Forward-looking statements

This presentation may contain forward-looking statements that are subject to the safe harbors created under the Securities Act of 1933, as amended, and the Securities Exchange Act of 1934, as amended. All statements other than statements of historical facts are statements that could be deemed forward-looking statements. These statements are based on current expectations, estimates, forecasts, and projections about the industries in which we operate and the beliefs and assumptions of our management based on the information currently available to us. Words such as “expects,” “anticipates,” “targets,” “goals,” “projects,” “intends,” “plans,” “believes,” “momentum,” “seeks,” “estimates,” “continues,” “endeavors,” “strives,” “may,” variations of such words, and similar expressions are intended to identify such forward-looking statements. In addition, any statements that refer to (1) our goals, commitments, and programs; (2) our business plans, initiatives, and objectives; and (3) our assumptions and expectations, including our expectations regarding our financial performance, products, technology, strategy, customers, markets, acquisitions and investments are forward-looking statements. 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. Readers are cautioned that these forward-looking statements are only predictions and are subject to risks, uncertainties, and assumptions that are difficult to predict, including those identified in the “Risk Factors” section of Cisco’s most recent report on Form 10-Q filed on May 21, 2024 and its most recent report on Form 10-K filed on September 7, 2023. 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 Cisco or Splunk, on Cisco or Splunk’s website or otherwise, it may not contain current or accurate information. Cisco and Splunk undertake no obligation to revise or update any forward-looking statements for any reason, except as required by law.

In addition, any information about new products, features, functionality or 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 or be relied upon in making a purchasing decision. We undertake no commitment, promise or 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. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

splunk> .conf24

A Definite Guide to Ingesting Syslog

PLA1726B

.conf24

splunk>

Bring on the future.

A Definite Guide to Ingesting Syslog

Bartlomiej Czarnecki

Senior Manager, Product Management
Splunk

Oskar Patnaik

Senior Product Manager
Splunk

splunk> .conf24

Table of Contents

About Syslog

Why should we care about
something that is 40 years old?

Our Relationship with Syslog

We Love Syslog for…

- Ubiquity
- Simplicity
- Efficiency
- Real-time logging

We Hate Syslog for…

- Limited security
- Reliability
- Lack of structure
- Scalability of Syslog ingestion

Syslog in Numbers

60%

Ubiquity

There are Splunk customers that get more than 60% of all their data through Syslog

250+

Scale

The largest deployments of Splunk Connect for Syslog have more than 250 instances

30%

Sourcetyping

Around 30% of Splunk deployments have some data indexed with sourcetype=syslog

Internal Splunk data.

How can you get the most out of Syslog?

Get off Syslog Protocol

- Collect Syslog messages as close to source as possible, to limit amount of networking between syslog client and server
- Use UDP for performance, TCP for reliability
- Do not overenginner, there will always be some data loss

Assign Sourcetype

- Choose the simplest sourcetyping strategy:
  - Port
  - IP
  - Host
  - Header/Content
- Splunk Connect for Syslog has hundreds of out-of-the-box filters

Filter/Route

- Send Syslog to Edge Processor HEC input
- Use rulesets to filter and route data

Splunk Connect for Syslog (SC4S)

Our recommended way to collect and sourcetype Syslog messages

About Splunk Connect for Syslog

Resilient

- Built on top of syslog-ng
- Containerised
- Disk buffer
- Defaults and fall-baks
- Monitoring dashboard
- Various deployment methods supported, including BYOE

Scalable

- Performance tests are part of CI/CD pipelines
- Vertical scaling up to 6TB/day on m5zn.2xlarge
- Horizontal scaling only with microK8s m5zn.2xlarge

Ready to Use

- Hundreds of filters for most common data sources
- Sourcetyping and other required metadata
- Defaults and fallbacks
- Works with TAs and CIM

Agile

- Open source software
- Main contributors is dedicated Splunk eng team
- Healthy cadence of releases - usually every 2 weeks
- Large and engaged community

