Splunk Cloud

Richer Insights for More Users

- Splunk Dashboards
- Analytics Workspace
- Field Access Controls
- Python 3.7

Expansive Data Access

- Search Performance
- Data Model Summaries
- Metrics Optimization

Splunk Cloud for More Customers

- Global Growth
- FedRAMP

Operability At Scale

- **Monitoring & Diagnostics**
- **Workload Management**
- **SmartStore**
- **Kubernetes Operator**

Operability: Monitoring and

Monitoring Console Gets Real (Time)

Visibility Drives Reliability

- Provides **holistic visibility** of Splunk Enterprise deployment and underlying infrastructure
- **Monitor and identify anomalies in real time** without running searches
- **Customize metrics** for richer, deployment-specific context

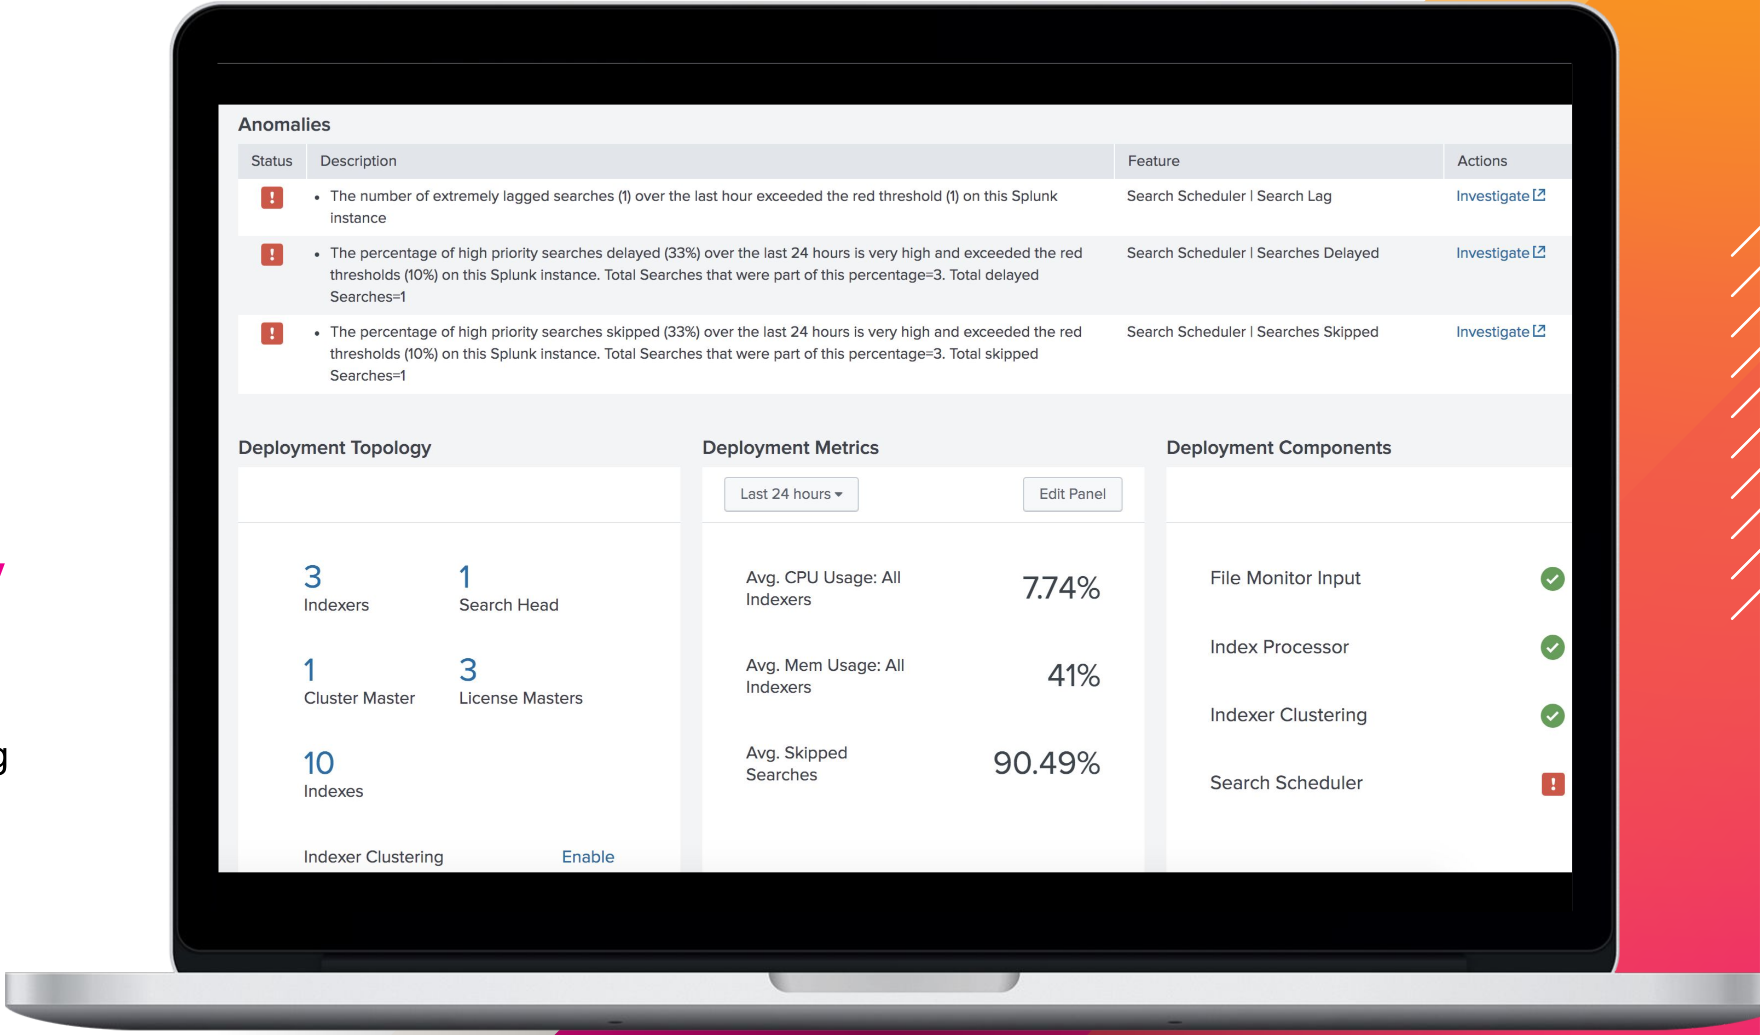

Anomalies

Status	Description	Feature	Actions
!	• The number of extremely lagged searches (1) over the last hour exceeded the red threshold (1) on this Splunk instance	Search Scheduler | Search Lag	Investigate
!	• The percentage of high priority searches delayed (33%) over the last 24 hours is very high and exceeded the red thresholds (10%) on this Splunk instance. Total Searches that were part of this percentage=3. Total delayed Searches=1	Search Scheduler | Searches Delayed	Investigate
!	• The percentage of high priority searches skipped (33%) over the last 24 hours is very high and exceeded the red thresholds (10%) on this Splunk instance. Total Searches that were part of this percentage=3. Total skipped Searches=1	Search Scheduler | Searches Skipped	Investigate

Deployment Topology

3 Indexers

1 Search Head

1 Cluster Master

3 License Masters

10 Indexes

Indexer Clustering Enable

Deployment Metrics

Avg. CPU Usage: All Indexers	7.74%
Avg. Mem Usage: All Indexers	41%
Avg. Skipped Searches	90.49%

Deployment Components

File Monitor Input	✓
Index Processor	✓
Indexer Clustering	✓
Search Scheduler	!

