Forward-Looking Statements

This presentation may contain forward-looking statements regarding future events, plans or the expected financial performance of our company, including our expectations regarding our products, technology, strategy, customers, markets, acquisitions and investments. These statements reflect management's current expectations, estimates and assumptions based on the information currently available to us. These forward-looking statements are not guarantees of future performance and involve significant risks, uncertainties and other factors that may cause our actual results, performance or achievements to be materially different from results, performance or achievements expressed or implied by the forward-looking statements contained in this presentation.

For additional information about factors that could cause actual results to differ materially from those described in the forward-looking statements made in this presentation, please refer to our periodic reports and other filings with the SEC, including the risk factors identified in our most recent quarterly reports on Form 10-Q and annual reports on Form 10-K, copies of which may be obtained by visiting the Splunk Investor Relations website at www.investors.splunk.com or the SEC's website at www.sec.gov. The forward-looking statements made in this presentation are made as of the time and date of this presentation. If reviewed after the initial presentation, even if made available by us, on our website or otherwise, it may not contain current or accurate information. We disclaim any obligation to update or revise any forward-looking statement based on new information, future events or otherwise, except as required by applicable law.

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. We undertake no obligation either to develop the features or functionalities described, in beta or in preview (used interchangeably), or to include any such feature or functionality in a future release.