SC4S Architecture

Syslog Devices

SC4S Architecture

Nothing

Syslog Devices

SC4S Architecture

SC4S on mK8s

Nothing

Syslog Devices

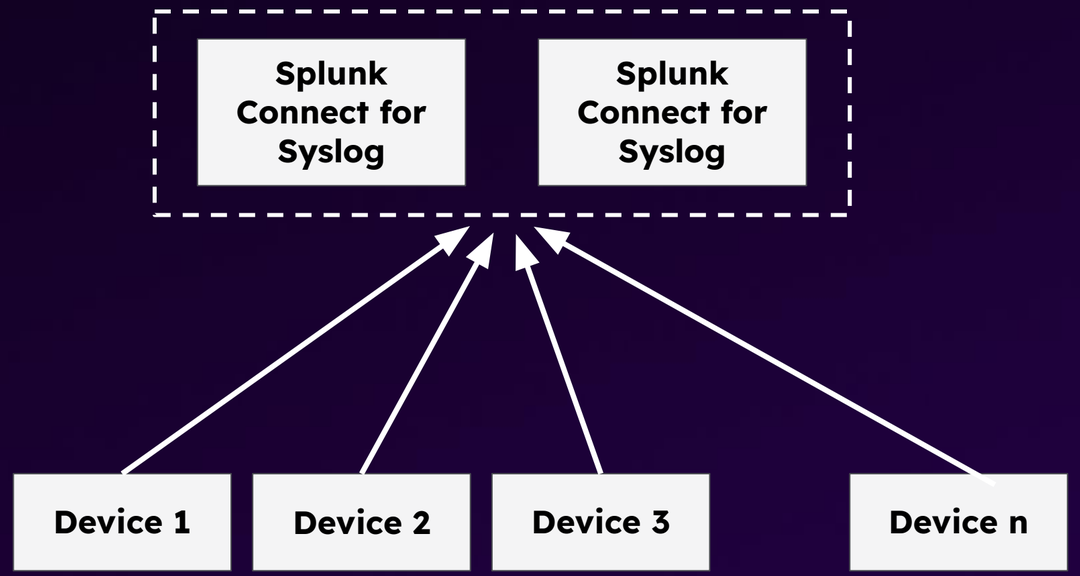

SC4S Architecture

Anything

SC4S on mK8s

Nothing

Syslog Devices

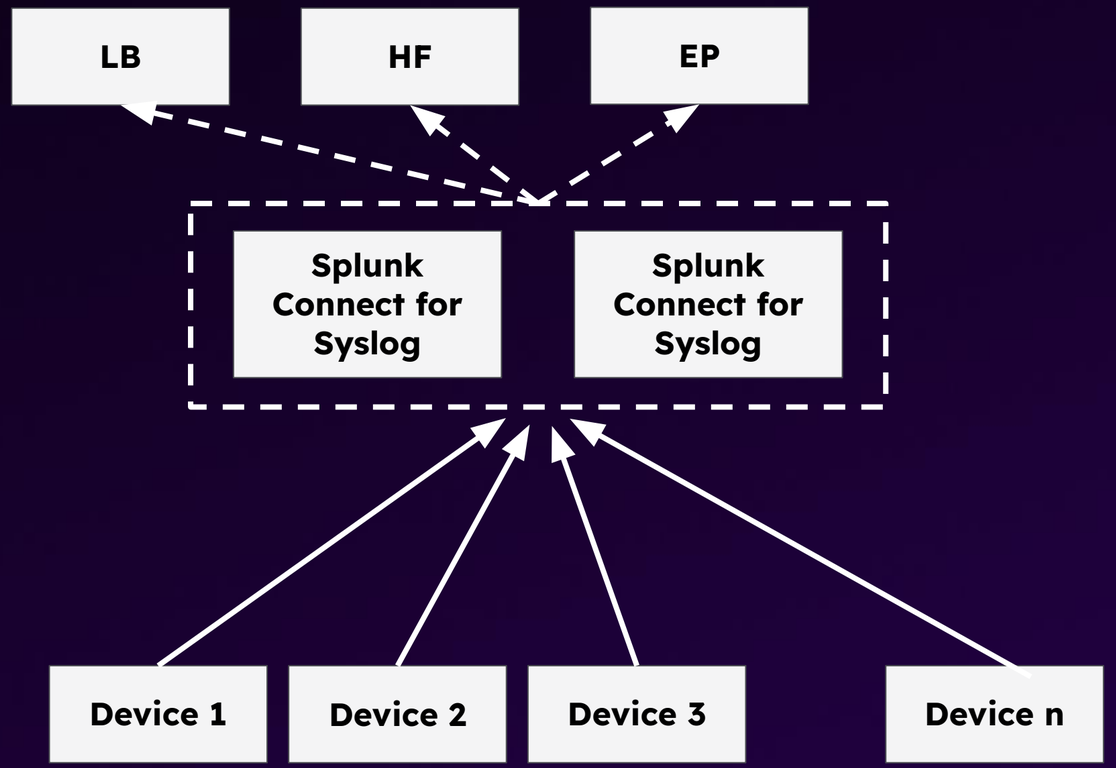

SC4S Architecture

HEC Input

Anything

SC4S on mK8s

Nothing

Syslog Devices