Splunk Enterprise

Splunk Cloud

splunk> .conf19

Monitoring Console Gets Proactive

Finds Your Faults. Just Like Mom.

- Health Reports dig deep into wellbeing of Indexer Cluster, Scheduler, and other components
- Real-time, deployment-wide health of most critical Splunk Indicators
- Health Checks integrated into Monitoring Console
- Downloadable, Splunk Support-curated content drives proactive problem identification

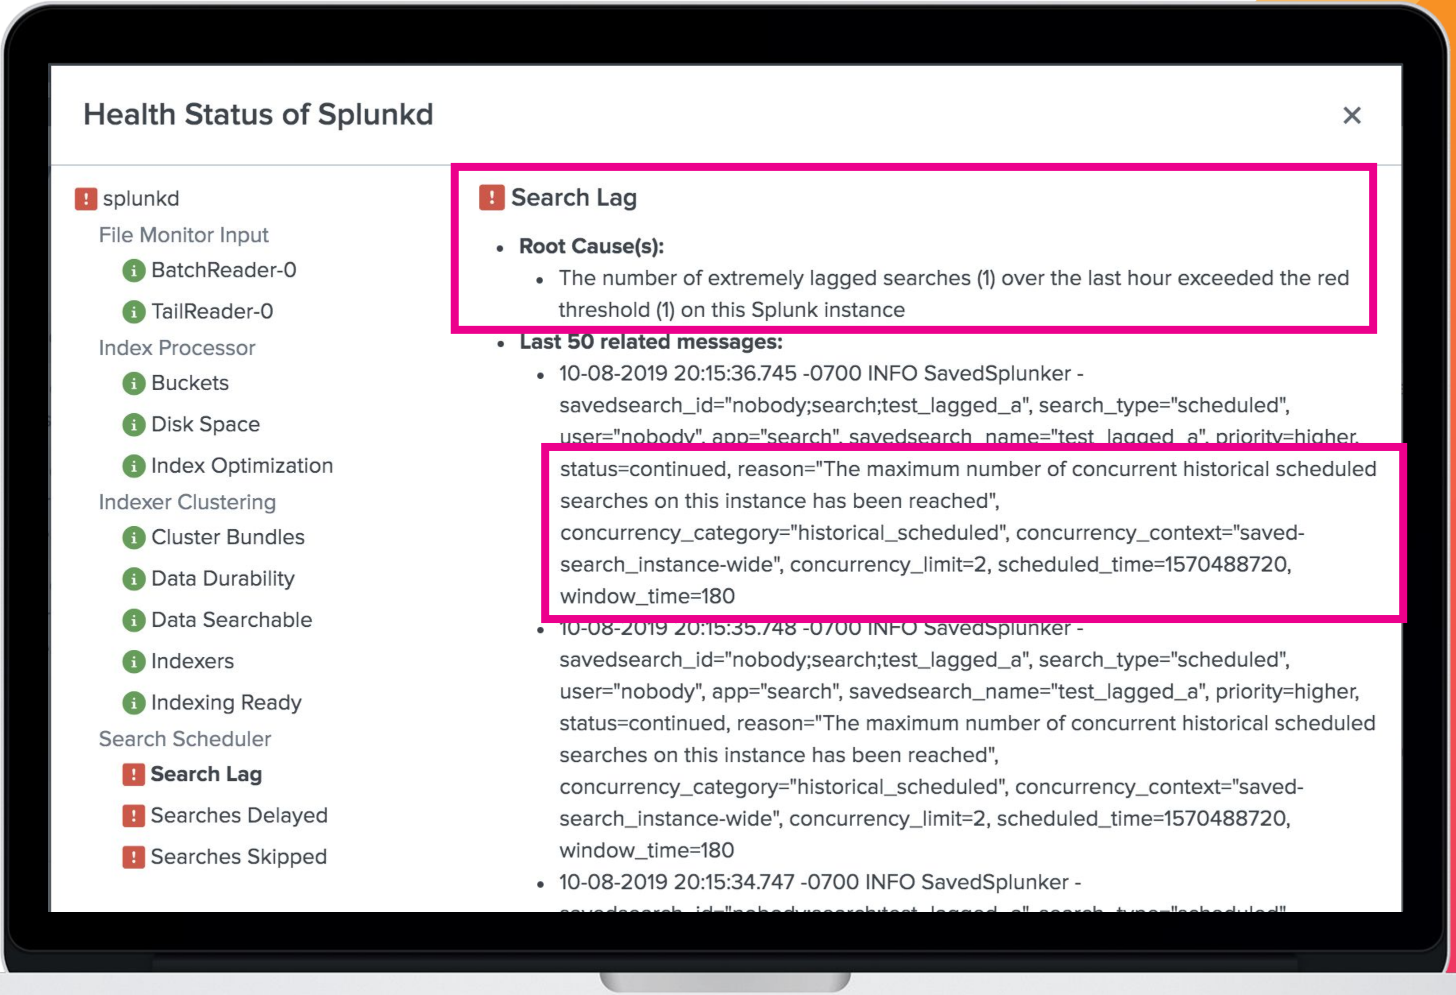

Health Status of Splunkd ×

! Search Lag

- **Root Cause(s):**
  - The number of extremely lagged searches (1) over the last hour exceeded the red threshold (1) on this Splunk instance
- **Last 50 related messages:**
  - 10-08-2019 20:15:36.745 -0700 INFO SavedSplunker - savedsearch_id="nobody;search;test_lagged_a", search_type="scheduled", user="nobody", app="search", savedsearch_name="test_lagged_a", priority=higher, status=continued, reason="The maximum number of concurrent historical scheduled searches on this instance has been reached", concurrency_category="historical_scheduled", concurrency_context="saved-search_instance-wide", concurrency_limit=2, scheduled_time=1570488720, window_time=180
  - 10-08-2019 20:15:35.748 -0700 INFO SavedSplunker - savedsearch_id="nobody;search;test_lagged_a", search_type="scheduled", user="nobody", app="search", savedsearch_name="test_lagged_a", priority=higher, status=continued, reason="The maximum number of concurrent historical scheduled searches on this instance has been reached", concurrency_category="historical_scheduled", concurrency_context="saved-search_instance-wide", concurrency_limit=2, scheduled_time=1570488720, window_time=180
  - 10-08-2019 20:15:34.747 -0700 INFO SavedSplunker - savedsearch_id="nobody;search;test_lagged_a", search_type="scheduled"

Splunk Enterprise

Splunk Cloud

Monitoring Console Goin' Mobile

rld's my home

- **Real-time alerting** for critical Splunk Health conditions
- Wherever and whenever
- Via **Splunk Mobile application** for iOS devices (Android in Beta)
- Or via web hooks to **any external application**

