Infrastructure and System Monitoring with Splunk and Telegraf

Lance O’Connor – TiVo
Nick Tankersley – Splunk

Lance O'Connor

Distinguished Engineer | TiVo, Inc.

Nick Tankersley

Principal Product Manager | Splunk, Inc.

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.

- 11 of the last 13 years at TiVo
- Focused on Splunk, metrics, monitoring and data science for the last 3 years
- Previously worked in everything from corporate IT and security, production operations, to engineering teams
- I’ve traveled all over the US and Europe doing large scale systems integrations and building out data centers
- Telegraf (1.12+) now includes a Fireboard input so I can track temperatures from my smoker… in Splunk of course

ronnocol

https://keybase.io/ronnocol

- 4 years at Splunk
- Product Manager for Splunk ITSI, Splunk App for Infrastructure
- Last 2 years focused on metrics based monitoring for Splunk
- Previously a developer who designed and built data as a service solutions for an evolving media company
- I will eat the food Lance makes with his smoker

IMDb imdb.com/name/nm2685199/

Metrics > Events

Metrics indexes are more performant than events indexes for time-series based data

Why the move to Metrics?

- Event based monitoring is slow and complex
  - Searching raw events for metric data is often slow and requires complex SPL and or transforms/props.
- Metrics indexes are significantly faster and smaller than event based indexes
- We got a new set of teammates that used Telegraf and Scalyr that were used to:
  - Time-Series indexes (optimized for metrics)
  - Point and Click interfaces for graphs (no need to learn SPL)
  - Rapid indications of changes of status (short sampling periods)
- We just couldn’t support the requirements and expectations of our new teammates without moving to the metrics index store.

Telegraf

Collecting, processing, aggregating, and writing metrics.

Telegraf by the Numbers

> 160

Inputs

> 20

Parsers, Processors, Aggregators, and Serializers

> 30

Outputs, 11 that support serializers

> 25

Deployment packages built nightly for 13 different platforms

The Telegraf Architecture

Telegraf has a modular, plugin based architecture

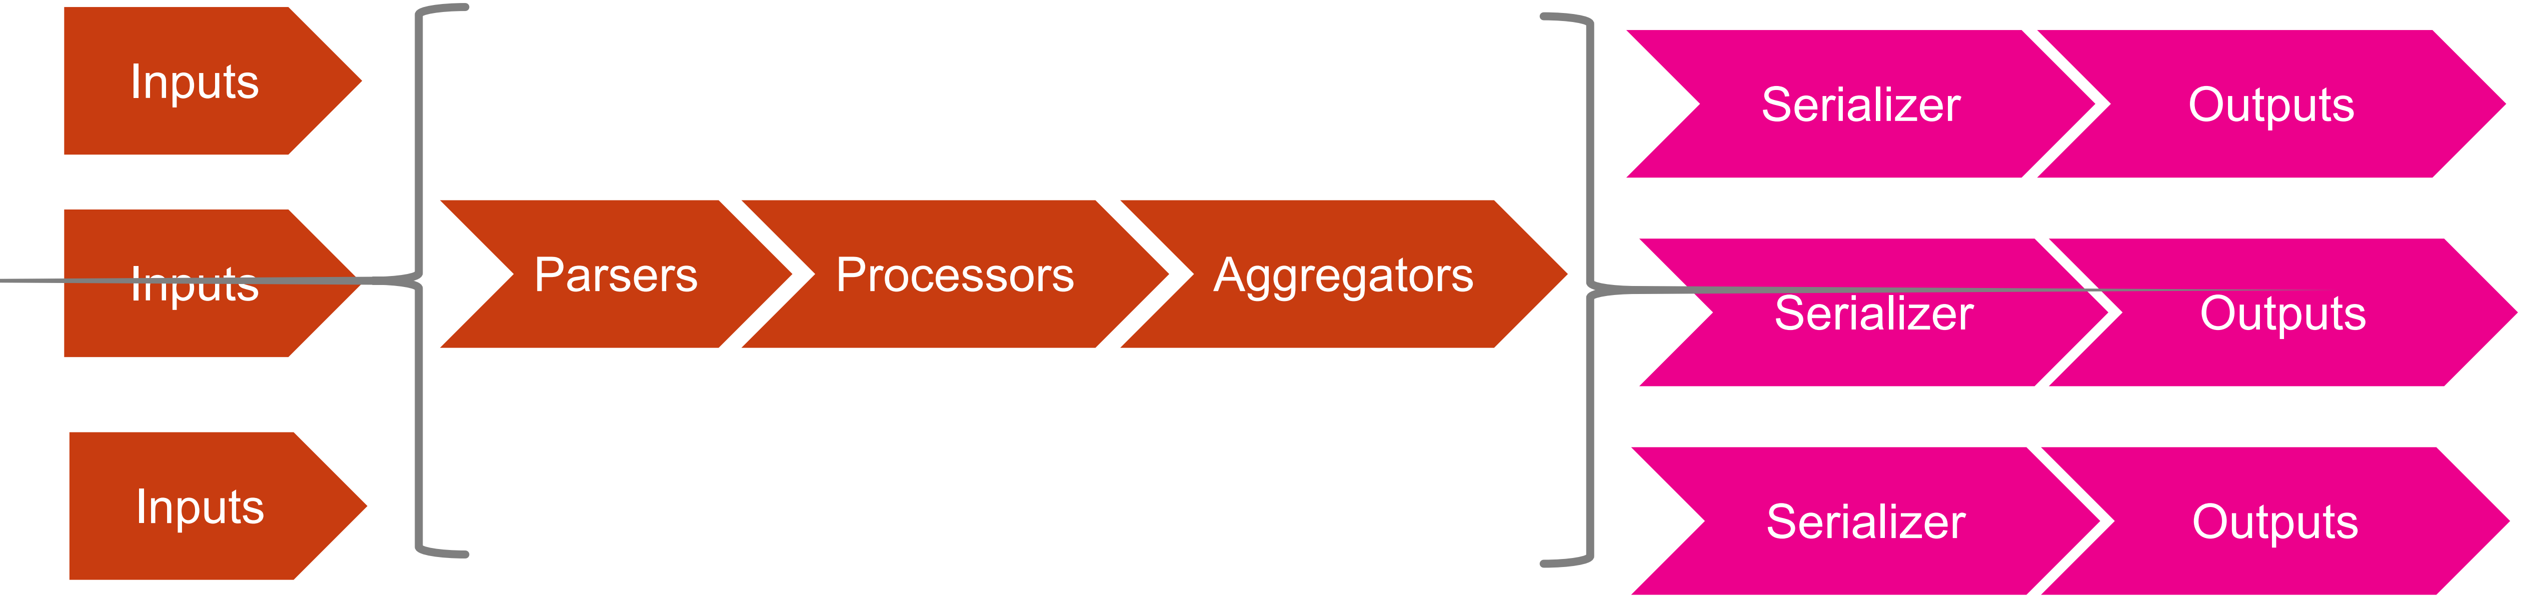

Connecting Telegraf to Splunk

- Originally I created a new output plugin
  - Knew how to send data to a Splunk HEC
  - Had data formatting and HTTP client code
  - Duplicated much of the existing HTTP output functionality

Outputs

Serializer

- Threw all of the code away and created a serializer instead
  - Only had to manage taking the metrics from the internal format and formatting it into a Splunk compatible fashion (e.g. expectations of JSON structure and field names… like _value)
  - Compatible with any output plugin that can use a data_format serializer
  - Significantly increased our ability to integrate Telegraf with our Splunk deployments

Telegraf Takeaways

1. Metrics based indexes are better for time-series based data
2. Telegraf has a small footprint that allows it to be installed on almost any platform of any size, from a raspberry pi to a multi-core server
3. Telegraf has input plugins for virtually any data source (with more being added every day)
4. Telegraf talks to Splunk natively

Deployablility

Why using a serializer was better

Standalone Deployment

Deploy Telegraf on its own

- This method doesn’t require any additional Splunk components to be installed
- Very small memory and processor resource requirements
- Talks directly to the HEC
- Allows for centralized management of metrics collectors from other tools (Ansible, Puppet, etc.) and decentralized from the Splunk deployment.

“Sidecar” Deployment

Deploy Telegraf alongside a Splunk forwarder

- Telegraf is installed alongside a universal or heavy forwarder
- Splunk is configured to read the file that Telegraf outputs
- Allows for Splunk admins to administer Splunk and System admins to administer Telegraf with very little overlap
- Writes out a file using
- Splunk has a monitor:// config to read the file that Telegraf generates

Splunk Application Deployment

Deploy Telegraf as a Splunk scripted input

- Telegraf is installed on a Universal or Heavy forwarder by a deployment server
- There can be different Telegraf configurations associated with the scripted input name
- Uses the Splunk forwarder’s already configured outputs to ingest the data from Telegraf
- Scripted input controls Telegraf’s configuration file
- Splunk starts Telegraf and ensures it continues to run

Telegraf Deployment Options

1. There's a variety of ways to deploy Telegraf

2. Can be used to meet any number of IT deployment toolkits or requirements

3. Can manage Telegraf and Splunk independently

4. Can use Splunk as a deployment and configuration engine

Splunk Integrations

How to use Telegraf metrics in Splunk

splunk> .conf19

Splunk App for Infrastructure

Fast and Easy Infrastructure Monitoring with Metrics and Logs

Splunk App for Infrastructure

Comprehensive infrastructure monitoring, alerting and investigation for Servers, OS, AWS, Kubernetes, Docker, OpenShift, Vmware

- Fast time-to-value: simple setup and data collection
- Guided investigations to quickly identify trends and root cause
- Detailed and flexible analysis spanning metrics and logs
- Splunk Cloud support
- Integrations with ITSI