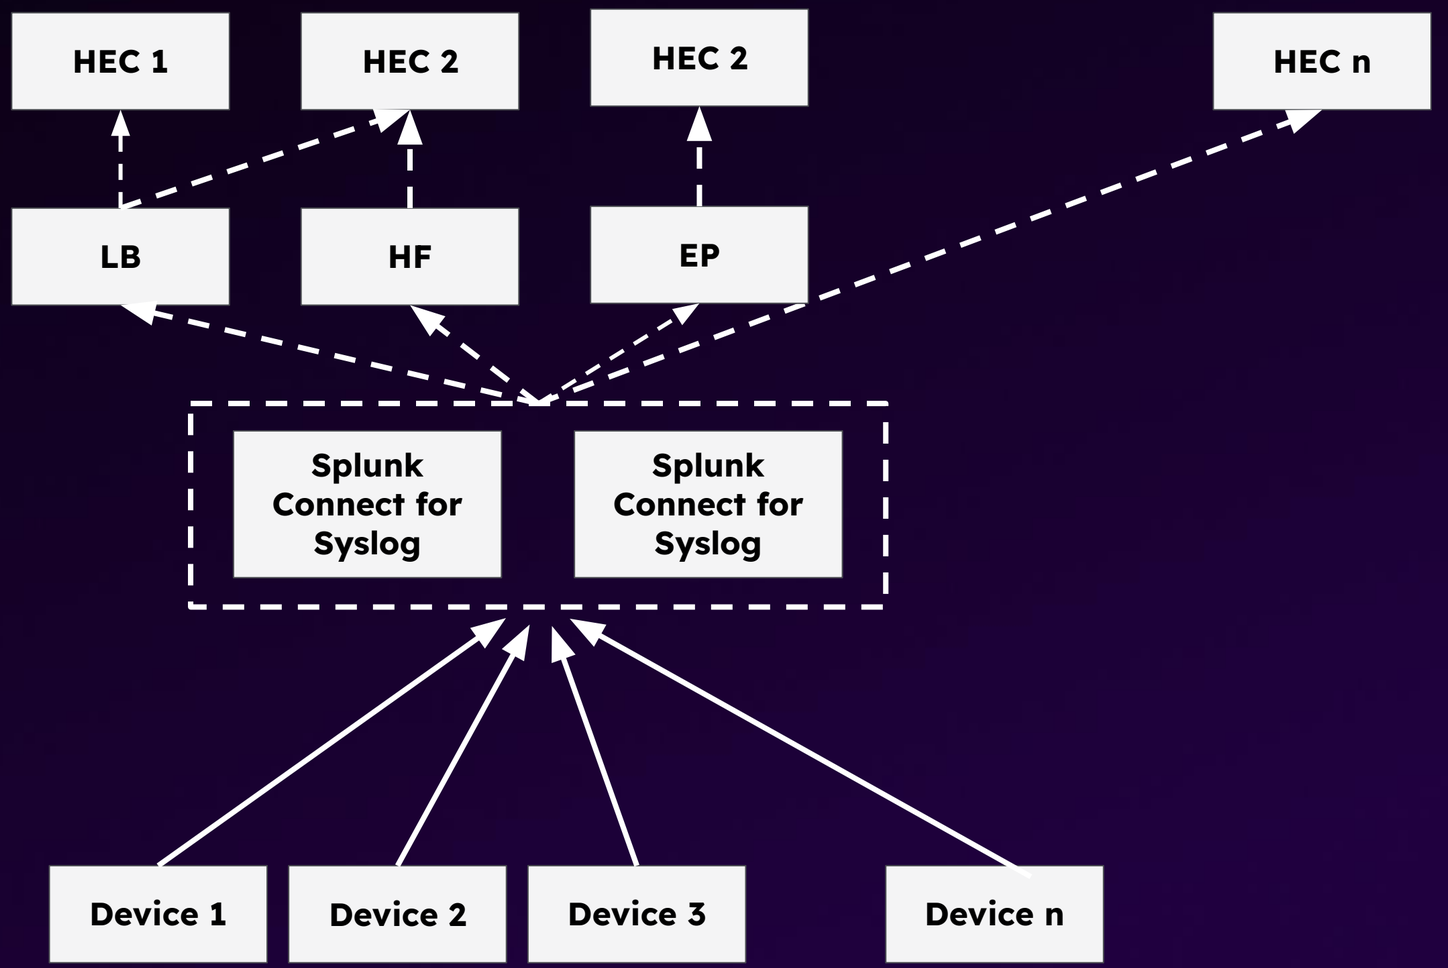

Optimize Syslog Data Ingestion

Plug SC4S into Splunk Edge Processor

Enable more use cases with Splunk® Edge Processor

Filtering

- Filter info and debug messages to Splunk, resulting in quicker problem identification
- Filter out excessive noise from firewall and load balancers to reduce Splunk's storage and compute cost