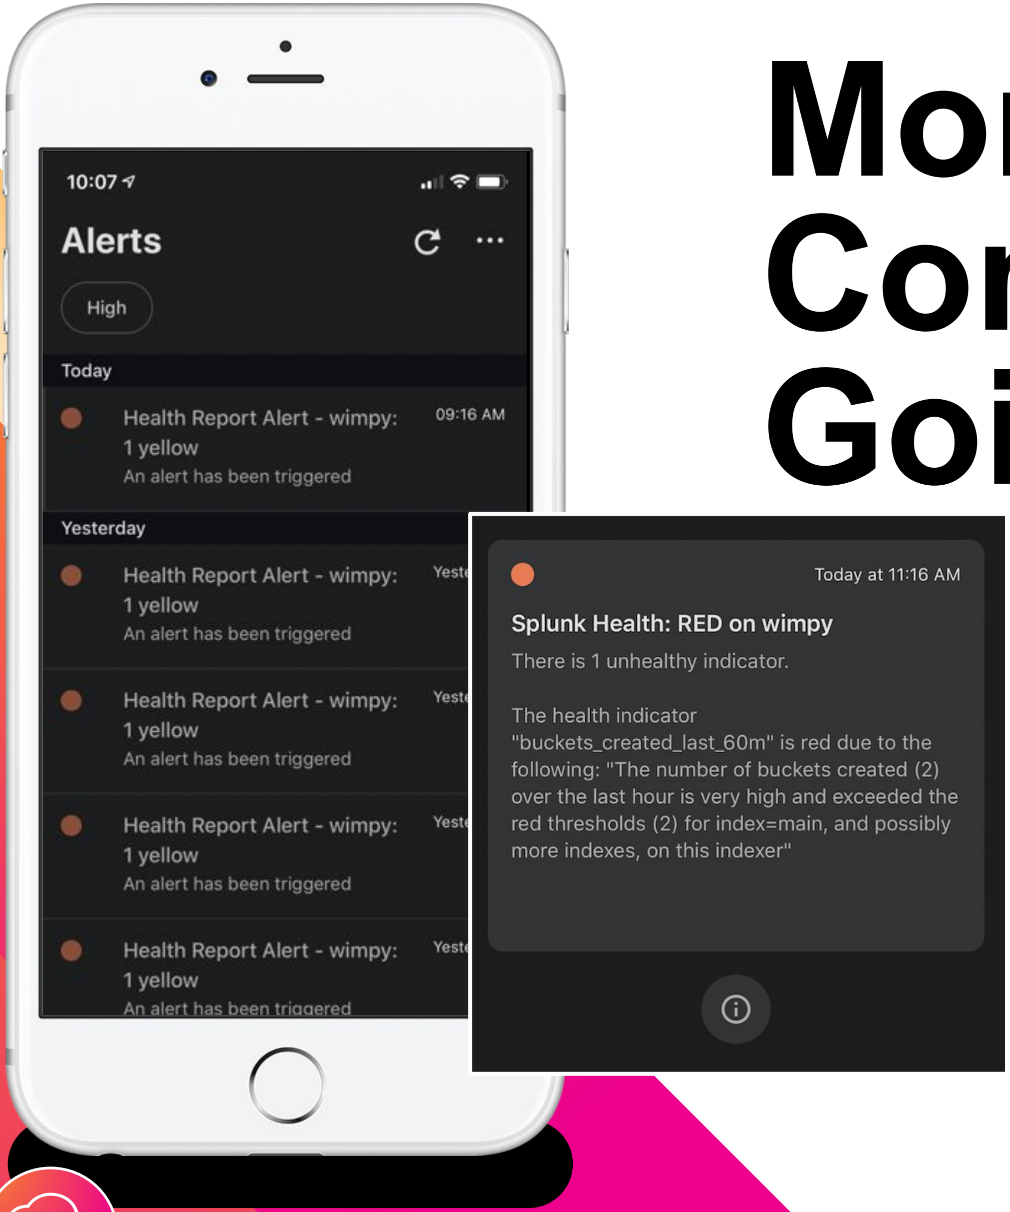

10:07

Alerts

High

Today

Health Report Alert - wimpy: 1 yellow 09:16 AM
An alert has been triggered

Yesterday

Health Report Alert - wimpy: 1 yellow Yeste
An alert has been triggered

Health Report Alert - wimpy: 1 yellow Yeste
An alert has been triggered

Health Report Alert - wimpy: 1 yellow Yeste
An alert has been triggered

Health Report Alert - wimpy: 1 yellow Yeste
An alert has been triggered

Today at 11:16 AM

Splunk Health: RED on wimpy

There is 1 unhealthy indicator.

The health indicator "buckets_created_last_60m" is red due to the following: "The number of buckets created (2) over the last hour is very high and exceeded the red thresholds (2) for index=main, and possibly more indexes, on this indexer"

Splunk Enterprise

Splunk Cloud

Cloud Monitoring Console

How The Other Half Lives

- **At-a-glance visibility** into Splunk Cloud service health, user, indexing and search activity, and other key metrics
- **Secure access** with role scoped to sc_admin
- **Prescriptive guidance** into storage utilization and search performance