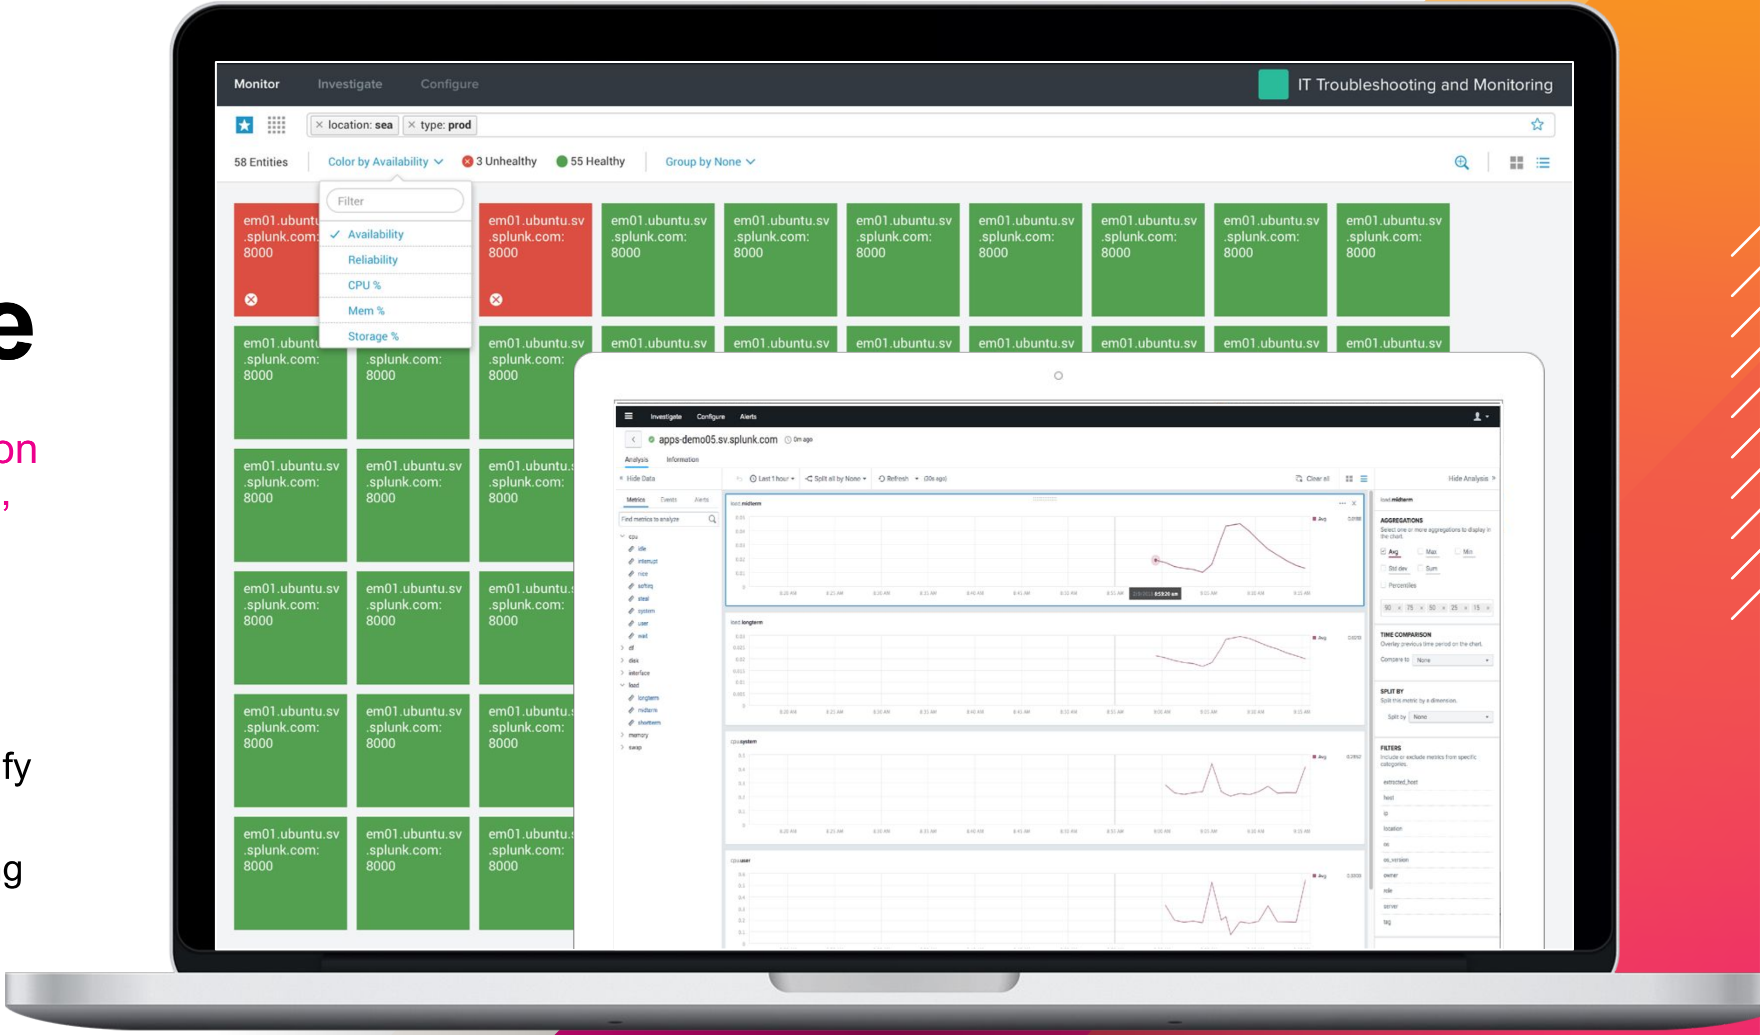

Prescriptive Data Collection

Onboard thousands of servers in the time it takes to install most other Enterprise monitoring tools

- Guided data onboarding for:
  - Linux/Unix/OSX Infrastructure
  - Windows Infrastructure
  - Amazon Web Services
  - Docker
  - Kubernetes
  - OpenShift
  - Vmware
- Easy install scripts deploys collection tools for metrics and logs in minutes
- Add custom metadata and tune data collection within the UI

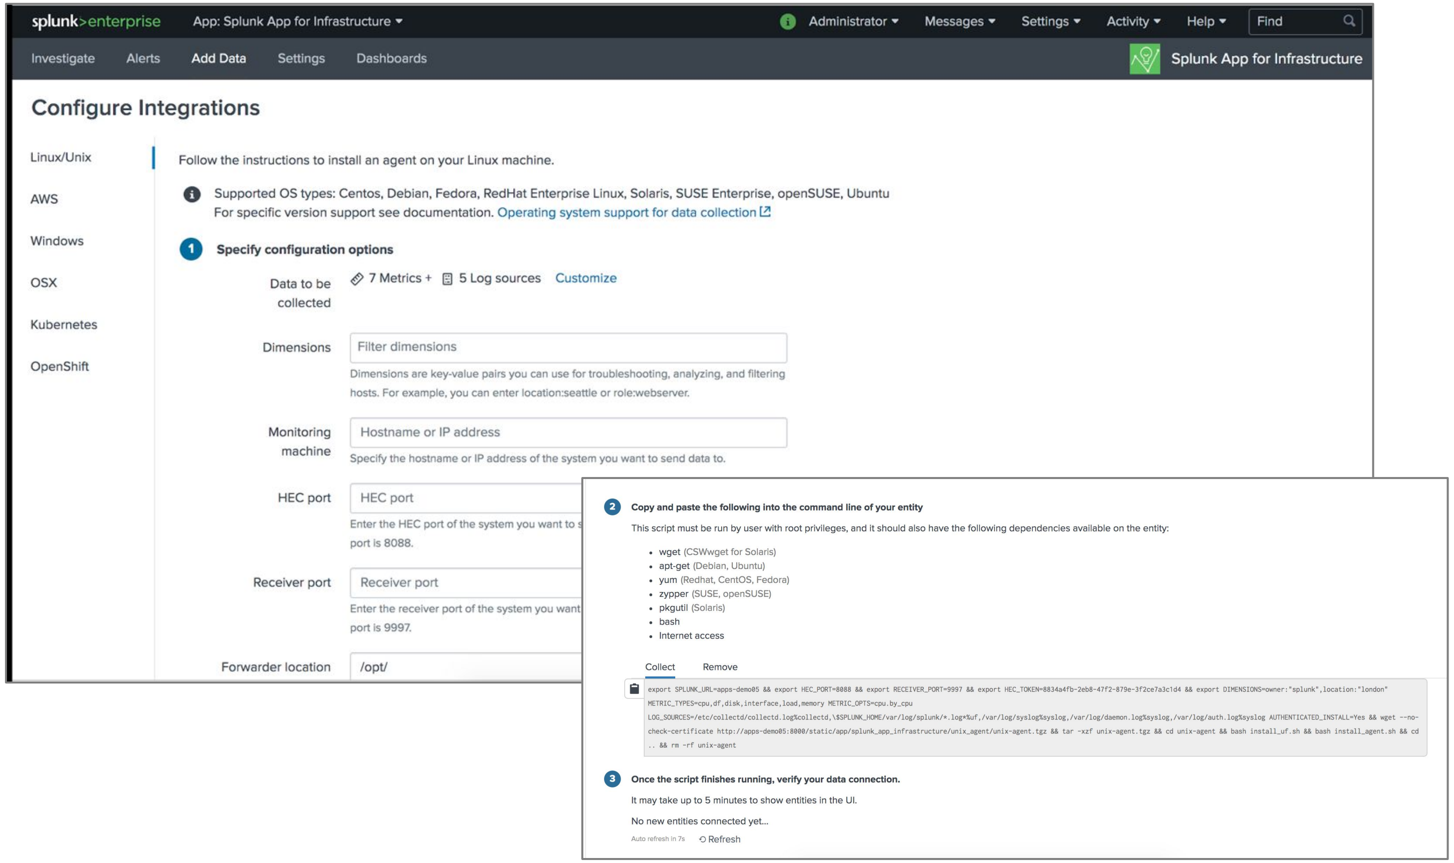

Configure Integrations

Follow the instructions to install an agent on your Linux machine.

Supported OS types: Centos, Debian, Fedora, RedHat Enterprise Linux, Solaris, SUSE Enterprise, openSUSE, Ubuntu
For specific version support see documentation. Operating system support for data collection

1 Specify configuration options

Data to be collected: 7 Metrics + 5 Log sources Customize

Dimensions: Filter dimensions
Dimensions are key-value pairs you can use for troubleshooting, analyzing, and filtering hosts. For example, you can enter location:seattle or role:webserver.

Monitoring machine: Hostname or IP address
Specify the hostname or IP address of the system you want to send data to.

HEC port: HEC port
Enter the HEC port of the system you want to s… port is 8088.

Receiver port: Receiver port
Enter the receiver port of the system you want… port is 9997.