Masking

- Mask sensitive data to be compliant with legal and/or regulatory requirements

Routing

- Route data to a destinations of choice (Splunk, AWS S3, Apache Kafka, etc)

Architecture Overview

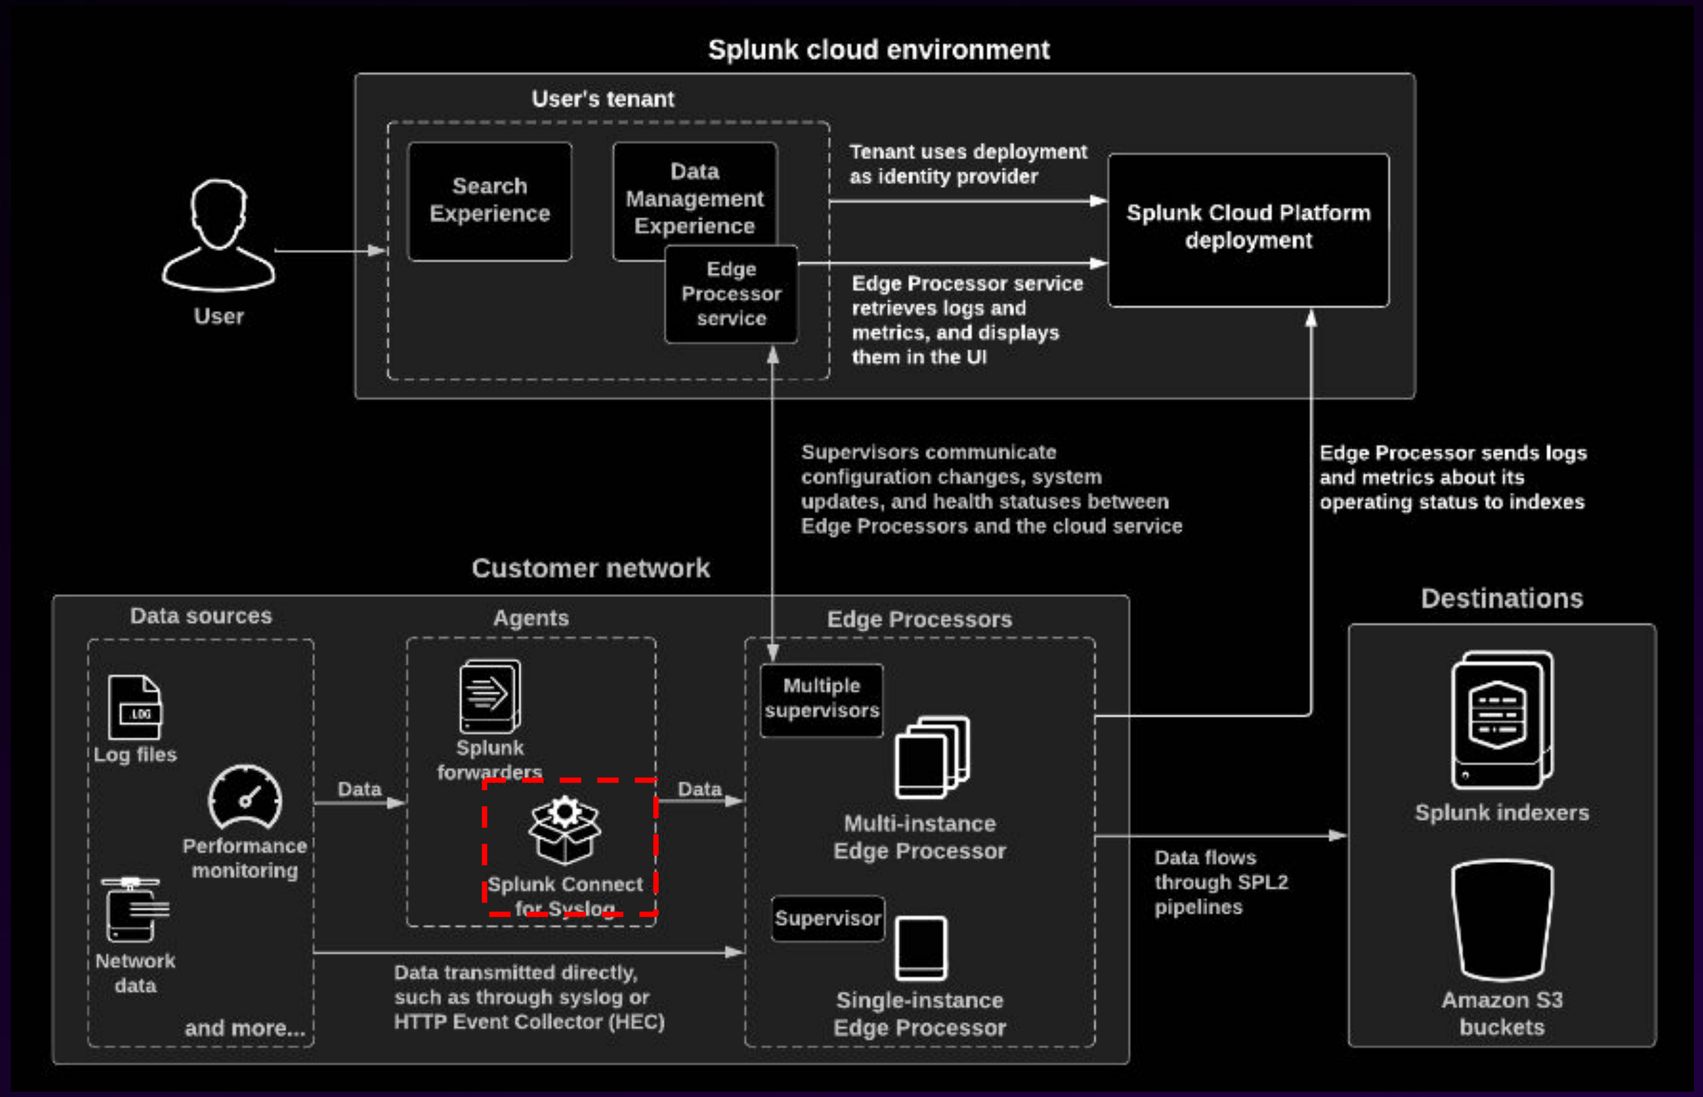

Architecture Overview

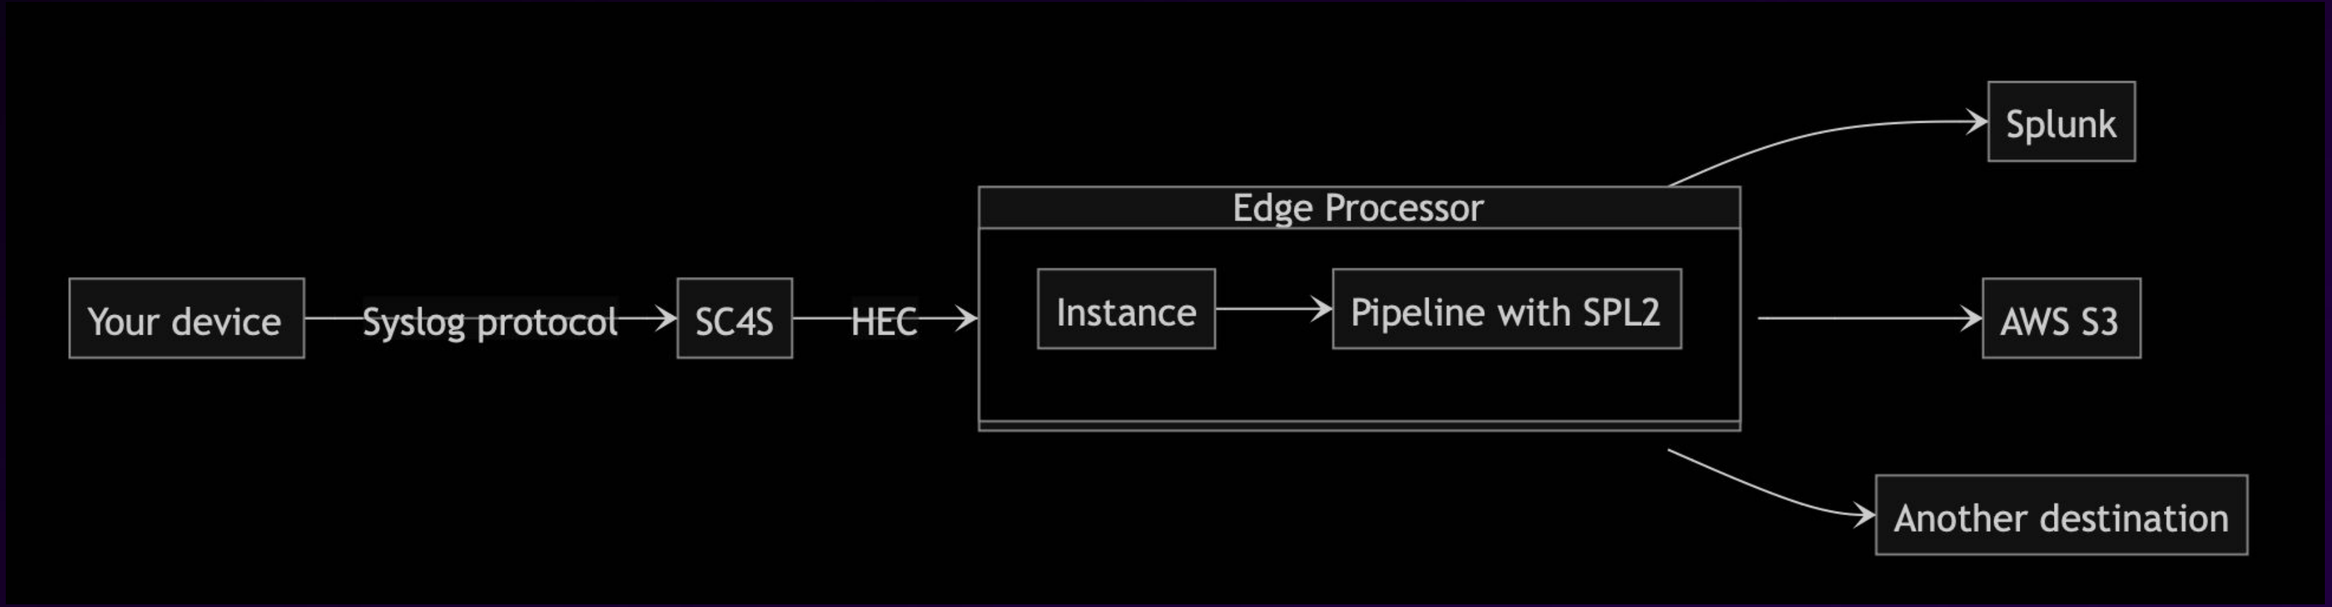

Who is Edge Processor for?