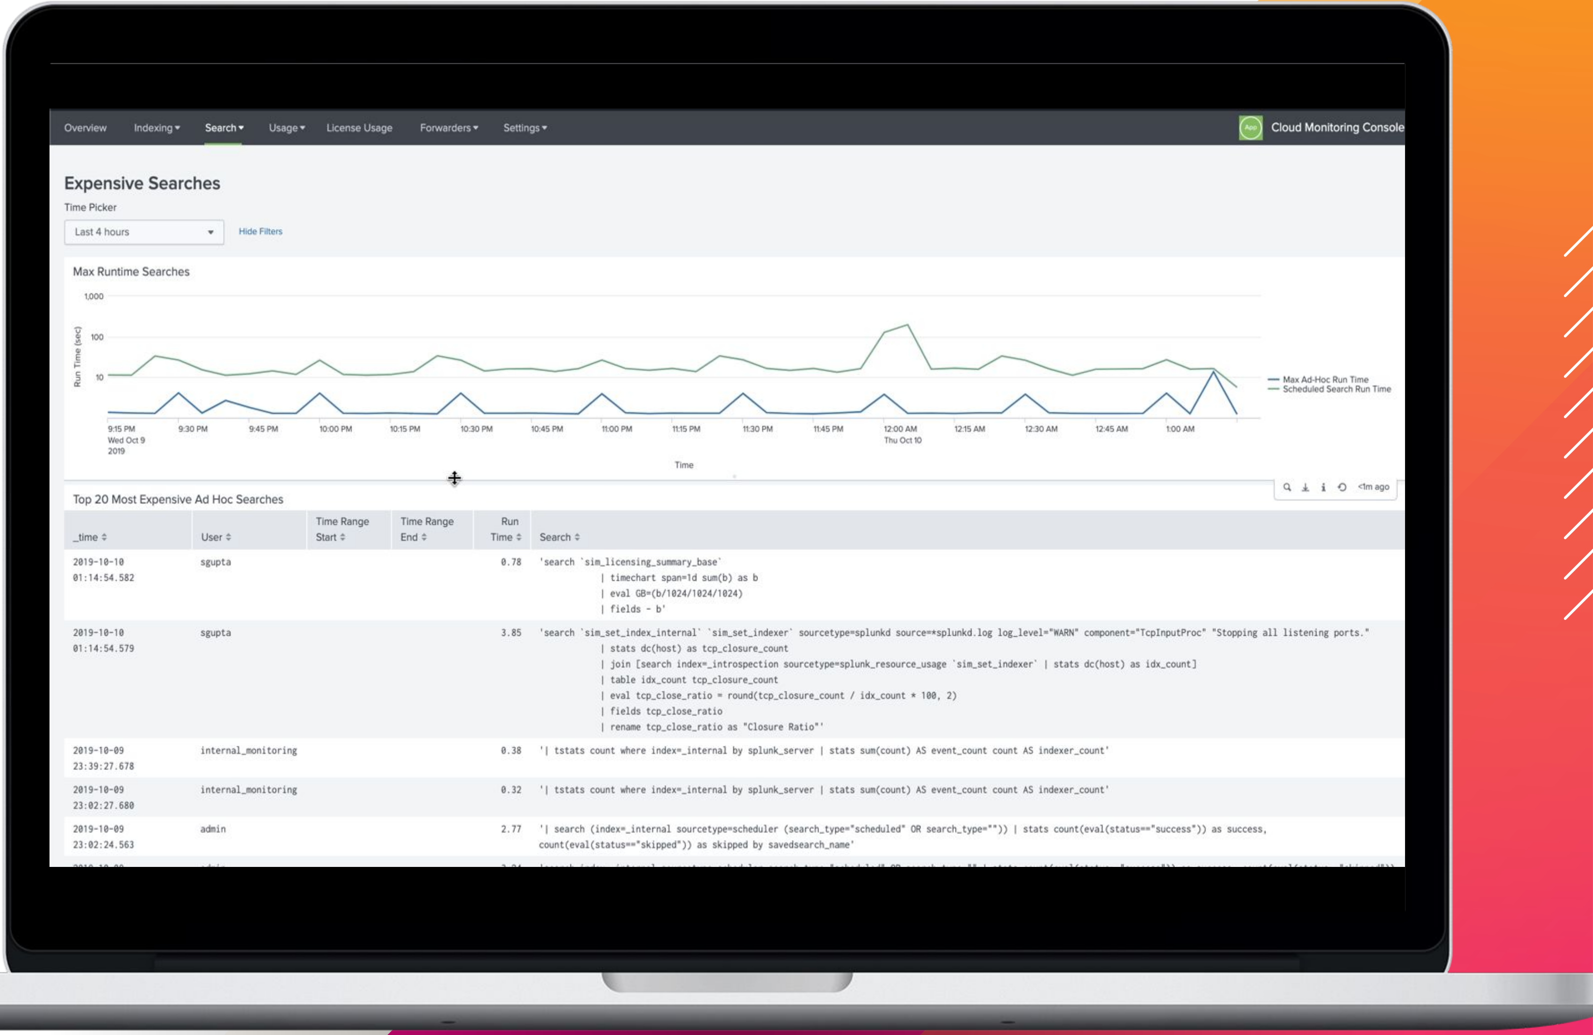

Expensive Searches

Top 20 Most Expensive Ad Hoc Searches

<table>
<tr><th>_time ⇕</th><th>User ⇕</th><th>Time Range Start ⇕</th><th>Time Range End ⇕</th><th>Run Time ⇕</th><th>Search ⇕</th></tr>
<tr><td>2019-10-10
01:14:54.582</td><td>sgupta</td><td></td><td></td><td>0.78</td><td>'search `sim_licensing_summary_base`
| timechart span=1d sum(b) as b
| eval GB=(b/1024/1024/1024)
| fields - b'</td></tr>
<tr><td>2019-10-10
01:14:54.579</td><td>sgupta</td><td></td><td></td><td>3.85</td><td>'search `sim_set_index_internal` `sim_set_indexer` sourcetype=splunkd source=*splunkd.log log_level="WARN" component="TcpInputProc" "Stopping all listening ports."
| stats dc(host) as tcp_closure_count
| join [search index=_introspection sourcetype=splunk_resource_usage `sim_set_indexer` | stats dc(host) as idx_count]
| table idx_count tcp_closure_count
| eval tcp_close_ratio = round(tcp_closure_count / idx_count * 100, 2)
| fields tcp_close_ratio
| rename tcp_close_ratio as "Closure Ratio"'</td></tr>
<tr><td>2019-10-09
23:39:27.678</td><td>internal_monitoring</td><td></td><td></td><td>0.38</td><td>'| tstats count where index=_internal by splunk_server | stats sum(count) AS event_count count AS indexer_count'</td></tr>
<tr><td>2019-10-09
23:02:27.680</td><td>internal_monitoring</td><td></td><td></td><td>0.32</td><td>'| tstats count where index=_internal by splunk_server | stats sum(count) AS event_count count AS indexer_count'</td></tr>
<tr><td>2019-10-09
23:02:24.563</td><td>admin</td><td></td><td></td><td>2.77</td><td>'| search (index=_internal sourcetype=scheduler (search_type="scheduled" OR search_type="")) | stats count(eval(status=="success")) as success, count(eval(status=="skipped")) as skipped by savedsearch_name'</td></tr>
</table>

Splunk Enterprise

Splunk Cloud

Workload Management

Conflicts. Avoided. Resolved.

- **Introduced in Enterprise 7.2** to **avoid contention** between indexing and search activity and **ensure predictable operation**
- **Automated remediation** of rogue or runaway searches
- **Enhanced rules framework** streamlines configuration
- **Schedule-based rules** manage workloads across peak and off-peak periods