Forwarder location: /opt/

2 Copy and paste the following into the command line of your entity

This script must be run by user with root privileges, and it should also have the following dependencies available on the entity:

- wget (CSWwget for Solaris)
- apt-get (Debian, Ubuntu)
- yum (Redhat, CentOS, Fedora)
- zypper (SUSE, openSUSE)
- pkgutil (Solaris)
- bash
- Internet access

```
export SPLUNK_URL=apps-demo05 && export HEC_PORT=8088 && export RECEIVER_PORT=9997 && export HEC_TOKEN=8834a4fb-2eb8-47f2-879e-3f2ce7a3c1d4 && export DIMENSIONS=owner:"splunk",location:"london"
METRIC_TYPES=cpu,df,disk,interface,load,memory METRIC_OPTS=cpu.by_cpu
LOG_SOURCES=/etc/collectd/collectd.log%collectd,\$SPLUNK_HOME/var/log/splunk/*.log*%uf,/var/log/syslog%syslog,/var/log/daemon.log%syslog,/var/log/auth.log%syslog AUTHENTICATED_INSTALL=Yes && wget --no-
check-certificate http://apps-demo05:8000/static/app/splunk_app_infrastructure/unix_agent/unix-agent.tgz && tar -xzf unix-agent.tgz && cd unix-agent && bash install_uf.sh && bash install_agent.sh && cd
.. && rm -rf unix-agent
```

3 Once the script finishes running, verify your data connection.

It may take up to 5 minutes to show entities in the UI.

No new entities connected yet...

Auto refresh in 7s

Instant Infrastructure Observability

Infrastructure components immediately available with no further configuration

- Automatically correlate entities, metadata, metrics and logs from servers, cloud platforms, virtualized environments and other infrastructure entities
- Easily isolate and investigate problem entities no matter source, data type or volume
- Create groups of entities to monitor, alert, and troubleshoot at scale

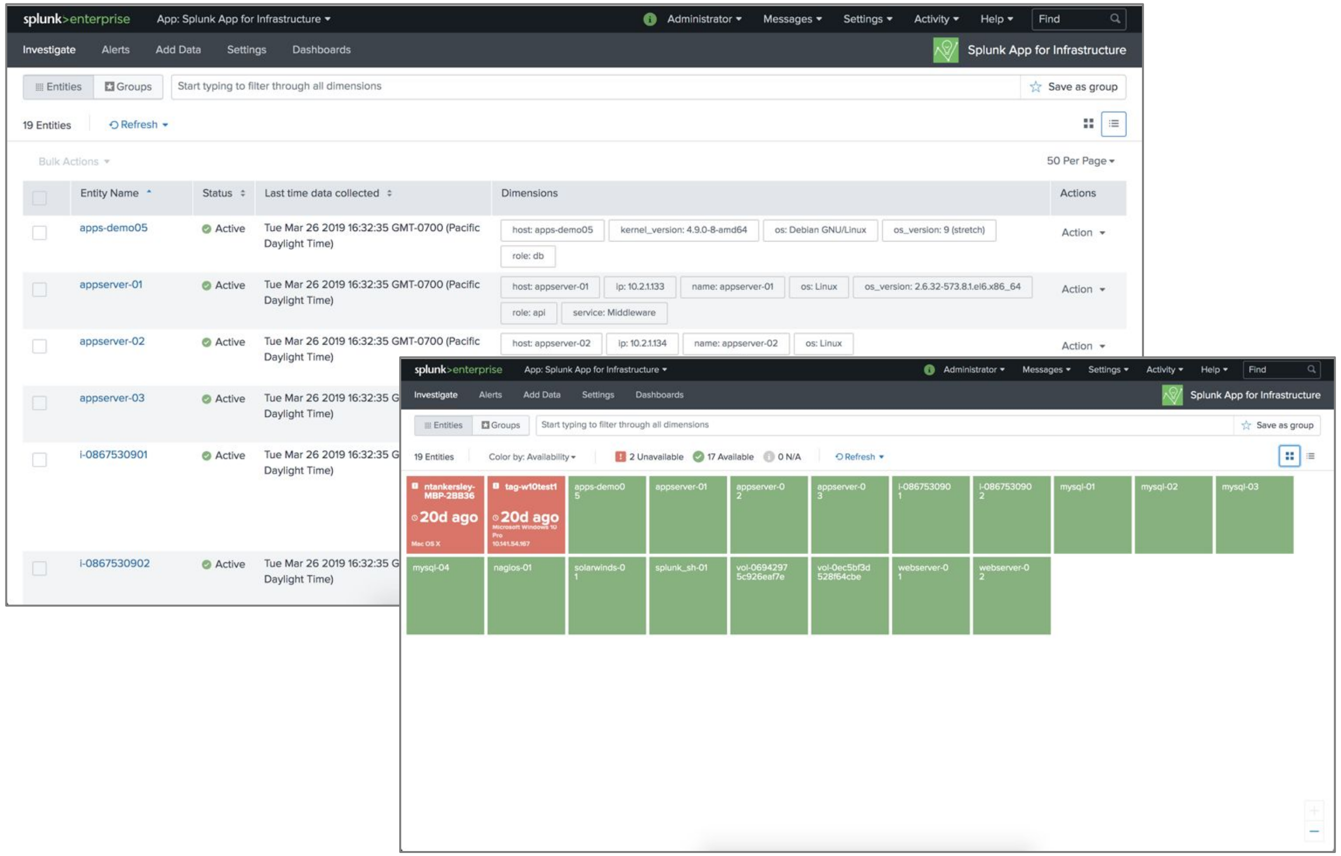

19 Entities

□	Entity Name	Status	Last time data collected	Dimensions	Actions
□	apps-demo05	Active	Tue Mar 26 2019 16:32:35 GMT-0700 (Pacific Daylight Time)	host: apps-demo05 kernel_version: 4.9.0-8-amd64 os: Debian GNU/Linux os_version: 9 (stretch) role: db	Action ▾
□	appserver-01	Active	Tue Mar 26 2019 16:32:35 GMT-0700 (Pacific Daylight Time)	host: appserver-01 ip: 10.2.1.133 name: appserver-01 os: Linux os_version: 2.6.32-573.8.1.el6.x86_64 role: api service: Middleware	Action ▾
□	appserver-02	Active	Tue Mar 26 2019 16:32:35 GMT-0700 (Pacific Daylight Time)	host: appserver-02 ip: 10.2.1.134 name: appserver-02 os: Linux	Action ▾
□	appserver-03	Active	Tue Mar 26 2019 16:32:35 G Daylight Time)		
□	i-0867530901	Active	Tue Mar 26 2019 16:32:35 G Daylight Time)		
□	i-0867530902	Active	Tue Mar 26 2019 16:32:35 G Daylight Time)		

19 Entities

Color by: Availability ▾

🟥 2 Unavailable ✓ 17 Available ● 0 N/A

ntankersley-MBP-2B836 20d ago Mac OS X	tag-w10test1 20d ago Microsoft Windows 10 Pro 10.141.54.167	apps-demo05	appserver-01	appserver-02	appserver-03	i-0867530901	i-0867530902	mysql-01	mysql-02	mysql-03
mysql-04	nagios-01	solarwinds-01	splunk_sh-01	vol-06942975c926eaf7e	vol-0ec5bf3d528f64cbe	webserver-01	webserver-02			

Single Experience for Multiple Activities

Monitor, alert, & investigate across your infrastructure in a single UI

- Use Workspaces to monitor a group or single entity in your infrastructure
- Quickly move from large groups of entities to focused views with rich out of the box content
- Visualize metrics and logs together – no expertise required
- Alert across thousands of entities or on individual entities, interfaces, disks or CPU

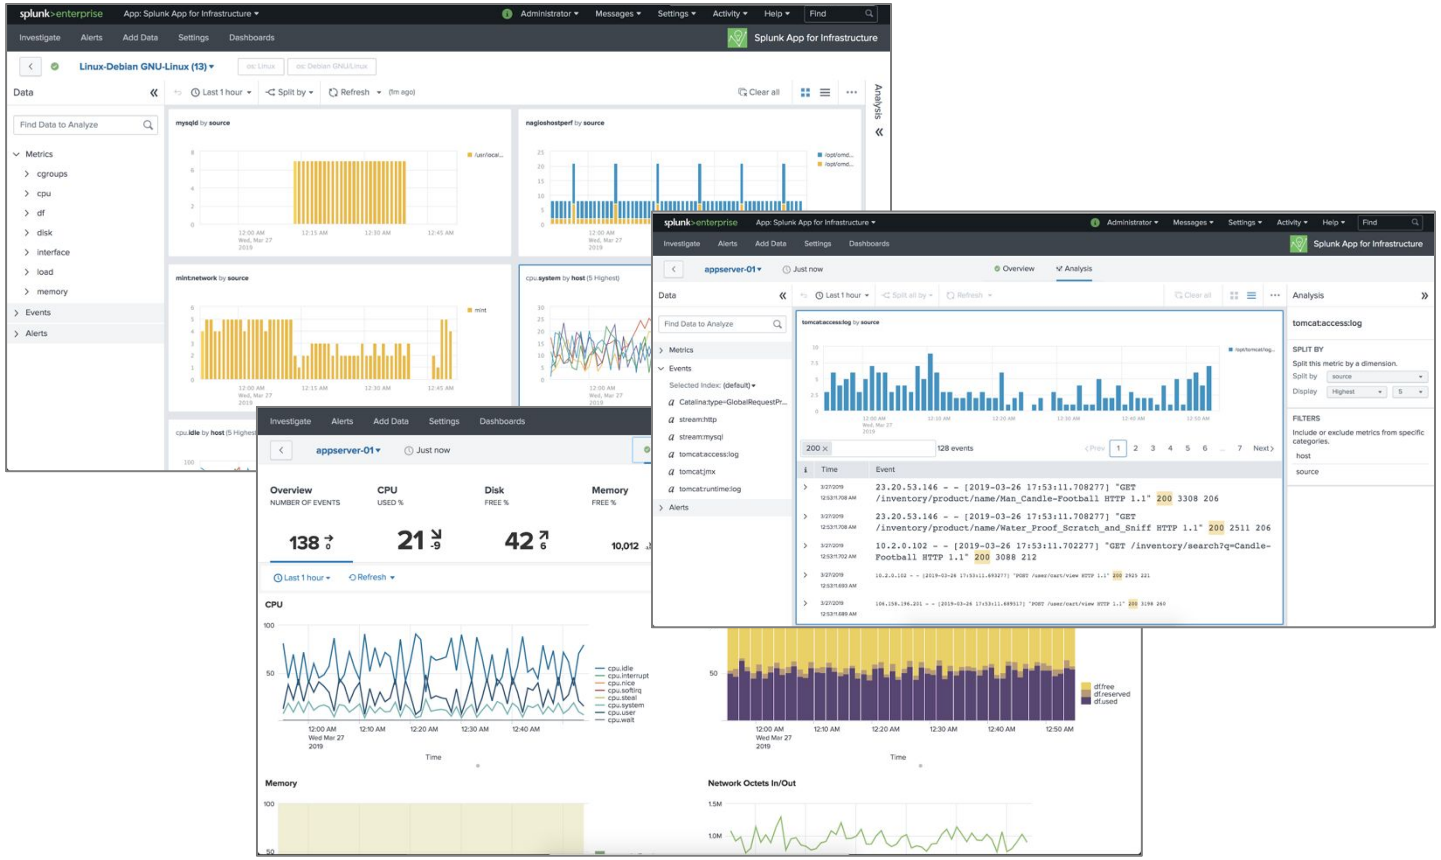

```
3/27/2019 12:53:11.708 AM  23.20.53.146 - - [2019-03-26 17:53:11.708277] "GET /inventory/product/name/Man_Candle-Football HTTP 1.1" 200 3308 206
3/27/2019 12:53:11.708 AM  23.20.53.146 - - [2019-03-26 17:53:11.708277] "GET /inventory/product/name/Water_Proof_Scratch_and_Sniff HTTP 1.1" 200 2511 206
3/27/2019 12:53:11.702 AM  10.2.0.102 - - [2019-03-26 17:53:11.702277] "GET /inventory/search?q=Candle-Football HTTP 1.1" 200 3088 212
3/27/2019 12:53:11.693 AM  10.2.0.102 - - [2019-03-26 17:53:11.693277] "POST /user/cart/view HTTP 1.1" 200 2925 221
3/27/2019 12:53:11.689 AM  106.158.196.201 - - [2019-03-26 17:53:11.689517] "POST /user/cart/view HTTP 1.1" 200 3198 260
```