Getting Data in More Efficiently Using the Splunk® Edge Processor

PLA1870A

Ben Ferguson

Senior Information Security Engineer | Principal Financial Group

Yogesh Sontakke

Director, Product Management | Splunk

splunk> .conf23

Ben Ferguson

Senior Information Security Engineer
Principal Financial Group

Yogesh Sontakke

Director, Product Management | Splunk

Edge Processor Overview

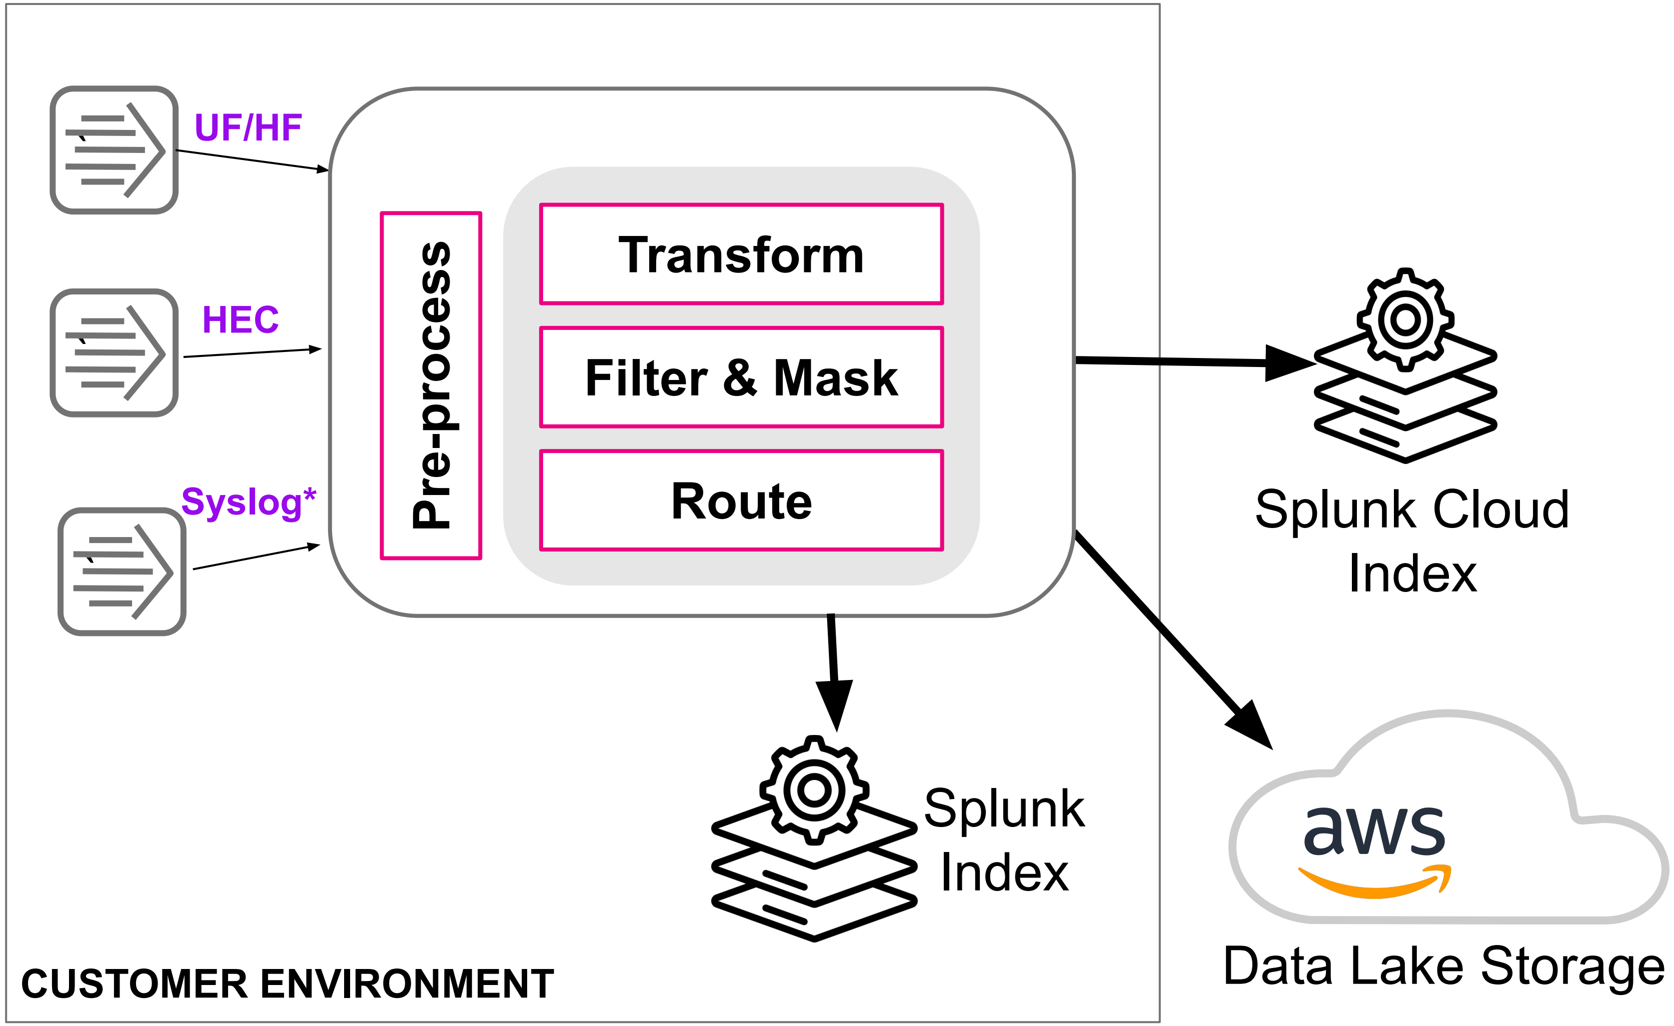

- Filter verbose or low-value sources, like DEBUG logs or other noisy data
- Extract just the critical data
- Route different “slices” of data to desired destinations

Amazon Web Services (AWS) are trademarks of Amazon.com, Inc. or its affiliates.

* = Not a GA Feature yet. Close to release.