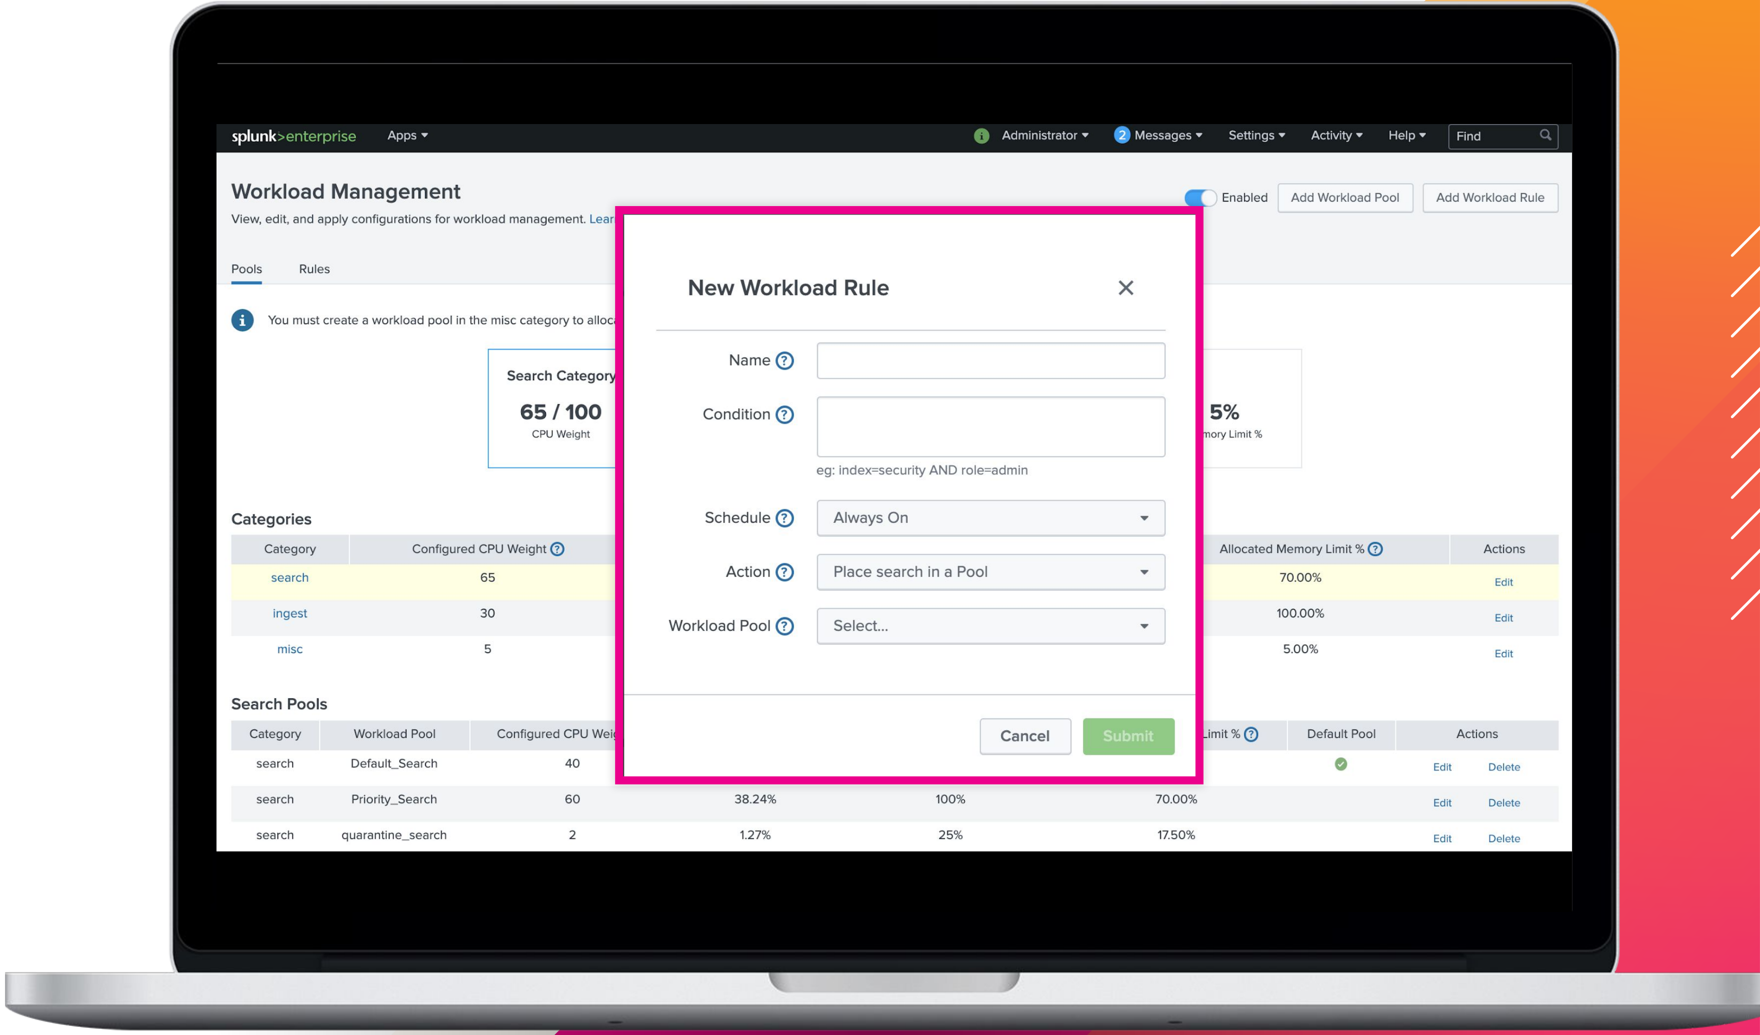

Workload Management

View, edit, and apply configurations for workload management. Lear

You must create a workload pool in the misc category to alloc

Search Category

65 / 100

CPU Weight

5%

mory Limit %

Categories

Category	Configured CPU Weight	Allocated Memory Limit %	Actions
search	65	70.00%	Edit
ingest	30	100.00%	Edit
misc	5	5.00%	Edit

Search Pools

Category	Workload Pool	Configured CPU Wei			Limit %	Default Pool	Actions
search	Default_Search	40				✓	Edit Delete
search	Priority_Search	60	38.24%	100%	70.00%		Edit Delete
search	quarantine_search	2	1.27%	25%	17.50%		Edit Delete

New Workload Rule ✕

Name ?

Condition ?

eg: index=security AND role=admin

Schedule ? Always On ▾

Action ? Place search in a Pool ▾

Workload Pool ? Select... ▾

Cancel

Submit

Splunk Enterprise

Splunk Cloud

Splunk Cloud Workload Management

Easy Is Good. Really Good.