Identify Root Cause in Two Clicks

Alert users to performance issues and lead them to the source in two clicks

- Toggle between group level and entity specific alerts
- Drill down to see the metrics affected and quickly begin finding root cause
- Alerts UI provides users with an easy to read summary of current status of groups and entities
- Focus on what's important not what's noisy

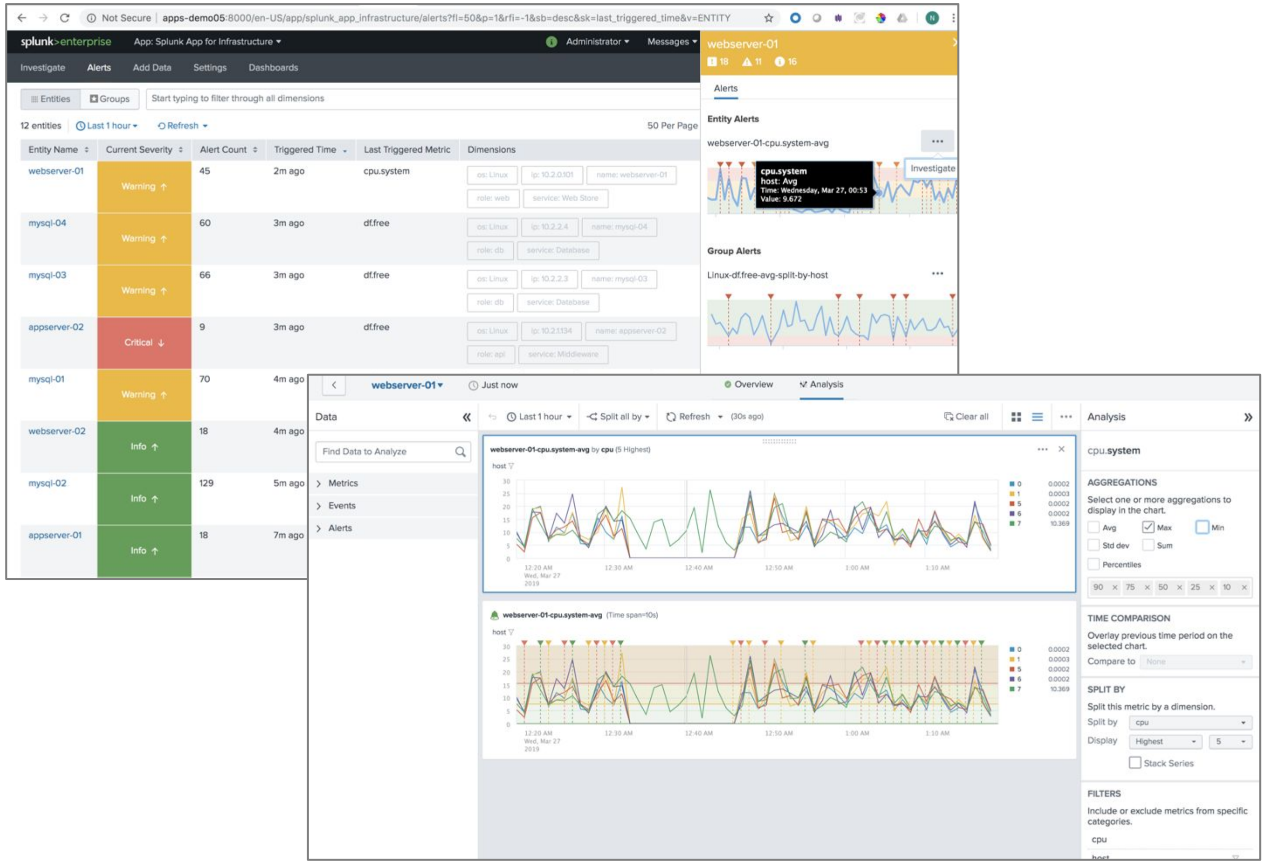

Entity Name	Current Severity	Alert Count	Triggered Time	Last Triggered Metric	Dimensions
webserver-01	Warning ↑	45	2m ago	cpu.system	os: Linux ip: 10.2.0.101 name: webserver-01 role: web service: Web Store
mysql-04	Warning ↑	60	3m ago	df.free	os: Linux ip: 10.2.2.4 name: mysql-04 role: db service: Database
mysql-03	Warning ↑	66	3m ago	df.free	os: Linux ip: 10.2.2.3 name: mysql-03 role: db service: Database
appserver-02	Critical ↓	9	3m ago	df.free	os: Linux ip: 10.2.1.134 name: appserver-02 role: api service: Middleware
mysql-01	Warning ↑	70	4m ago		
webserver-02	Info ↑	18	4m ago		
mysql-02	Info ↑	129	5m ago		
appserver-01	Info ↑	18	7m ago		

cpu.system
host: Avg
Time: Wednesday, Mar 27, 00:53
Value: 9.672

cpu.system

AGGREGATIONS

Select one or more aggregations to display in the chart.

□ Avg ☑ Max □ Min
□ Std dev □ Sum
□ Percentiles
90 × 75 × 50 × 25 × 10 ×

TIME COMPARISON

Overlay previous time period on the selected chart.

Compare to: None

SPLIT BY

Split this metric by a dimension.

Split by: cpu
Display: Highest 5
□ Stack Series

FILTERS

Include or exclude metrics from specific categories.

cpu
host

splunk> .conf19

Telegraf Dashboards in SAI 2.0

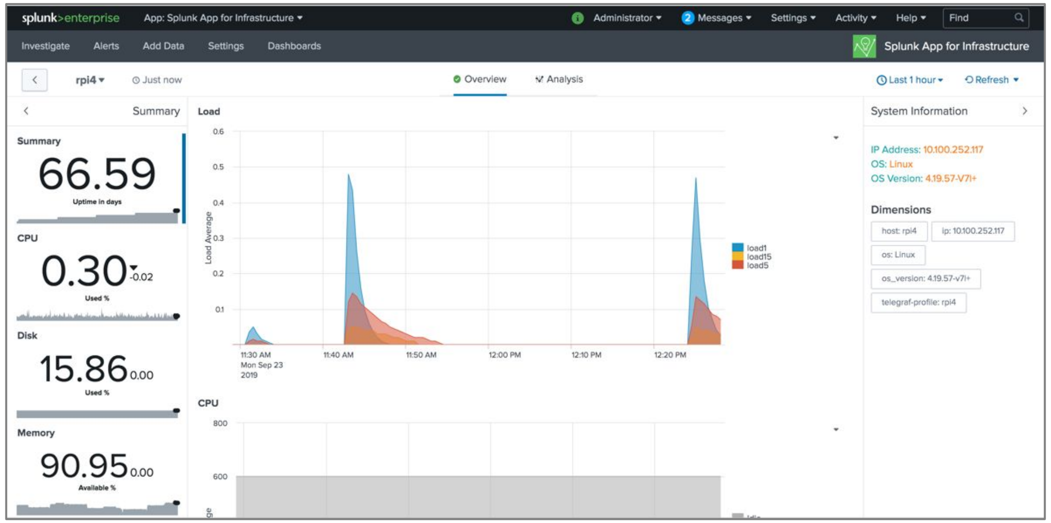

System Information

IP Address: 10.100.252.117
OS: Linux
OS Version: 4.19.57-V7l+

Dimensions

host: rpi4 | ip: 10.100.252.117 | os: Linux | os_version: 4.19.57-v7l+ | telegraf-profile: rpi4

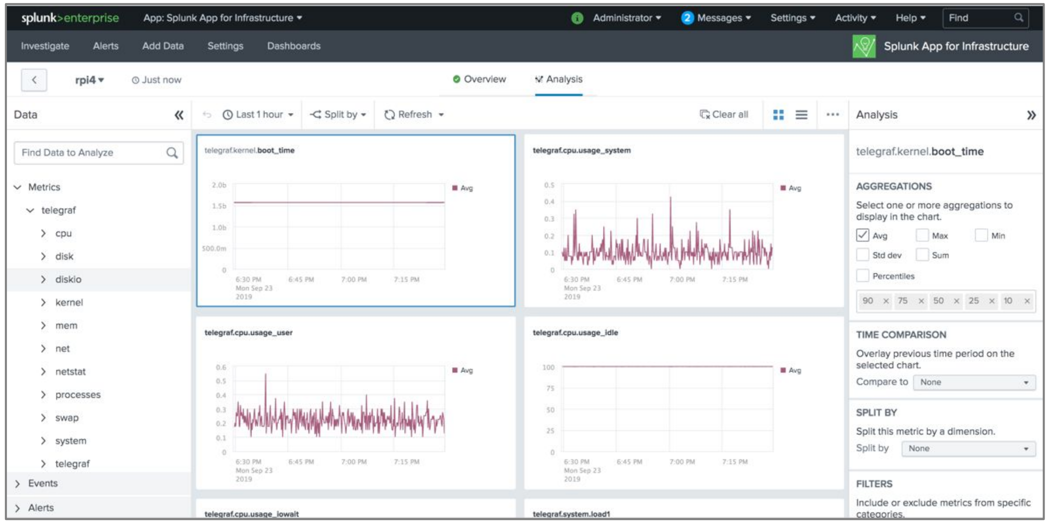

Analysis

telegraf.kernel.boot_time

AGGREGATIONS

Select one or more aggregations to display in the chart.