Edge Processor Architecture

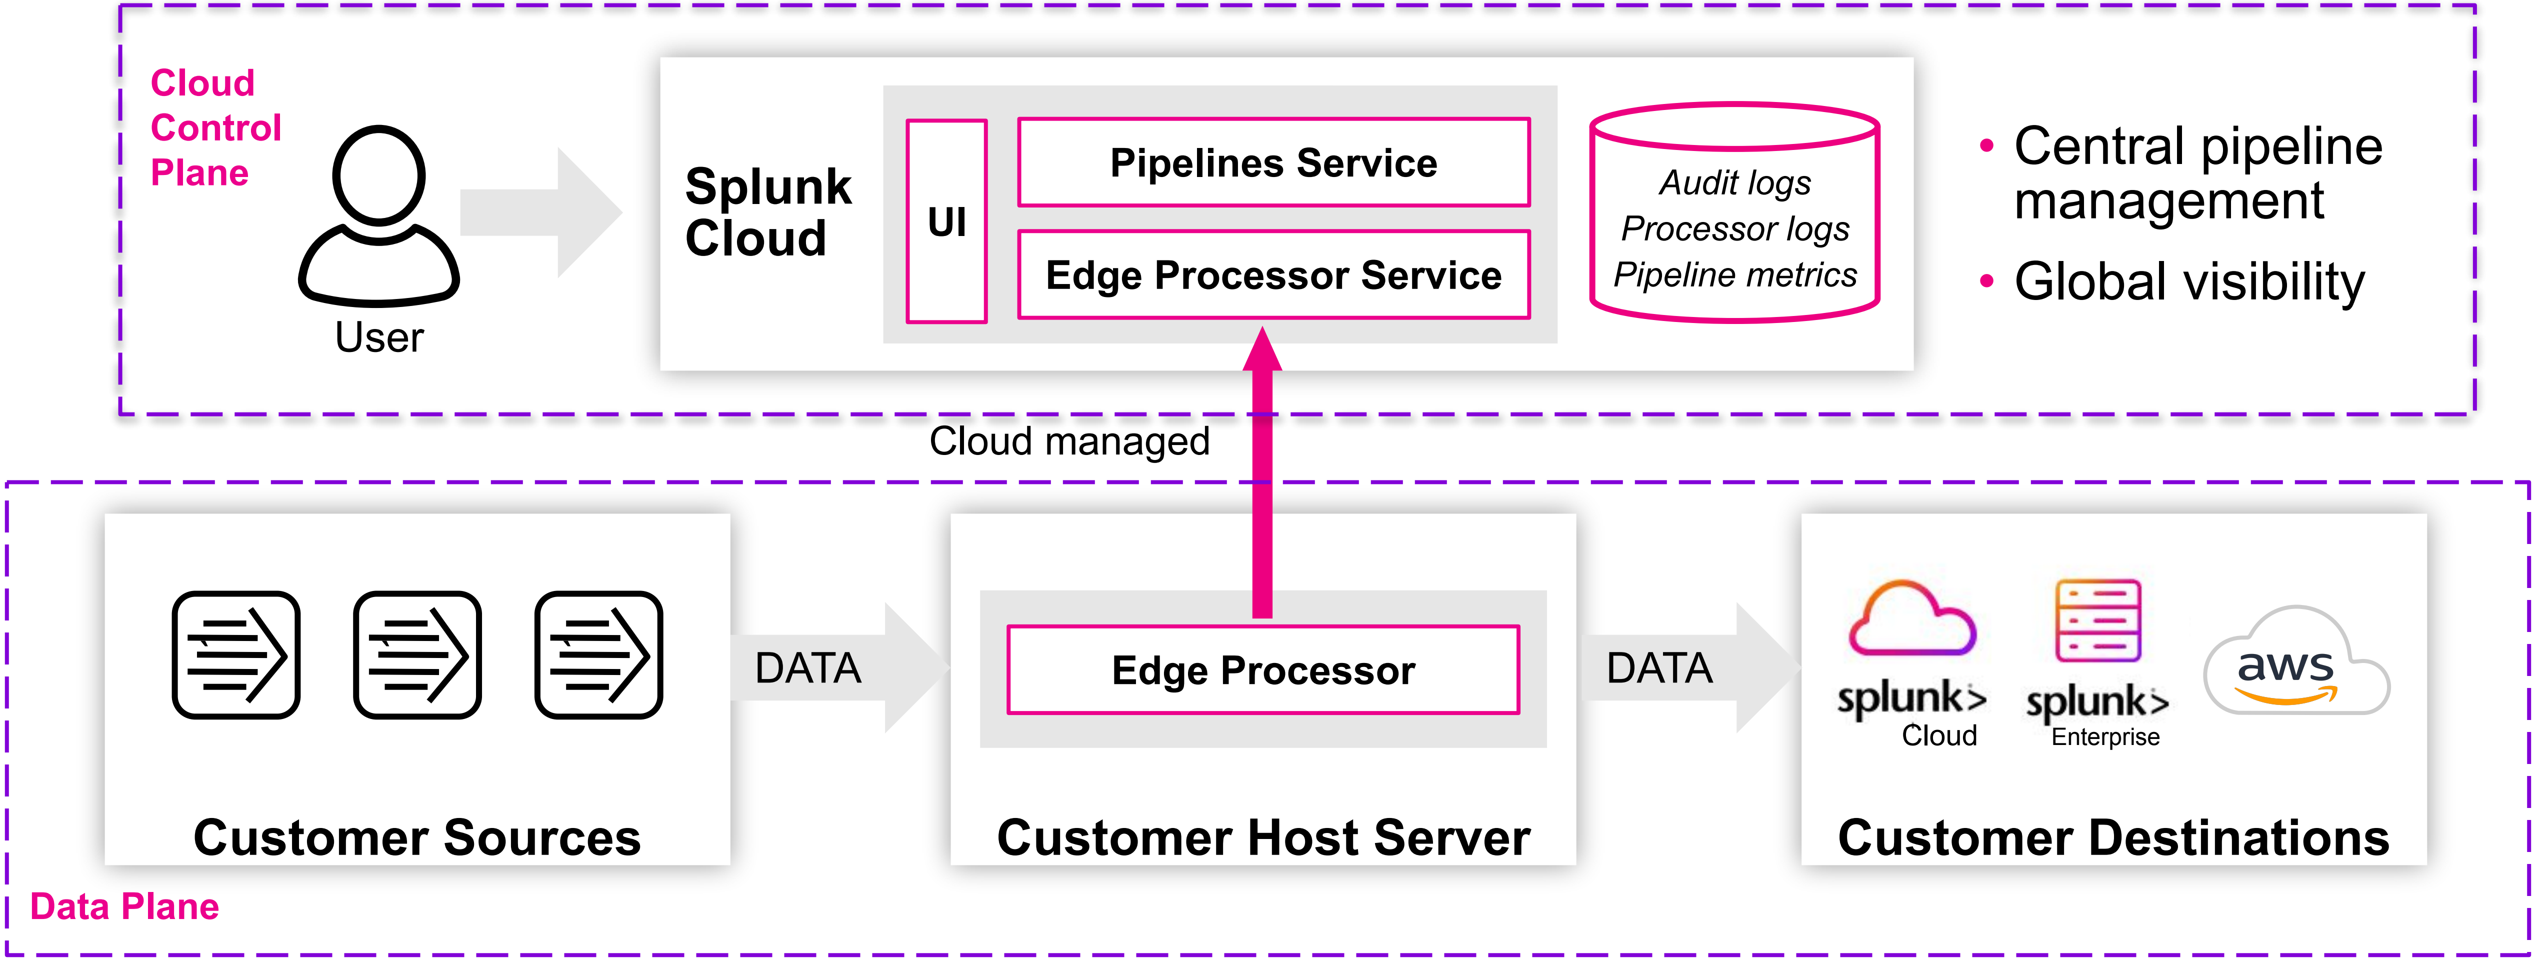

- Central pipeline management
- Global visibility

Who is Edge Processor for?

Admins and Architects

Administer and Manage Edge Processor nodes/clusters and overall deployments across multiple regions from a one-stop-shop Data Management Console