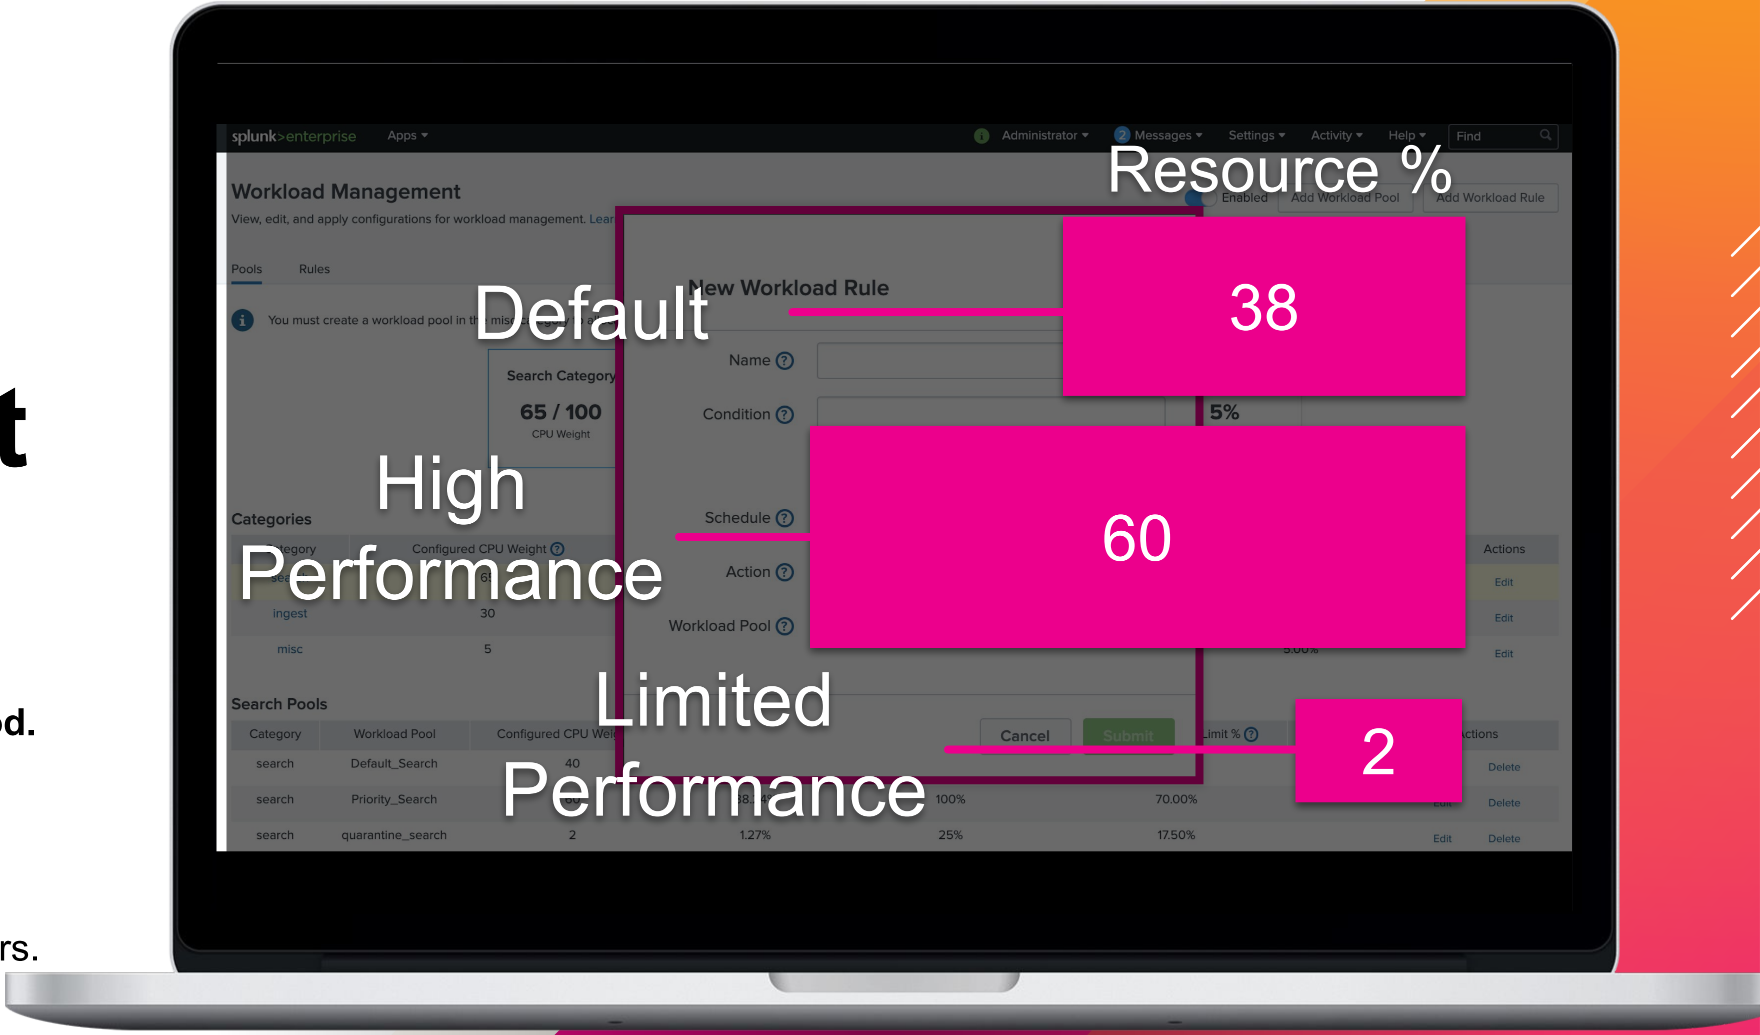

Workload Management

View, edit, and apply configurations for workload management. Lear

You must create a workload pool in th misc category to al

New Workload Rule

5%

Categories

Category	Configured CPU Weight
search	65
ingest	30
misc	5

Search Pools

Category	Workload Pool	Configured CPU Weight
search	Default_Search	40
search	Priority_Search	60
search	quarantine_search	2

Ingest resources are protected. Period.

Three search resource pools:

- **Default.** Everything starts here.
- **High Performance**. For priority workloads and otherwise-favored users.
- **Limited Performance.** AKA The Island of Misfit Toys.

Splunk Enterprise

Splunk Cloud

Operability: Bundle Management

Bundle Management

Like A Ferrari Pickup Truck

New Cascading Bundle Replication is up to **3x faster** and **scales to hundreds of indexers**

- **High-throughput pipeline** from search heads to peers
- Drive **consistent and efficient search execution** using most current knowledge objects

Improved **Deployer Bundle Push** is up to **10x faster**

- Streamlines configuration pushes to search heads

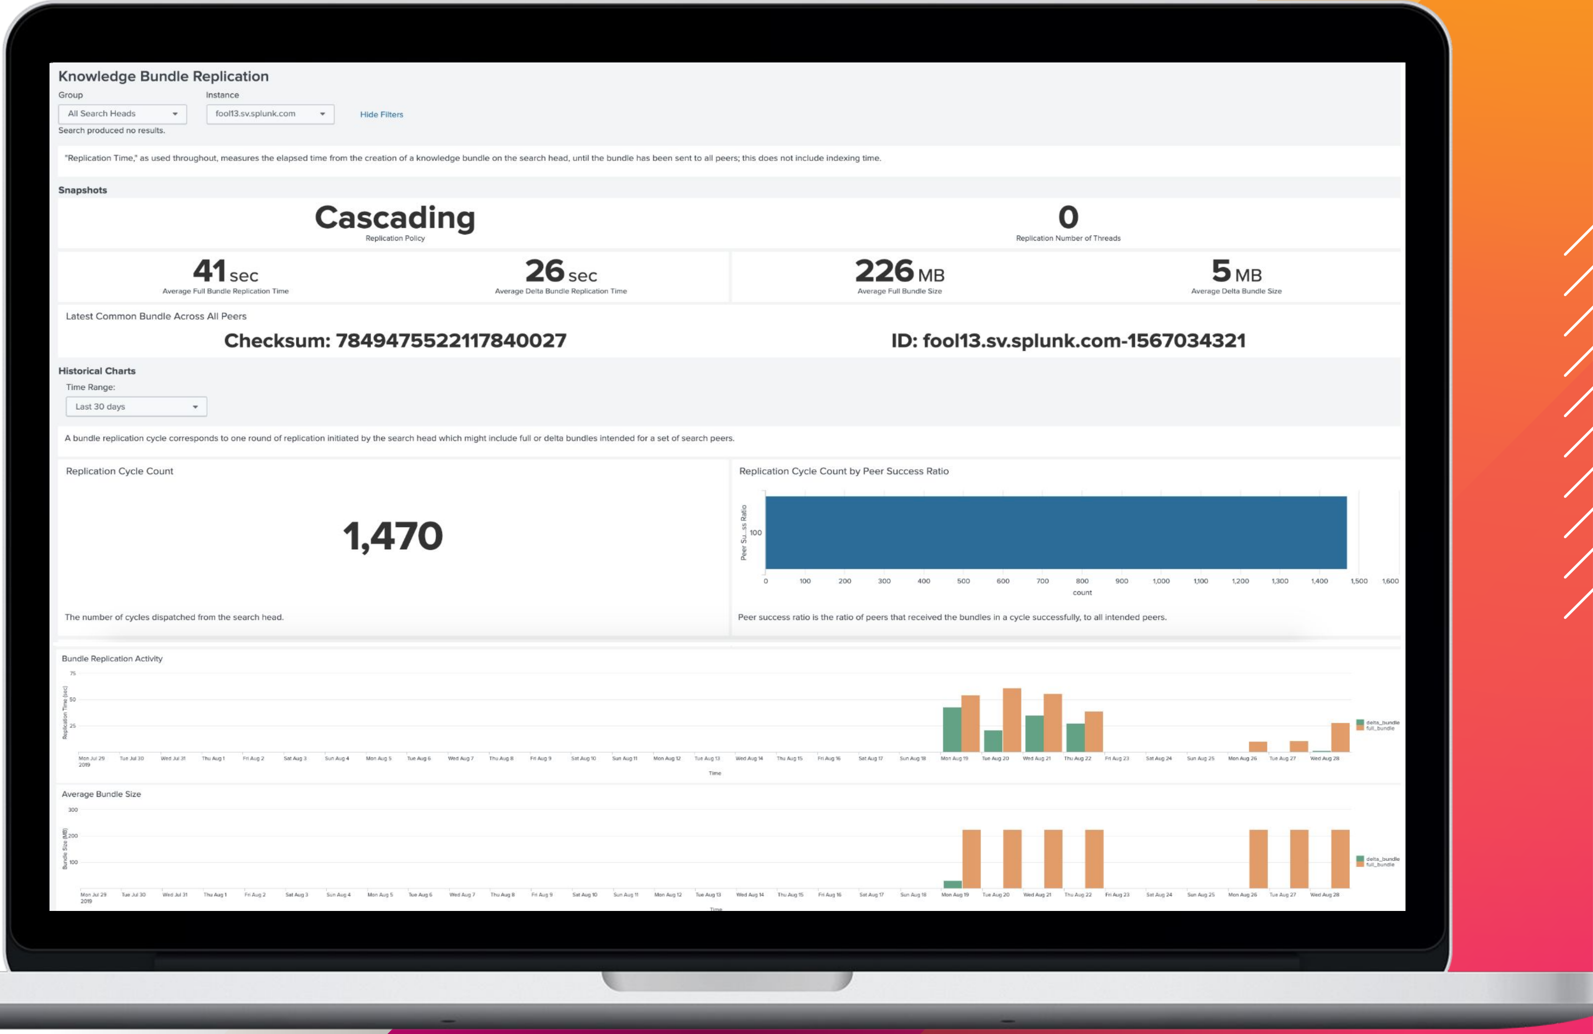

Knowledge Bundle Replication

Group: All Search Heads

Instance: fool13.sv.splunk.com

Hide Filters

Search produced no results.