☑ Avg ☐ Max ☐ Min
☐ Std dev ☐ Sum
☐ Percentiles
90 × 75 × 50 × 25 × 10 ×

TIME COMPARISON

Overlay previous time period on the selected chart.

Compare to: None

SPLIT BY

Split this metric by a dimension.

Split by: None

FILTERS

Include or exclude metrics from specific categories.

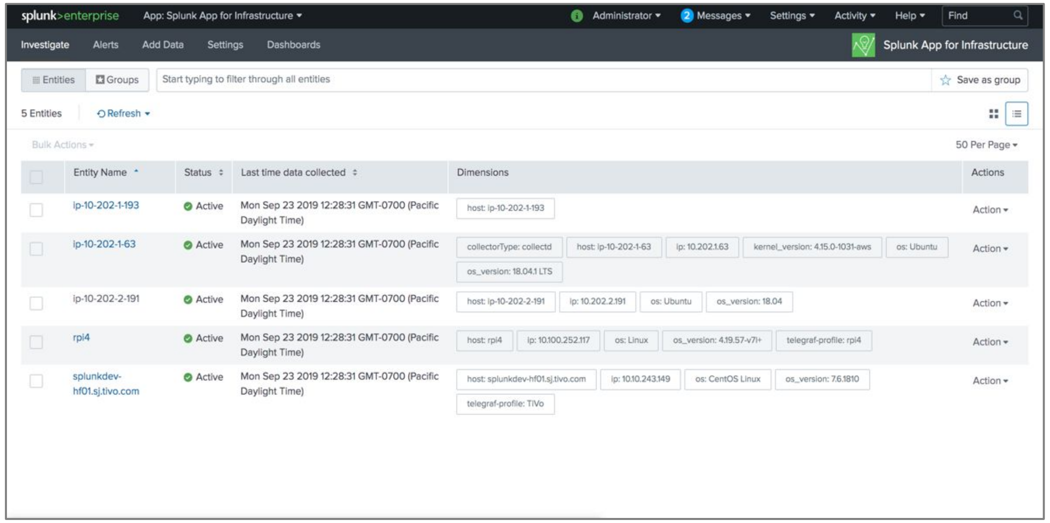

5 Entities

Entity Name	Status	Last time data collected	Dimensions	Actions
ip-10-202-1-193	Active	Mon Sep 23 2019 12:28:31 GMT-0700 (Pacific Daylight Time)	host: ip-10-202-1-193	Action ▾
ip-10-202-1-63	Active	Mon Sep 23 2019 12:28:31 GMT-0700 (Pacific Daylight Time)	collectorType: collectd host: ip-10-202-1-63 ip: 10.202.1.63 kernel_version: 4.15.0-1031-aws os: Ubuntu os_version: 18.04.1 LTS	Action ▾
ip-10-202-2-191	Active	Mon Sep 23 2019 12:28:31 GMT-0700 (Pacific Daylight Time)	host: ip-10-202-2-191 ip: 10.202.2.191 os: Ubuntu os_version: 18.04	Action ▾
rpi4	Active	Mon Sep 23 2019 12:28:31 GMT-0700 (Pacific Daylight Time)	host: rpi4 ip: 10.100.252.117 os: Linux os_version: 4.19.57-v7l+ telegraf-profile: rpi4	Action ▾
splunkdev-hf01.sj.tivo.com	Active	Mon Sep 23 2019 12:28:31 GMT-0700 (Pacific Daylight Time)	host: splunkdev-hf01.sj.tivo.com ip: 10.10.243.149 os: CentOS Linux os_version: 7.6.1810 telegraf-profile: TiVo	Action ▾

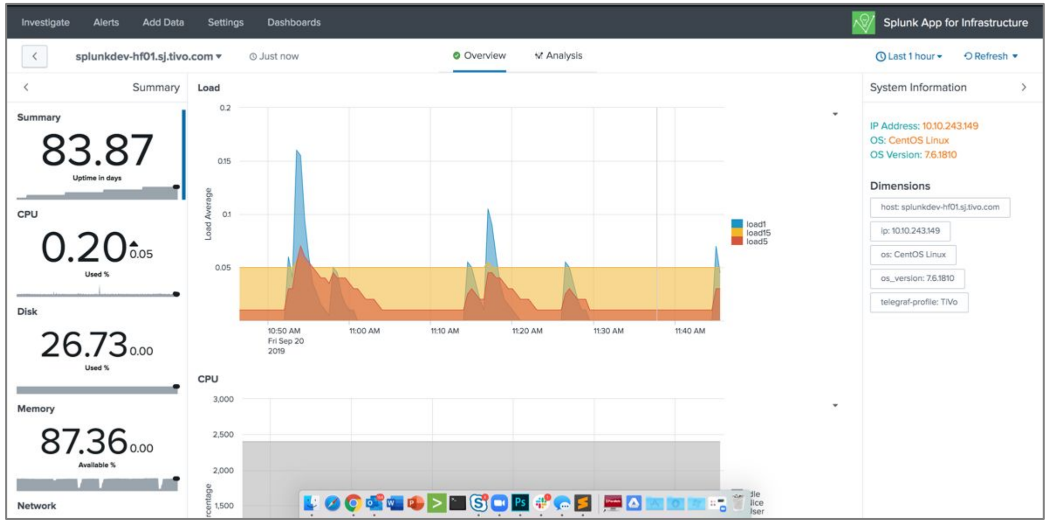

System Information

IP Address: 10.10.243.149
OS: CentOS Linux
OS Version: 7.6.1810

Dimensions

host: splunkdev-hf01.sj.tivo.com | ip: 10.10.243.149 | os: CentOS Linux | os_version: 7.6.1810 | telegraf-profile: TiVo