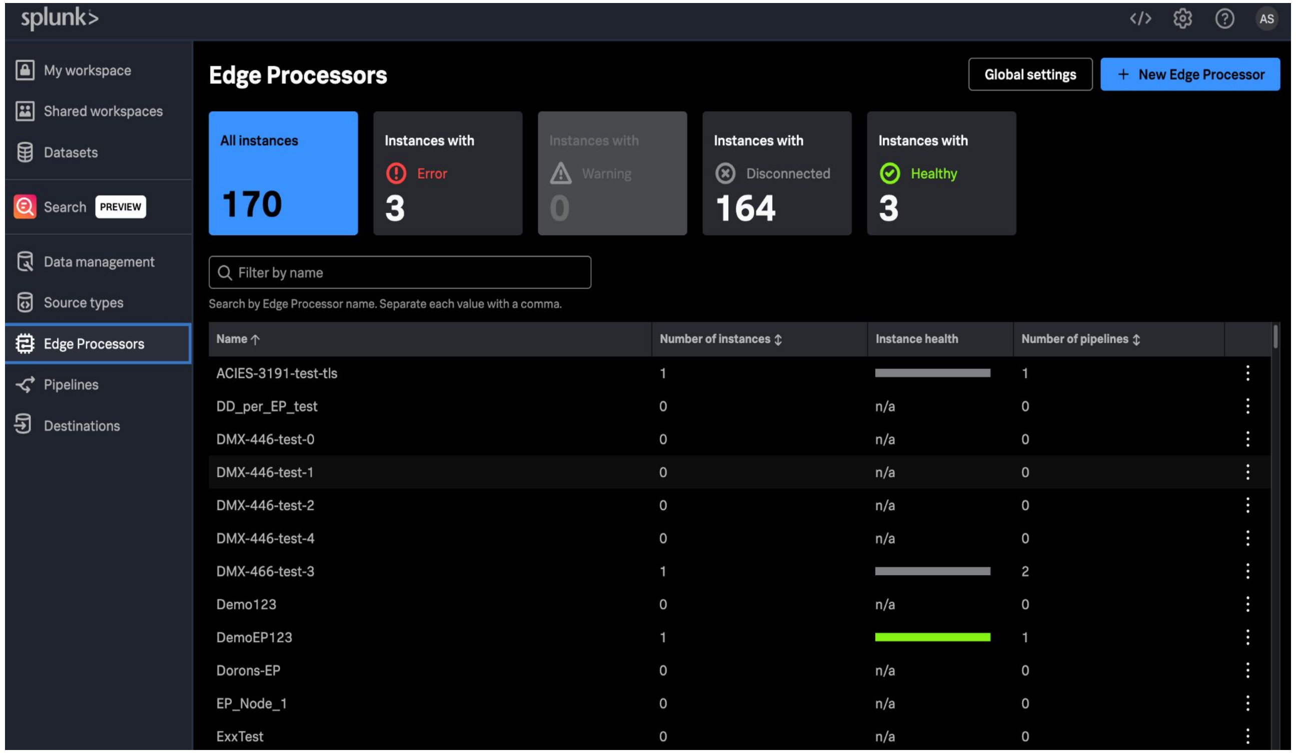

Edge Processors

All instances 170

Instances with Error 3

Instances with Warning 0

Instances with Disconnected 164

Instances with Healthy 3

Filter by name

Search by Edge Processor name. Separate each value with a comma.

Name ↑	Number of instances ↕	Instance health	Number of pipelines ↕
ACIES-3191-test-tls	1	▬	1
DD_per_EP_test	0	n/a	0
DMX-446-test-0	0	n/a	0
DMX-446-test-1	0	n/a	0
DMX-446-test-2	0	n/a	0
DMX-446-test-4	0	n/a	0
DMX-466-test-3	1	▬	2
Demo123	0	n/a	0
DemoEP123	1	🟩	1
Dorons-EP	0	n/a	0
EP_Node_1	0	n/a	0
ExxTest	0	n/a	0