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

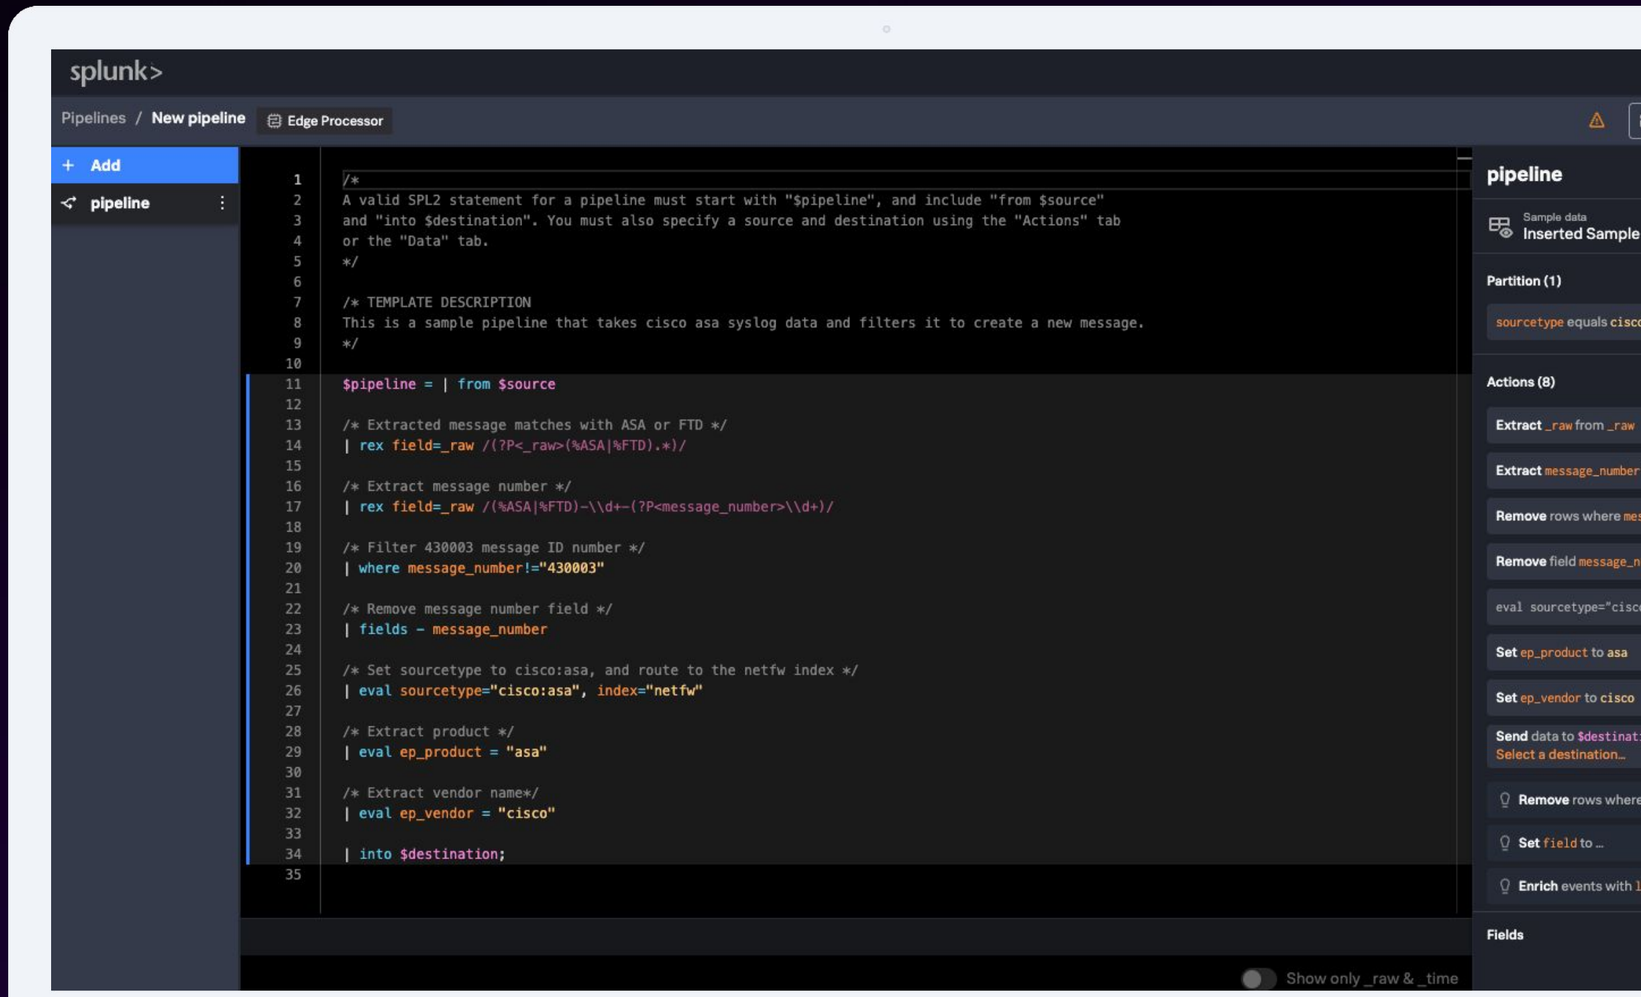

Benefits of Using