SAI 1.0

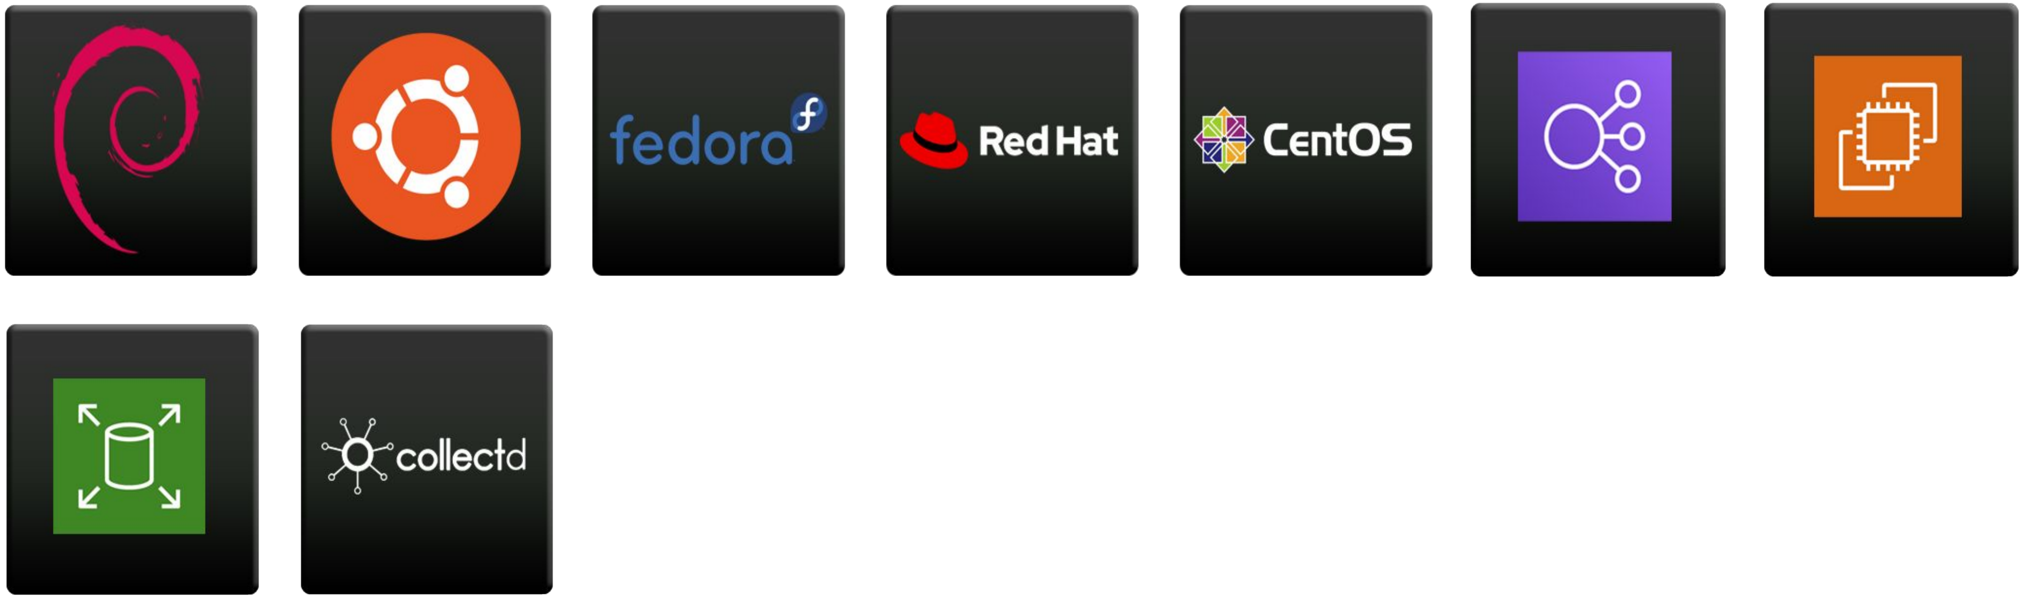

Metrics LOG Logs Configs/Metadata

SAI 1.1

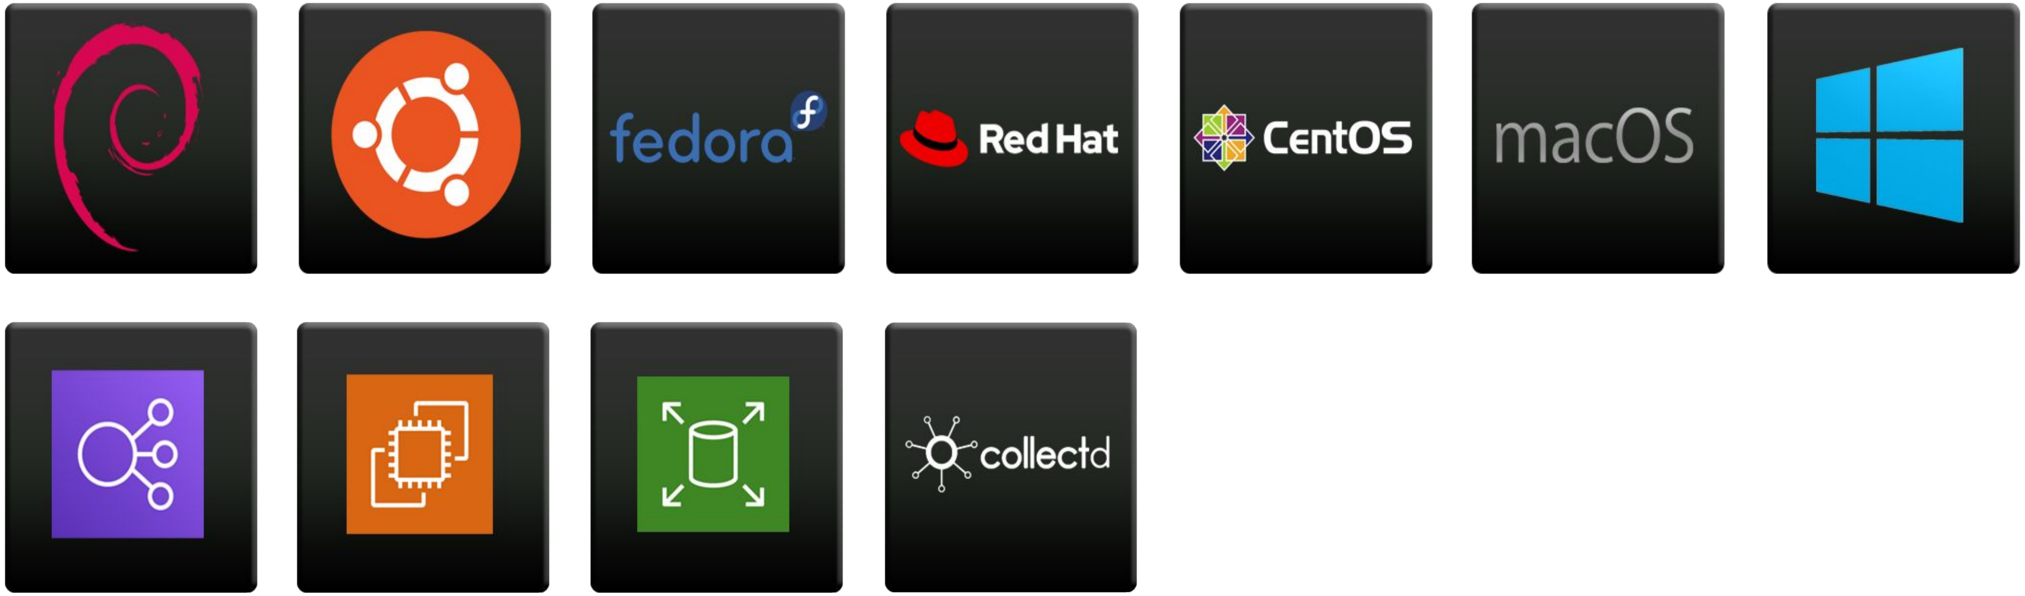

Metrics Logs Configs/Metadata

SAI 1.2

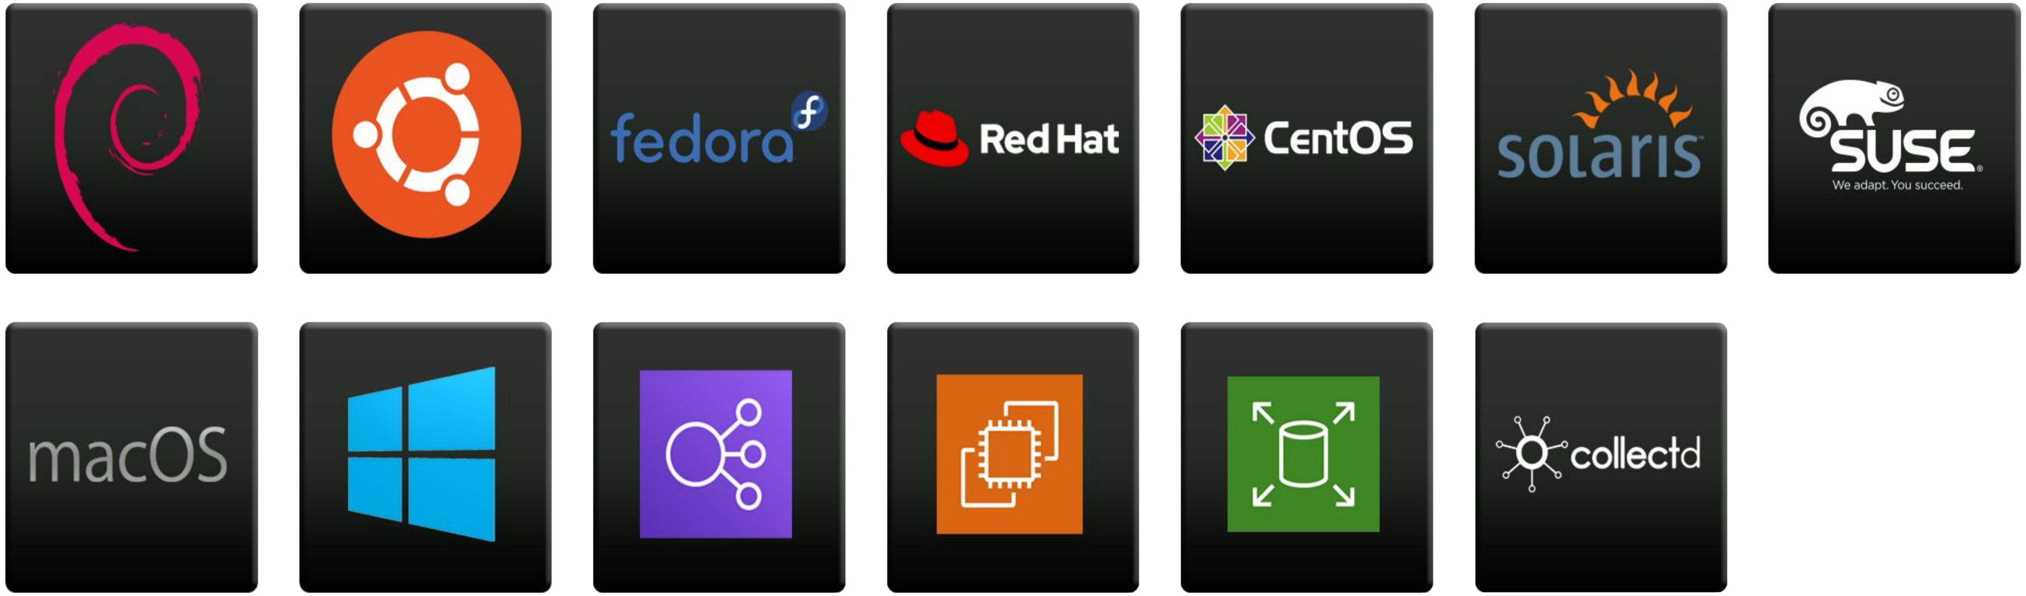

Metrics LOG Logs Configs/Metadata

SAI 1.4

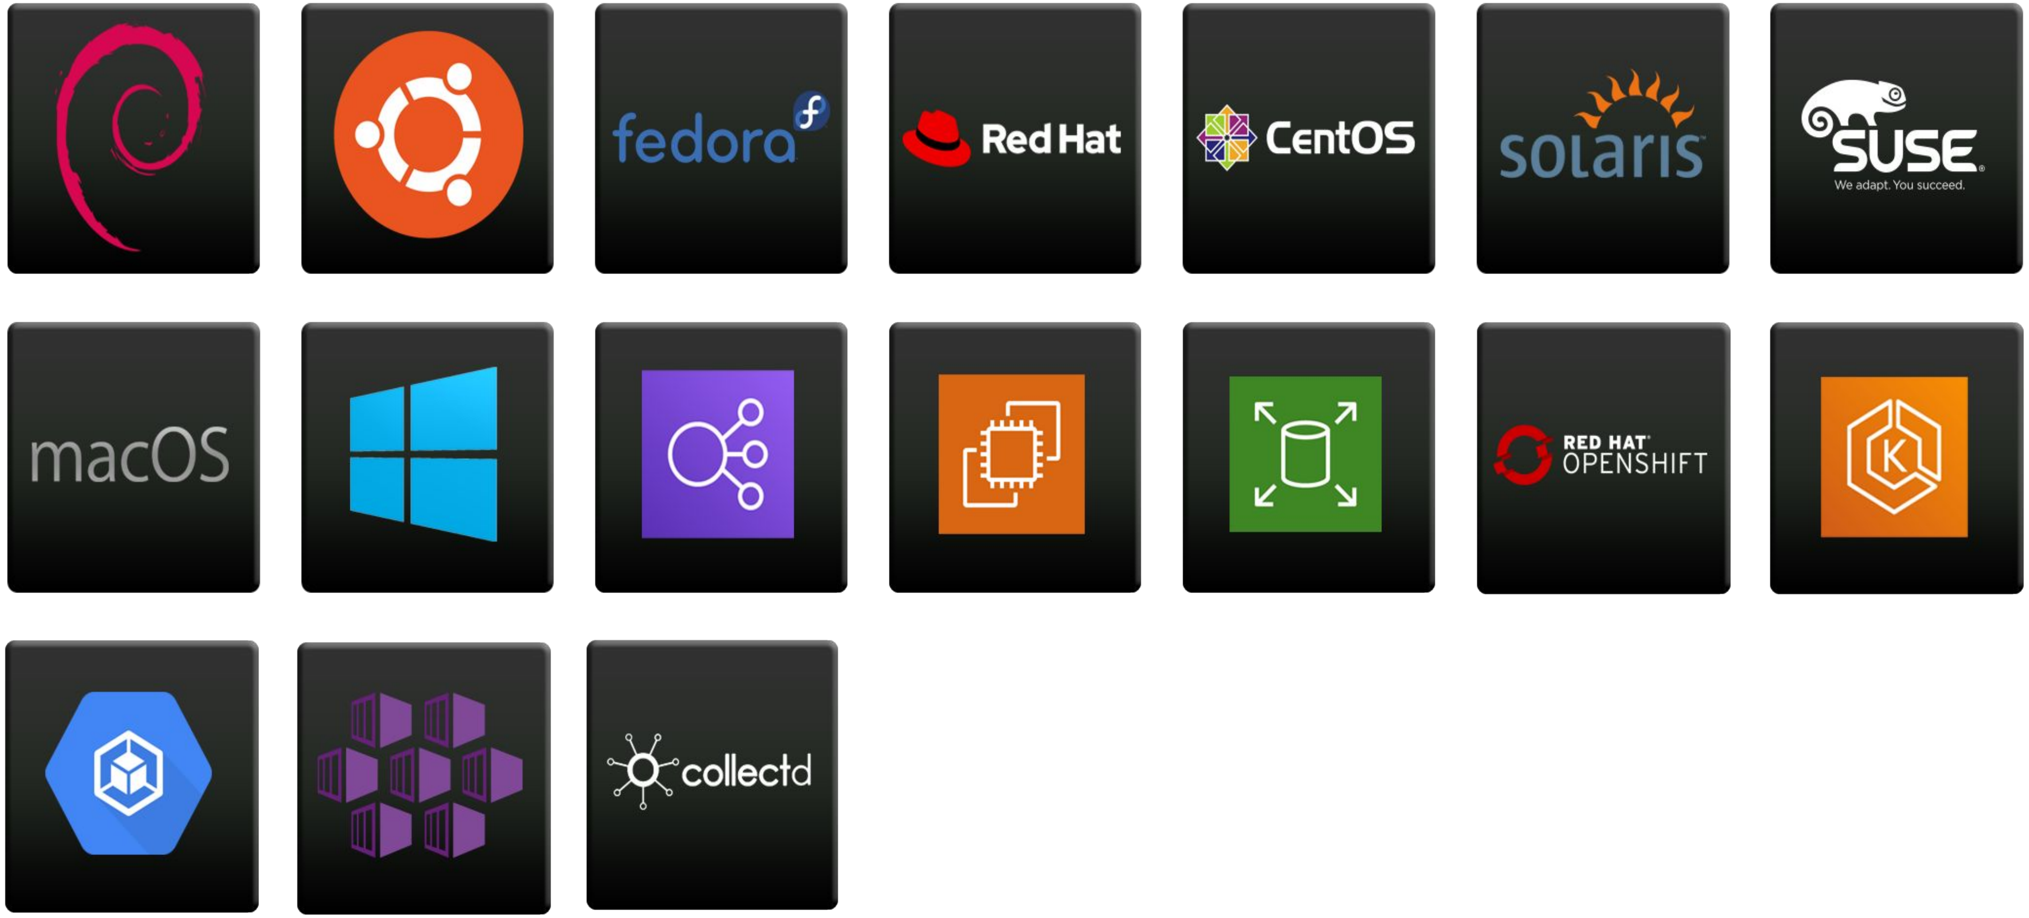

{ } Metrics LOG Logs Configs/Metadata