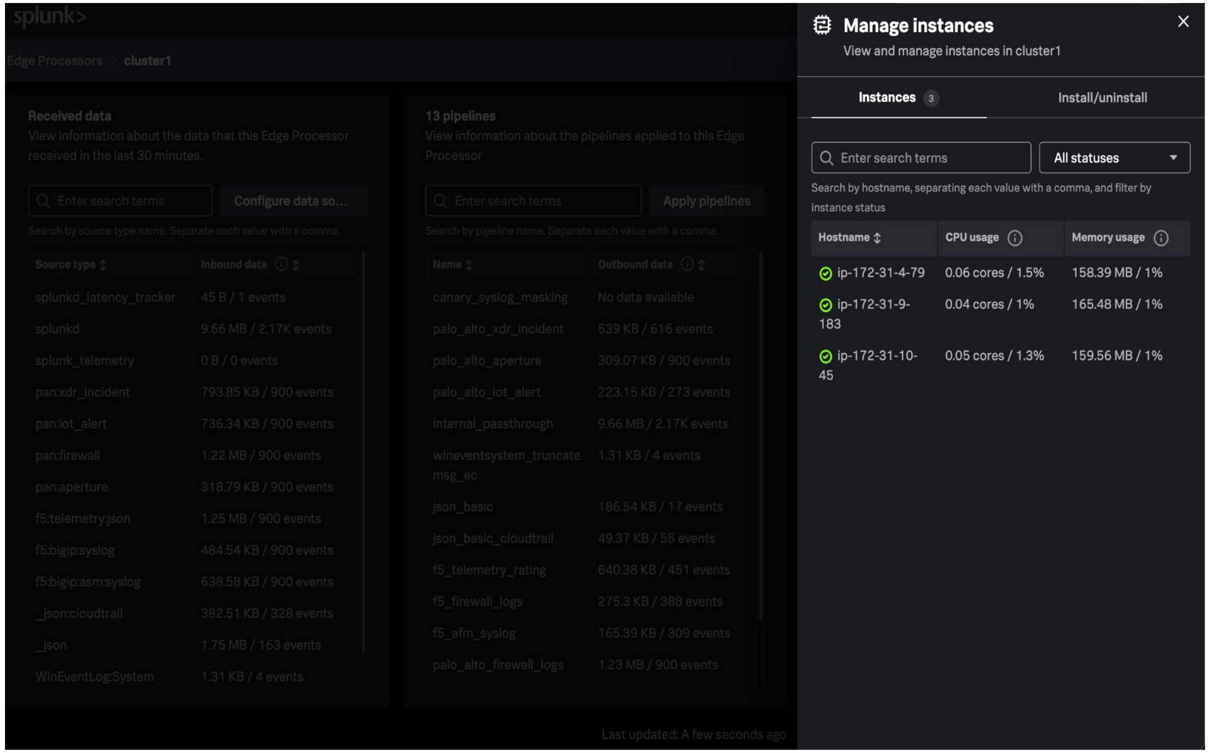

Received data

View information about the data that this Edge Processor received in the last 30 minutes.

Enter search terms

Search by source type name. Separate each value with a comma.

Source type ↕	Inbound data ⓘ ↕
splunkd_latency_tracker	45 B / 1 events
splunkd	9.66 MB / 2.17K events
splunk_telemetry	0 B / 0 events
pan:xdr_incident	793.85 KB / 900 events
pan:iot_alert	736.34 KB / 900 events
pan:firewall	1.22 MB / 900 events
pan:aperture	318.79 KB / 900 events
f5:telemetry:json	1.25 MB / 900 events
f5:bigip:syslog	484.54 KB / 900 events
f5:bigip:asm:syslog	638.58 KB / 900 events
_json:cloudtrail	382.51 KB / 328 events
_json	1.75 MB / 163 events
WinEventLog:System	1.31 KB / 4 events

13 pipelines

View information about the pipelines applied to this Edge Processor

Enter search terms

Search by pipeline name. Separate each value with a comma.

Name ↕	Outbound data ⓘ ↕
canary_syslog_masking	No data available
palo_alto_xdr_incident	539 KB / 616 events
palo_alto_aperture	309.07 KB / 900 events
palo_alto_iot_alert	223.15 KB / 273 events
internal_passthrough	9.66 MB / 2.17K events
wineventsystem_truncatemsg_ec	1.31 KB / 4 events
json_basic	186.54 KB / 17 events
json_basic_cloudtrail	49.37 KB / 55 events
f5_telemetry_rating	640.38 KB / 451 events
f5_firewall_logs	275.3 KB / 388 events
f5_afm_syslog	165.39 KB / 309 events
palo_alto_firewall_logs	1.23 MB / 900 events

Last updated: A few seconds ago

Manage instances

View and manage instances in cluster1

Enter search terms

Search by hostname, separating each value with a comma, and filter by instance status

Hostname ↕	CPU usage ⓘ	Memory usage ⓘ
✓ ip-172-31-4-79	0.06 cores / 1.5%	158.39 MB / 1%
✓ ip-172-31-9-183	0.04 cores / 1%	165.48 MB / 1%
✓ ip-172-31-10-45	0.05 cores / 1.3%	159.56 MB / 1%

Who is Edge Processor for?

Subject Matter Expert/SPL2 Content Creator

Author, test (in real-time) and share SPL2 content from the Data Management Console before deploying to your Edge Processor nodes

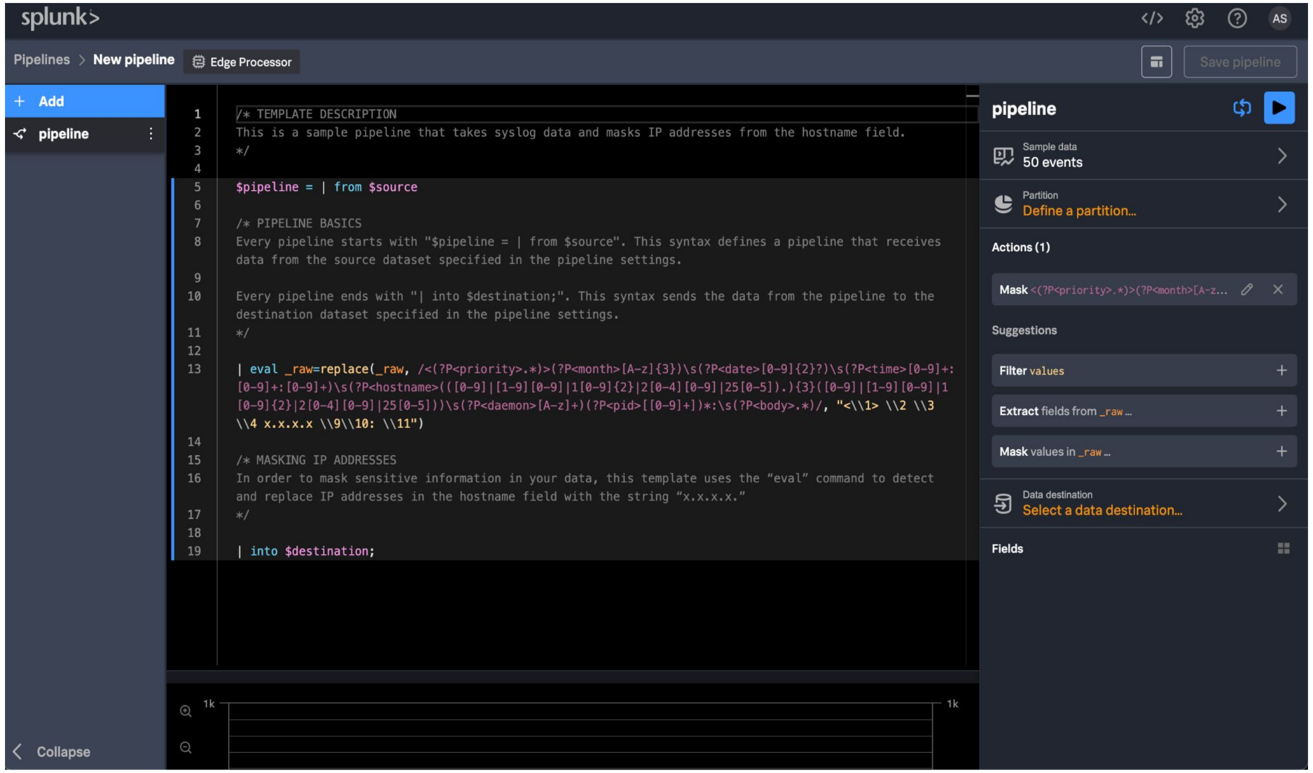

```
/* TEMPLATE DESCRIPTION
This is a sample pipeline that takes syslog data and masks IP addresses from the hostname field.
*/

$pipeline = | from $source

/* PIPELINE BASICS
Every pipeline starts with "$pipeline = | from $source". This syntax defines a pipeline that receives
data from the source dataset specified in the pipeline settings.

Every pipeline ends with "| into $destination;". This syntax sends the data from the pipeline to the
destination dataset specified in the pipeline settings.
*/

| eval _raw=replace(_raw, /<(?P<priority>.*)>(?P<month>[A-z]{3})\s(?P<date>[0-9]{2}?)\s(?P<time>[0-9]+:
[0-9]+:[0-9]+)\s(?P<hostname>(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5]).){3}([0-9]|[1-9][0-9]|1
[0-9]{2}|2[0-4][0-9]|25[0-5]))\s(?P<daemon>[A-z]+)(?P<pid>[[0-9]+])*:\s(?P<body>.*)/, "<\\1> \\2 \\3
\\4 x.x.x.x \\9\\10: \\11")

/* MASKING IP ADDRESSES
In order to mask sensitive information in your data, this template uses the “eval” command to detect
and replace IP addresses in the hostname field with the string “x.x.x.x.”
*/

| into $destination;
```