"Replication Time," as used throughout, measures the elapsed time from the creation of a knowledge bundle on the search head, until the bundle has been sent to all peers; this does not include indexing time.

Snapshots

Cascading Replication Policy	**0** Replication Number of Threads

41 sec Average Full Bundle Replication Time	**26** sec Average Delta Bundle Replication Time	**226** MB Average Full Bundle Size	**5** MB Average Delta Bundle Size

Latest Common Bundle Across All Peers

Checksum: 7849475522117840027

ID: fool13.sv.splunk.com-1567034321

Historical Charts

Time Range: Last 30 days

A bundle replication cycle corresponds to one round of replication initiated by the search head which might include full or delta bundles intended for a set of search peers.

Replication Cycle Count

1,470

The number of cycles dispatched from the search head.

Replication Cycle Count by Peer Success Ratio

Peer success ratio is the ratio of peers that received the bundles in a cycle successfully, to all intended peers.

splunk> .conf19

SmartStore

Fast. Cheap. Robust. Pick Any Three.

- **Splunk Cloud hardened,** then **introduced** in Enterprise 7.2 and **enhanced** in 7.3
- Uses a **separate tier of object storage** for warm and cold buckets
- Indexer **direct-attached storage** used for ingest hot buckets and search results
- **Independently scale** compute and storage
- **Lower TCO** and **improved operability**
- Enterprise 8.0 adds **improved operability**

Object Storage

Splunk Enterprise

Splunk Cloud

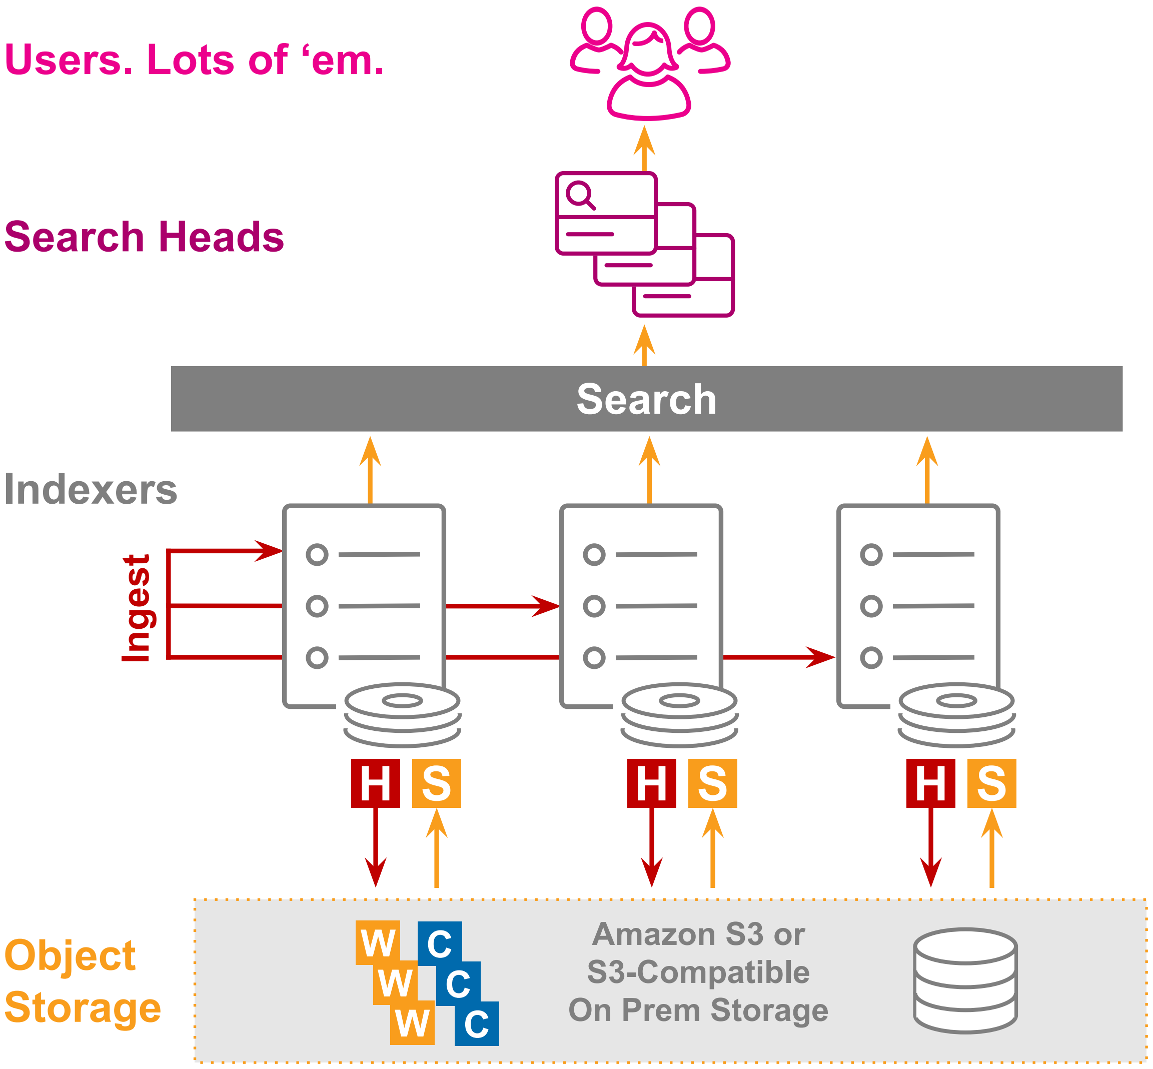

Splunk Cloud Data Retention Options

This One Goes to 11

Dynamic Data: Active Archive

- Move less-frequently accessed data to cost-effective, Splunk-managed data archive
- Easily restore data to Splunk Cloud

Dynamic Data: Self-Storage

- Move data from Splunk Cloud to your own Amazon S3 environment
- Data no longer accessible via k Cloud