SAI 2.0

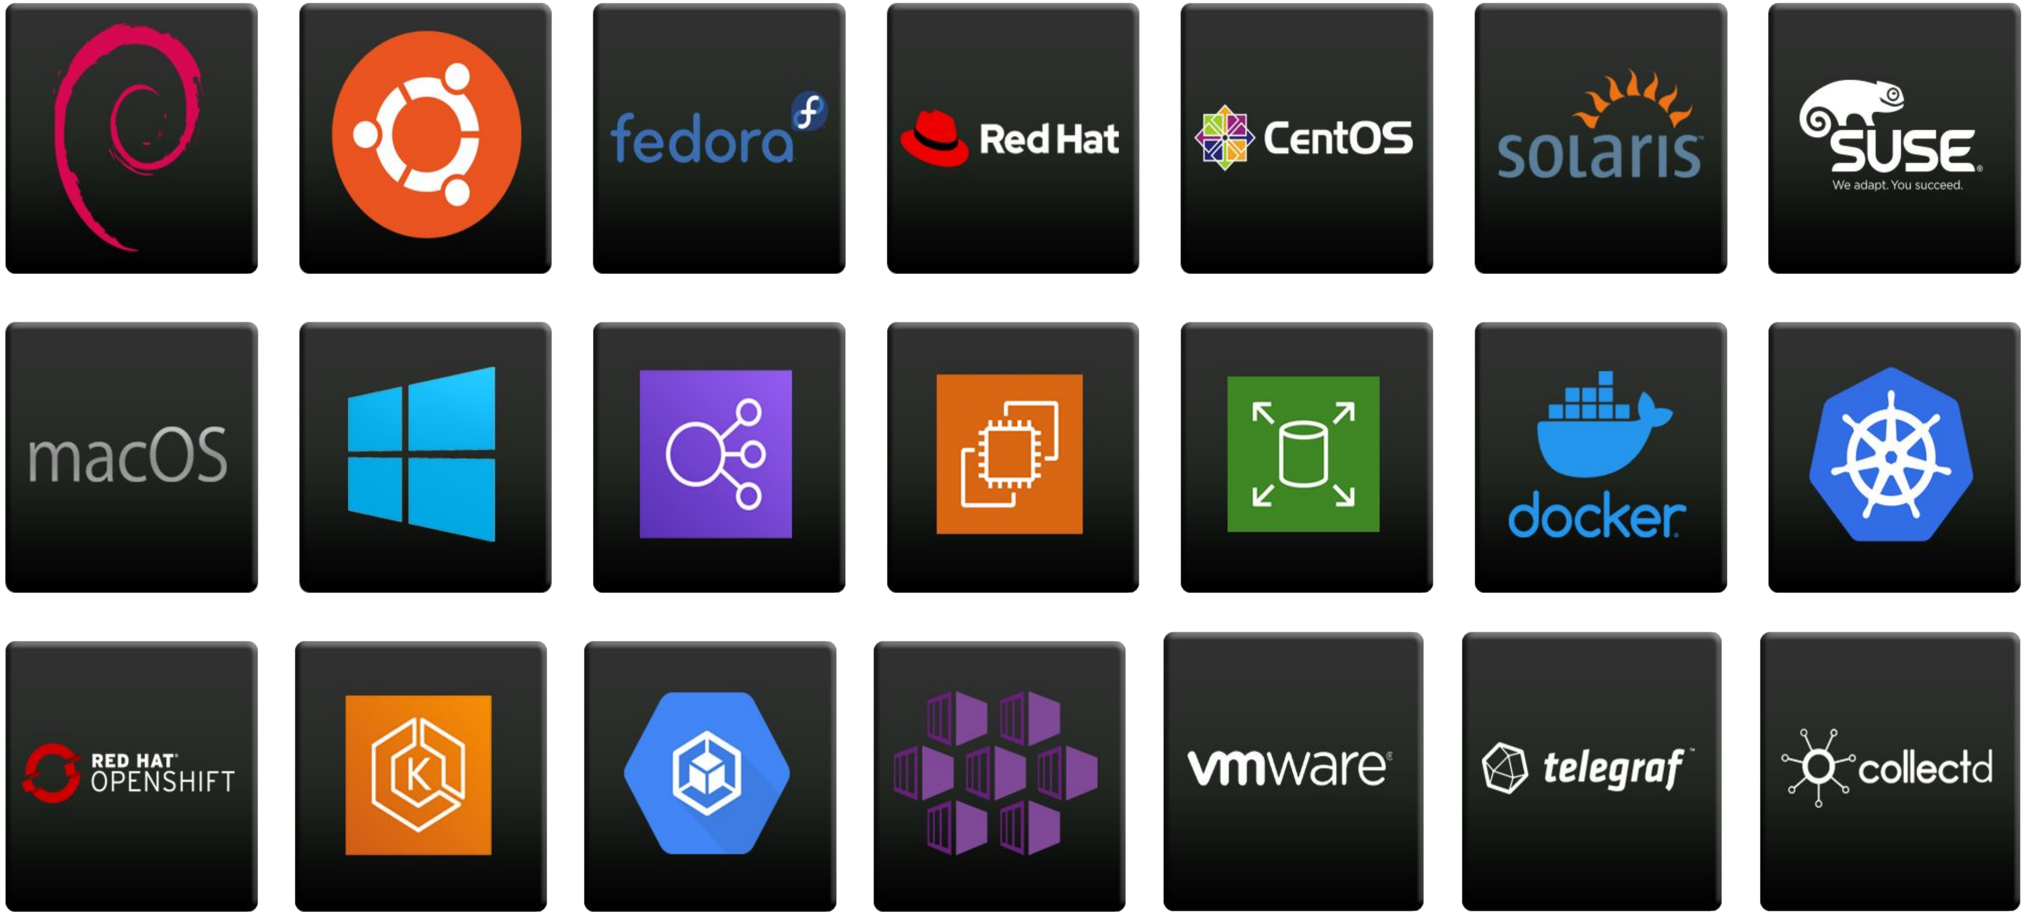

Metrics LOG Logs Configs/Metadata

ITSI Integrations

- ITSI 4.2.0 introduced integrations with SAI
  - SAI entities are able to be shared with ITSI
  - SAI exports several KPI Base searches for use by ITSI
- Telegraf based entities will be available in ITSI