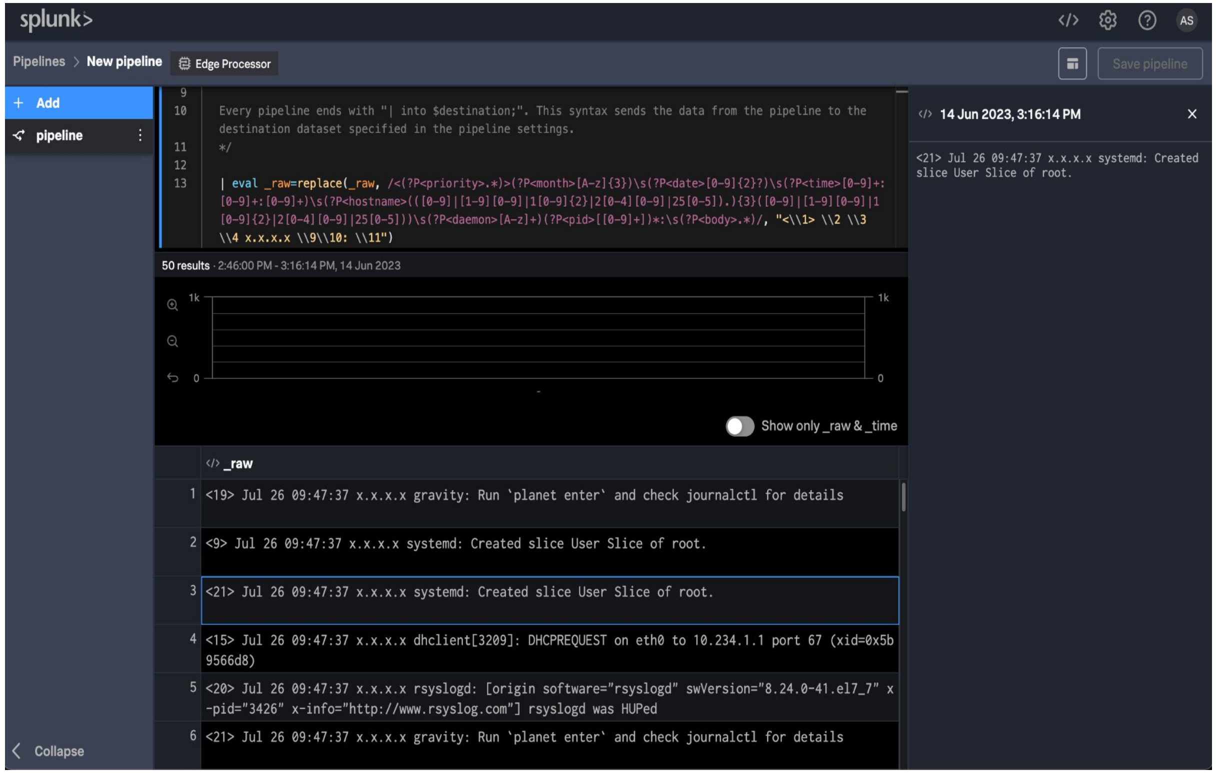

```
Every pipeline ends with "| into $destination;". This syntax sends the data from the pipeline to the
destination dataset specified in the pipeline settings.
*/

| eval _raw=replace(_raw, /<(?P<priority>.*)>(?P<month>[A-z]{3})\s(?P<date>[0-9]{2}?)\s(?P<time>[0-9]+:
[0-9]+:[0-9]+)\s(?P<hostname>(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5]).){3}([0-9]|[1-9][0-9]|1
[0-9]{2}|2[0-4][0-9]|25[0-5]))\s(?P<daemon>[A-z]+)(?P<pid>[[0-9]+])*:\s(?P<body>.*)/, "<\\1> \\2 \\3
\\4 x.x.x.x \\9\\10: \\11")
```