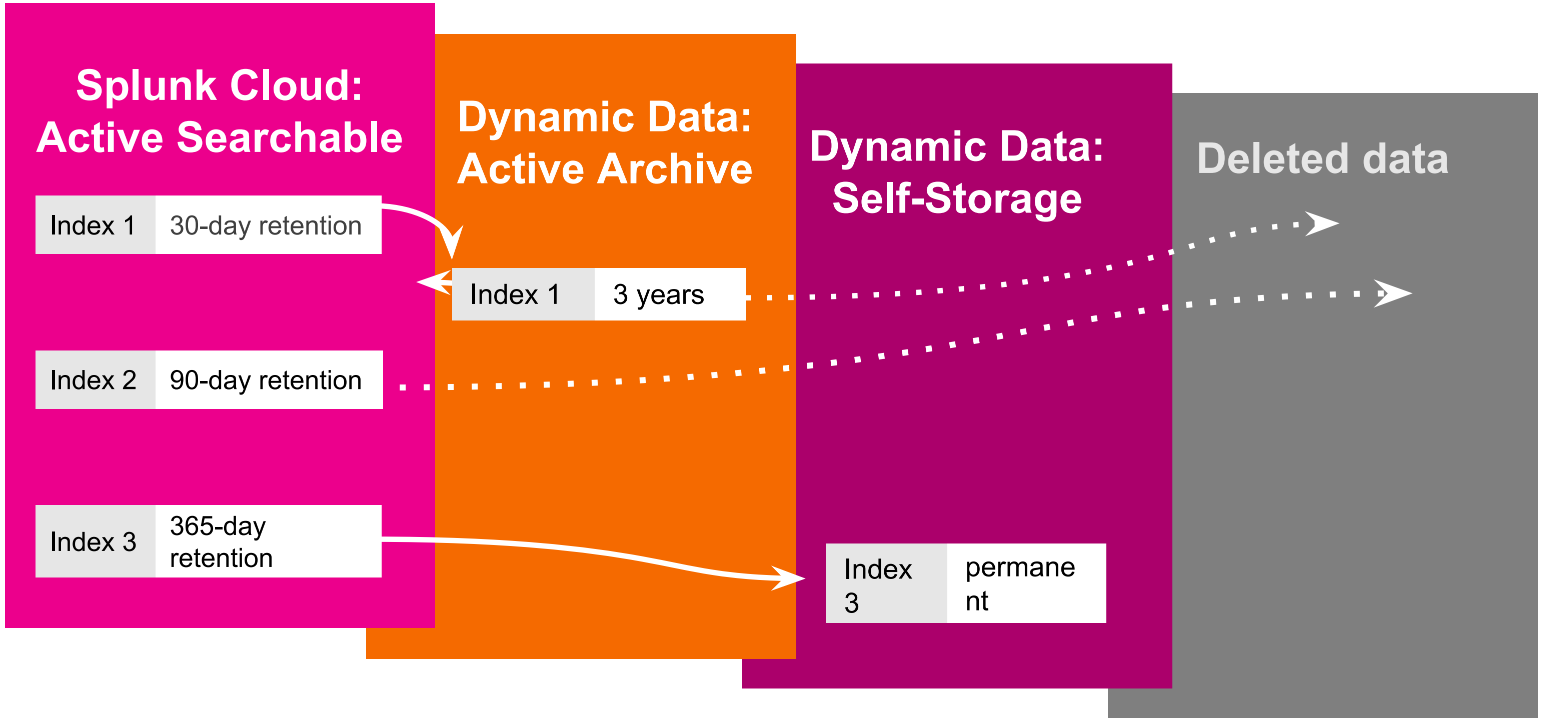

Splunk Enterprise Kubernetes Operator

Deployment Time, Meet Horror Movie

```
apiVersion:
enterprise.splunk.com/v1alpha1
kind: SplunkEnterprise
metadata:
 name: withdfs
 finalizers:
   -
enterprise.splunk.com/delete-
spec:
 enableDFS: true
 licenseUrl: https://my-license
topology:
 indexers: 3
 searchHeads: 3
```

- New (alpha) Operator **streamlines deployment and operation** of Splunk Enterprise using Docker Container
- **Automatically configures** new deployments according to best practices, including: app installation, Search Head and Indexer clustering
- **Automates upgrades**, scaling, rolling restarts, recovery from crashes and hardware failures
- Tested on **Amazon** EKS, **Google** GKE, **Red Hat** OpenShift, **Docker Enterprise, Open Source** Kubernetes
- Single-command deployment: **kubectl apply -f https://tiny.cc/splunk-operator-install**
- Open Source release - look forward to sharing the journey

Other Cool Stuff

Expanding The Splunk Platform

Expansive Data Access | **Richer Insights...** | **... For More Users**

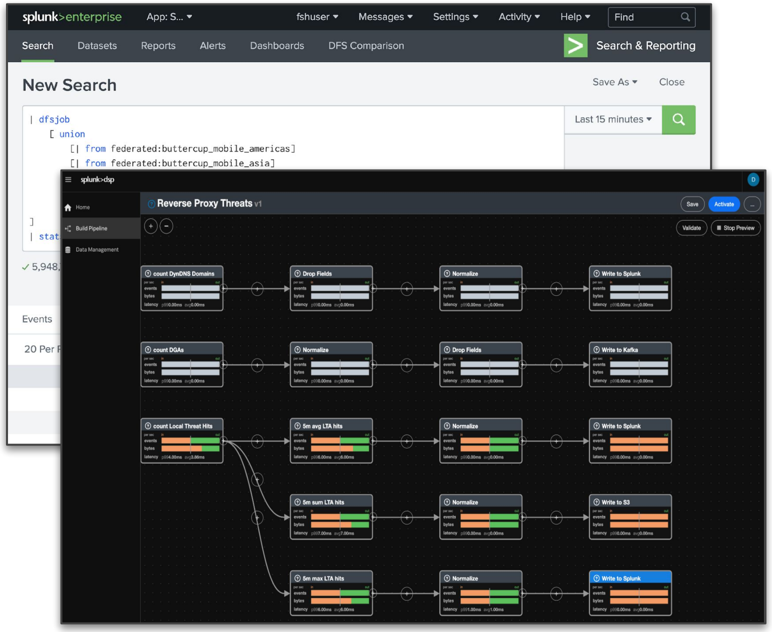

```
| dfsjob
    [ union
        [| from federated:buttercup_mobile_americas]
        [| from federated:buttercup_mobile_asia]
```

Data Fabric Search
Data Stream
Processor

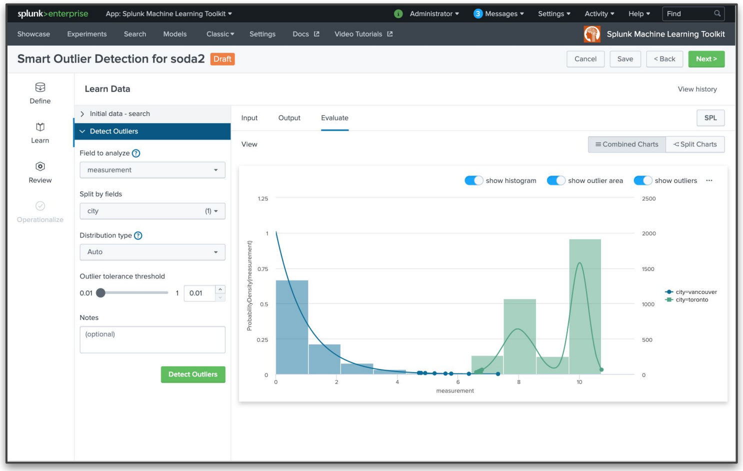

Machine Learning
Toolkit

12:03
Alerts
Critical
High
Med
Today
Storage Free Space: 5%
The system mysql-02 has less than 10% free disk space.
Account Deleted
Host With Multiple Infections (10.11.36.9)
Personal Identifiable Information Detected
Yesterday
Insecure Or Cleartext Authentication Detected
Alerts
Dashboards
Search
Backend Health
10:09
Latency by Request
41
Alert List Request
Circuit + LTE + Door Status Metrics
Master
Left1
Right1
Left2
Right2
Left3

Splunk Mobile - AR - Watch

Next Stop: More Information

And/Or Give Us Feedback….

source=*Pavilion | Apps Showcase | Foundations & Platforms

21	Workload Management & Monitoring
22	@Scale Arch, SmartStore and Clustering
23	Splunk Data Ingestion
24	Search
25	Splunk Stream
26	Enterprise Dashboards & Visualizations
27	Metrics & Analytics Workspace
30	Splunk Cloud

.conf19

splunk>

Thank You!

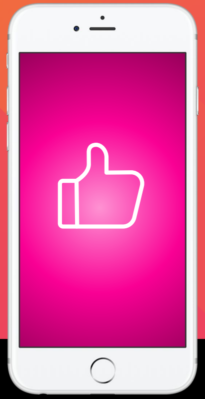

Go to the .conf19 mobile app to

RATE THIS SESSION