ITSI Integration

A Telegraf and collectd entity under the same service

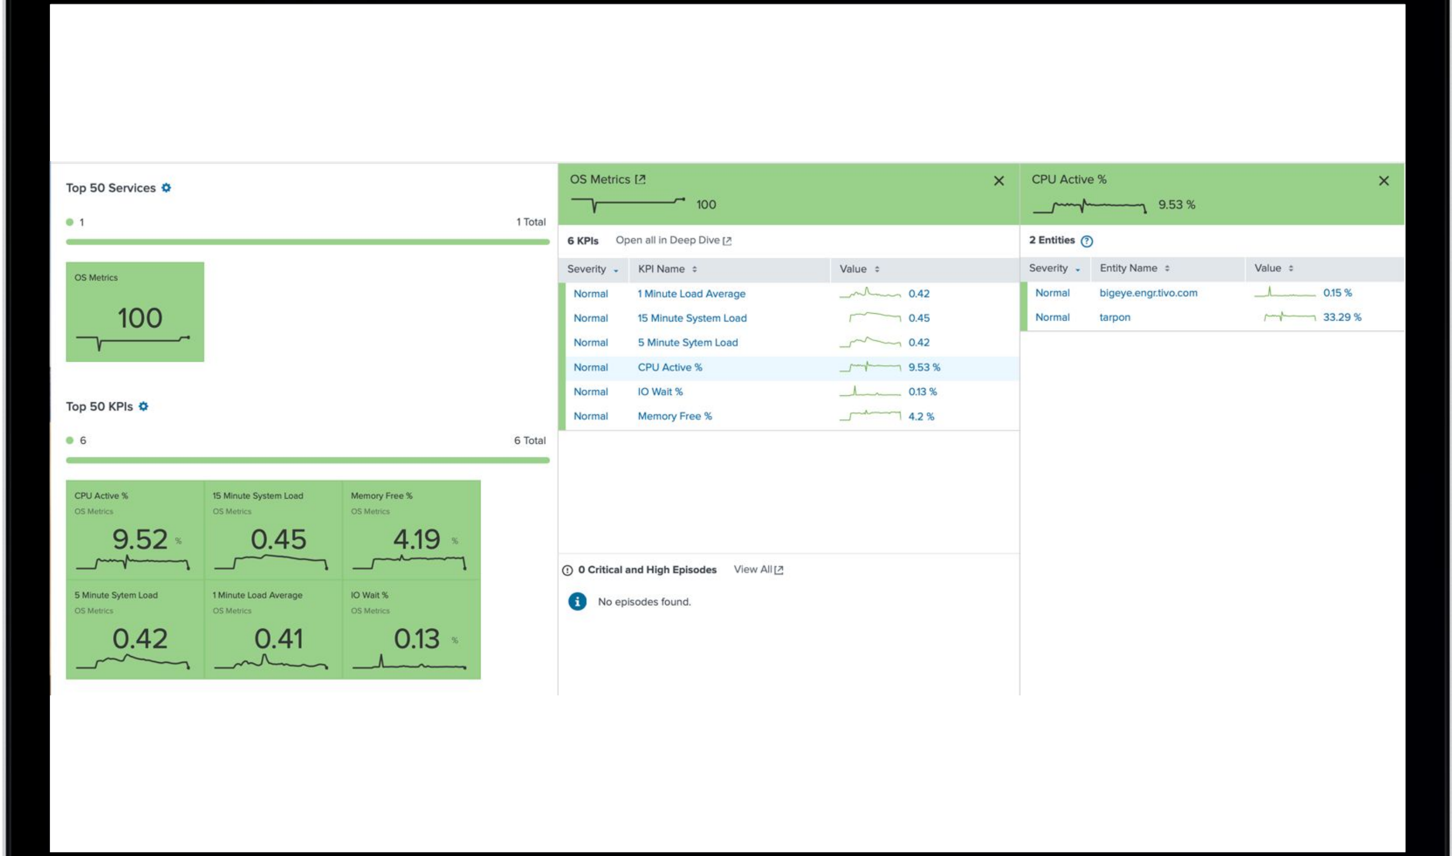

Top 50 Services

● 1 1 Total

OS Metrics
100

Top 50 KPIs

● 6 6 Total

CPU Active %
OS Metrics
9.52 %

15 Minute System Load
OS Metrics
0.45

Memory Free %
OS Metrics
4.19 %

5 Minute Sytem Load
OS Metrics
0.42

1 Minute Load Average
OS Metrics
0.41

IO Wait %
OS Metrics
0.13 %

OS Metrics
100

6 KPIs Open all in Deep Dive

Severity	KPI Name	Value
Normal	1 Minute Load Average	0.42
Normal	15 Minute System Load	0.45
Normal	5 Minute Sytem Load	0.42
Normal	CPU Active %	9.53 %
Normal	IO Wait %	0.13 %
Normal	Memory Free %	4.2 %

0 Critical and High Episodes View All

No episodes found.

CPU Active %
9.53 %

2 Entities

Severity	Entity Name	Value
Normal	bigeye.engr.tivo.com	0.15 %
Normal	tarpon	33.29 %

SPL Usage

As you would expect, once the data is in Splunk, it can be used just like any other metric data. You can make detailed graphs of all your metrics on in your existing dashboards.

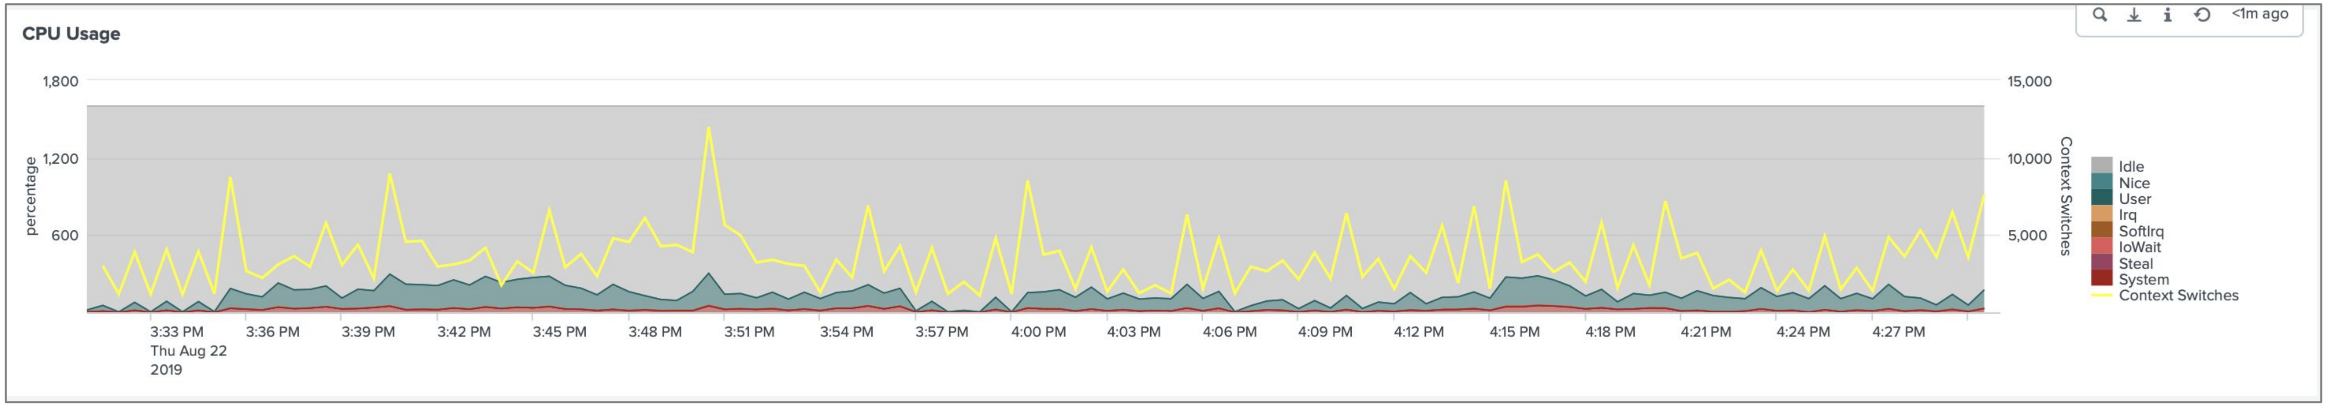

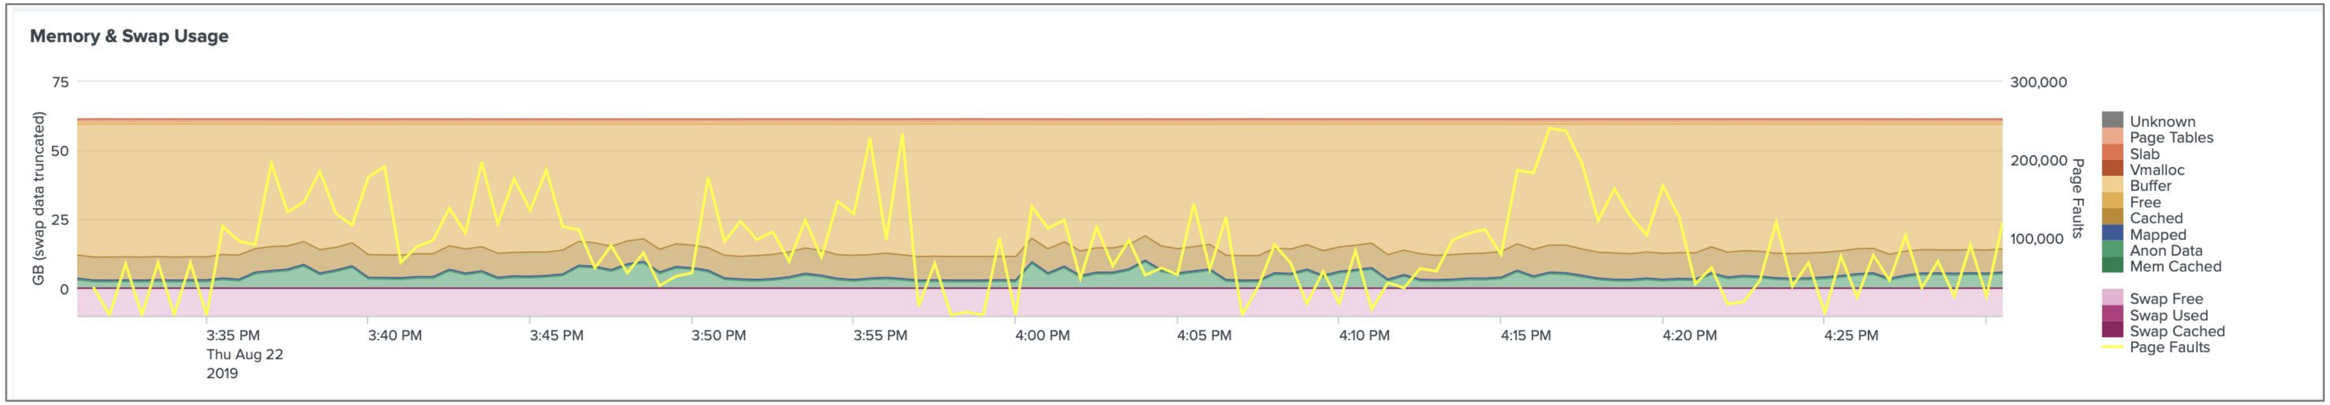

Splunk Integrations

1. Telegraf metrics are recognized by Splunk Application for Infrastructure

2. SAI 2.0 introduces native Telegraf collectors and dashboards

3. Telegraf based entities and data are recognized by ITSI using the SAI integration

4. Telegraf data can be used directly using mstats and friends for detailed metrics reporting

Learn More About Splunk App for Infrastructure with these Sessions

1. IT1766 – Monitoring your VMware vSphere Environment with Splunk.
   - **Wednesday, October 23, 11:15am-12:00pm**

2. IT2001 – Monitoring and troubleshooting workloads running on public cloud infrastructure made easy.
   - **Wednesday, October 23, 11:15am-12:00pm**

Q&A

splunk> .conf19

.conf19

splunk>

Thank You !

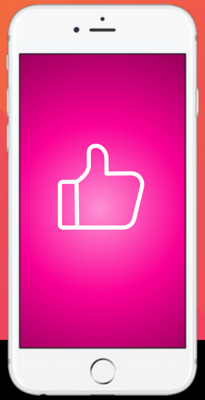

Go to the .conf19 mobile app to

RATE THIS SESSION