50 results · 2:46:00 PM - 3:16:14 PM, 14 Jun 2023

14 Jun 2023, 3:16:14 PM

```
<21> Jul 26 09:47:37 x.x.x.x systemd: Created
slice User Slice of root.
```

```
1 <19> Jul 26 09:47:37 x.x.x.x gravity: Run `planet enter` and check journalctl for details
2 <9> Jul 26 09:47:37 x.x.x.x systemd: Created slice User Slice of root.
3 <21> Jul 26 09:47:37 x.x.x.x systemd: Created slice User Slice of root.
4 <15> Jul 26 09:47:37 x.x.x.x dhclient[3209]: DHCPREQUEST on eth0 to 10.234.1.1 port 67 (xid=0x5b
  9566d8)
5 <20> Jul 26 09:47:37 x.x.x.x rsyslogd: [origin software="rsyslogd" swVersion="8.24.0-41.el7_7" x
  -pid="3426" x-info="http://www.rsyslog.com"] rsyslogd was HUPed
6 <21> Jul 26 09:47:37 x.x.x.x gravity: Run `planet enter` and check journalctl for details
```

Deployment with Edge Processor

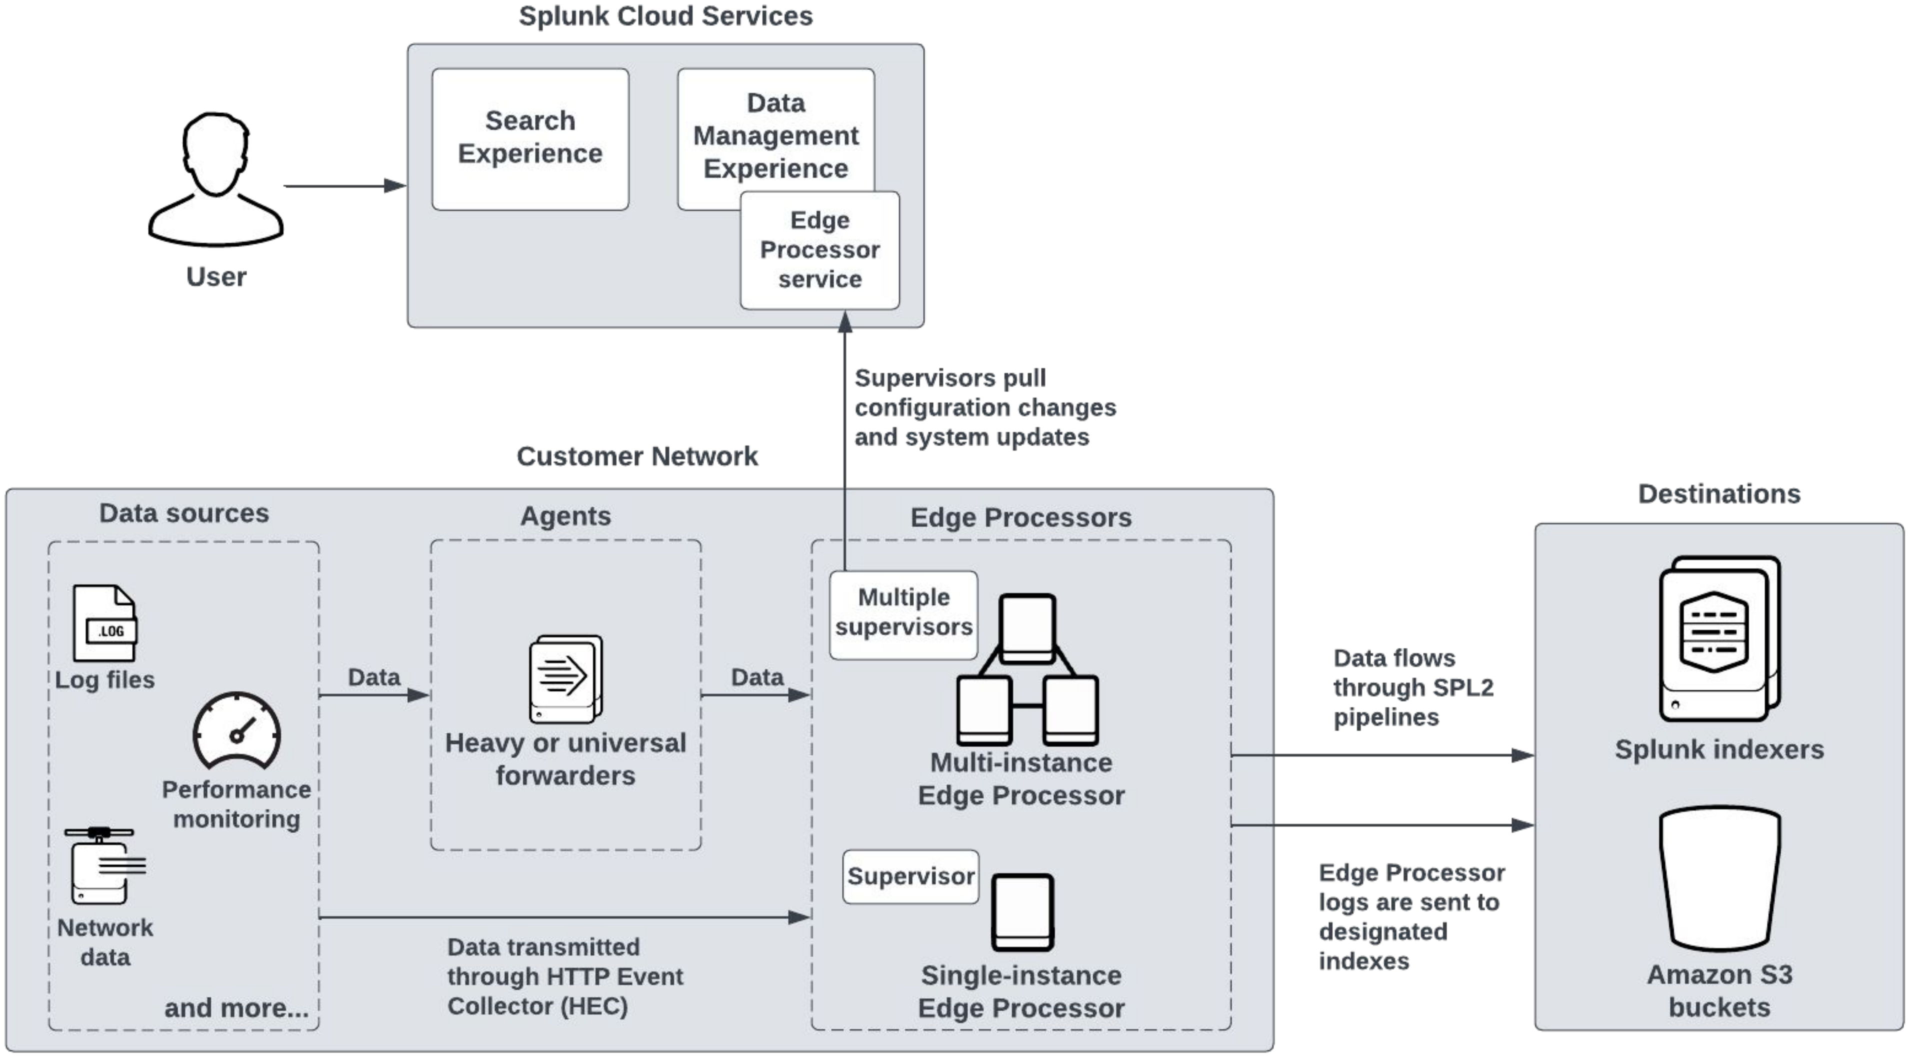
Splunk Cloud Services
Search Experience
Data Management Experience
Edge Processor service
User
Supervisors pull configuration changes and system updates
Customer Network
Data sources
Log files
.LOG
Performance monitoring
Network data
and more...
Data
Agents
Heavy or universal forwarders
Data
Edge Processors
Multiple supervisors
Multi-instance Edge Processor
Supervisor
Single-instance Edge Processor
Data transmitted through HTTP Event Collector (HEC)
Data flows through SPL2 pipelines
Edge Processor logs are sent to designated indexes
Destinations
Splunk indexers
Amazon S3 buckets

Demo - Overview

Questions?

splunk> .conf23

Filtering and Routing with Splunk Edge Processor

Security Engineer:

Filter excessive firewall and load balancer logs, in the event of a rule change or during period of bursts, to reduce storage and compute costs and route only relevant events to Splunk

Recreating Raw Events with Splunk Edge Processor

Application Owner:

Drop unnecessary fields from Windows Application log events, to trim down event size and reduce time to identify relevant information quickly

Masking with Splunk Edge Processor

Security Engineer:

Mask sensitive data, such as PII, to be compliant with legal and regulatory requirements

Use Cases

Filtering

Application Owner: Filter info and debug messages to Splunk,resulting in quicker problem identification

Security Engineer: Filter out excessive noise from firewall and load balancers to reduce Splunk's storage and compute cost

Masking

Security Engineer: Mask sensitive data to be compliant with legal and/or regulatory requirements

Cloud Migration

Splunk Admin: Easy and smooth migration from On-prem to Cloud

Monitoring

Splunk Admin: Single pane of glass to track and monitor all Edge Processor nodes

Questions?

splunk> .conf23

Still on the Fence?

Edge Processor is a great addition to ***<u>any</u>*** Splunk deployment. It allows you to achieve a higher signal-to-noise ratio which yields the following benefits:

1) Lower storage costs of data
2) Lower ingest-time processing costs
3) Lower search-time processing costs
4) Lower query execution times in turn returning faster results
5) Quicker and easier identification of key signals and the ‘needle’ in the haystack

And many more…

Resources

1) Go to **https://px.scs.splunk.com/<your tenant name>/data-management/** to try out Edge Processor (or contact *edgeprocessor@splunk.com* if you'd like it activated in your tenant)

2) Geek out on SPL2™ at **PLA1430A Workshop at 3pm PT** on Tuesday, 18th July

3) Check out the **Edge Processor (Getting Data In) Booth**

4) Edge Processor doc to get started. **Scan here!**

5) Reach out to me - **asaboowala@splunk.com**, if you are interested in learning more about Edge Processor

Thank You