EP for Syslog Traffic

Use case example

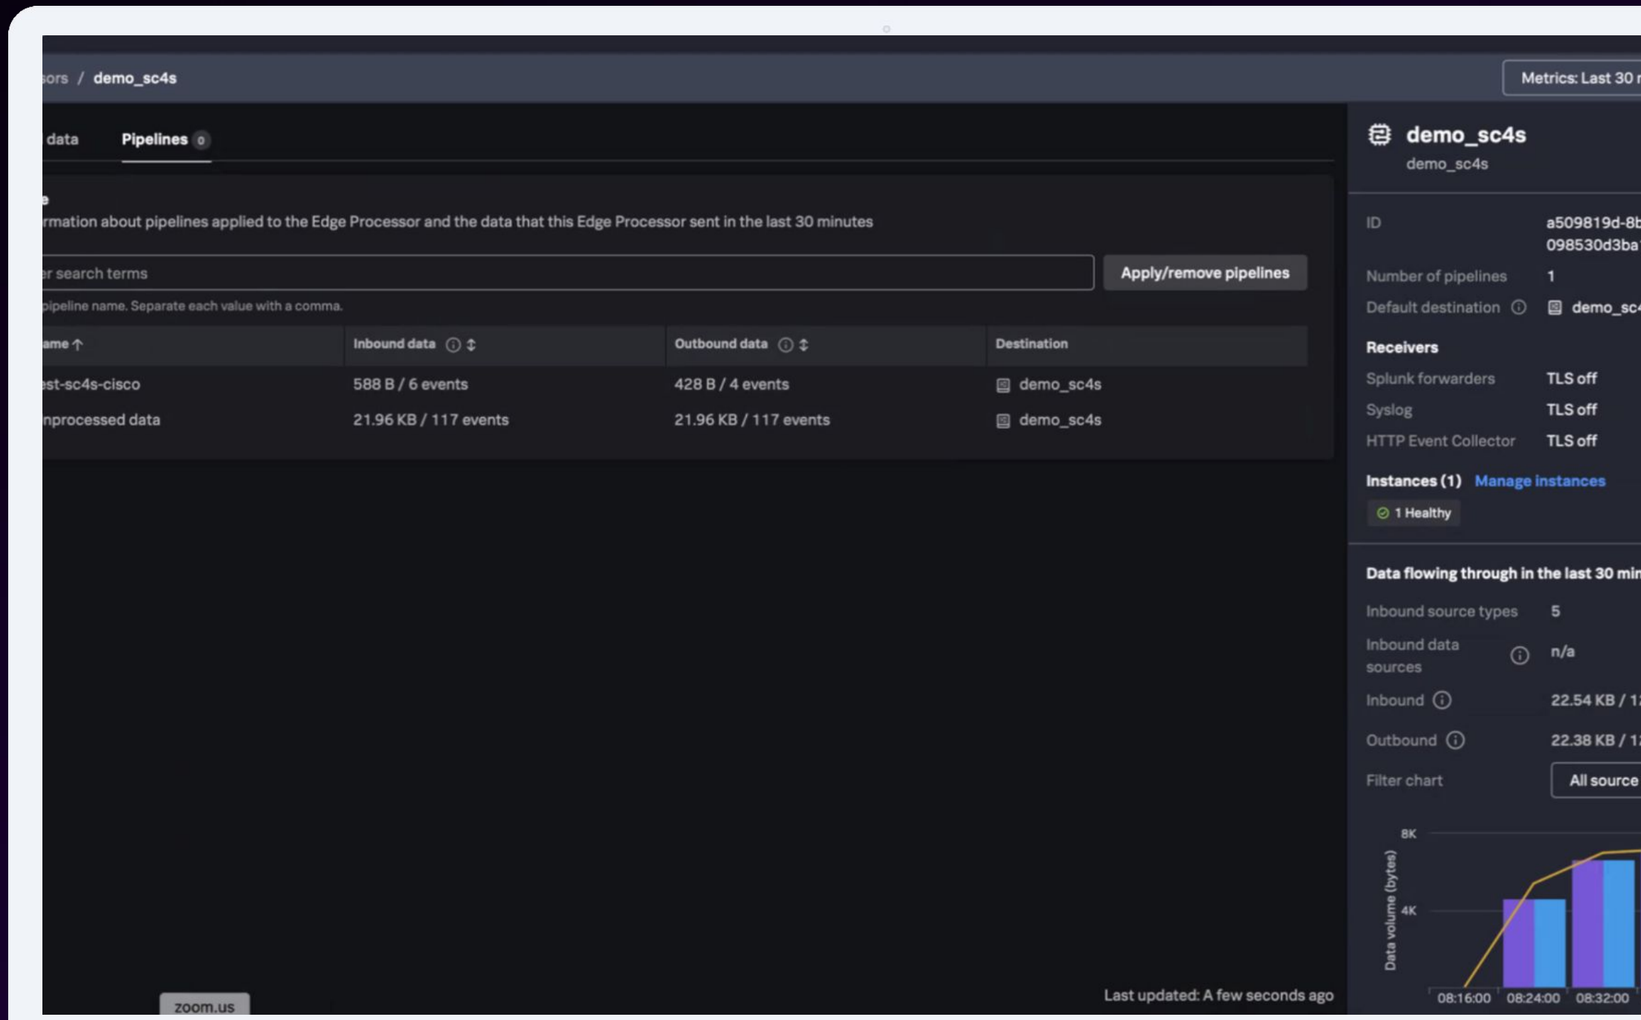

Benefits of Using EP for Syslog Traffic

Step-by-step

Next Step

One simple thing you should do now

Check for sourcetype=syslog in Your Splunk

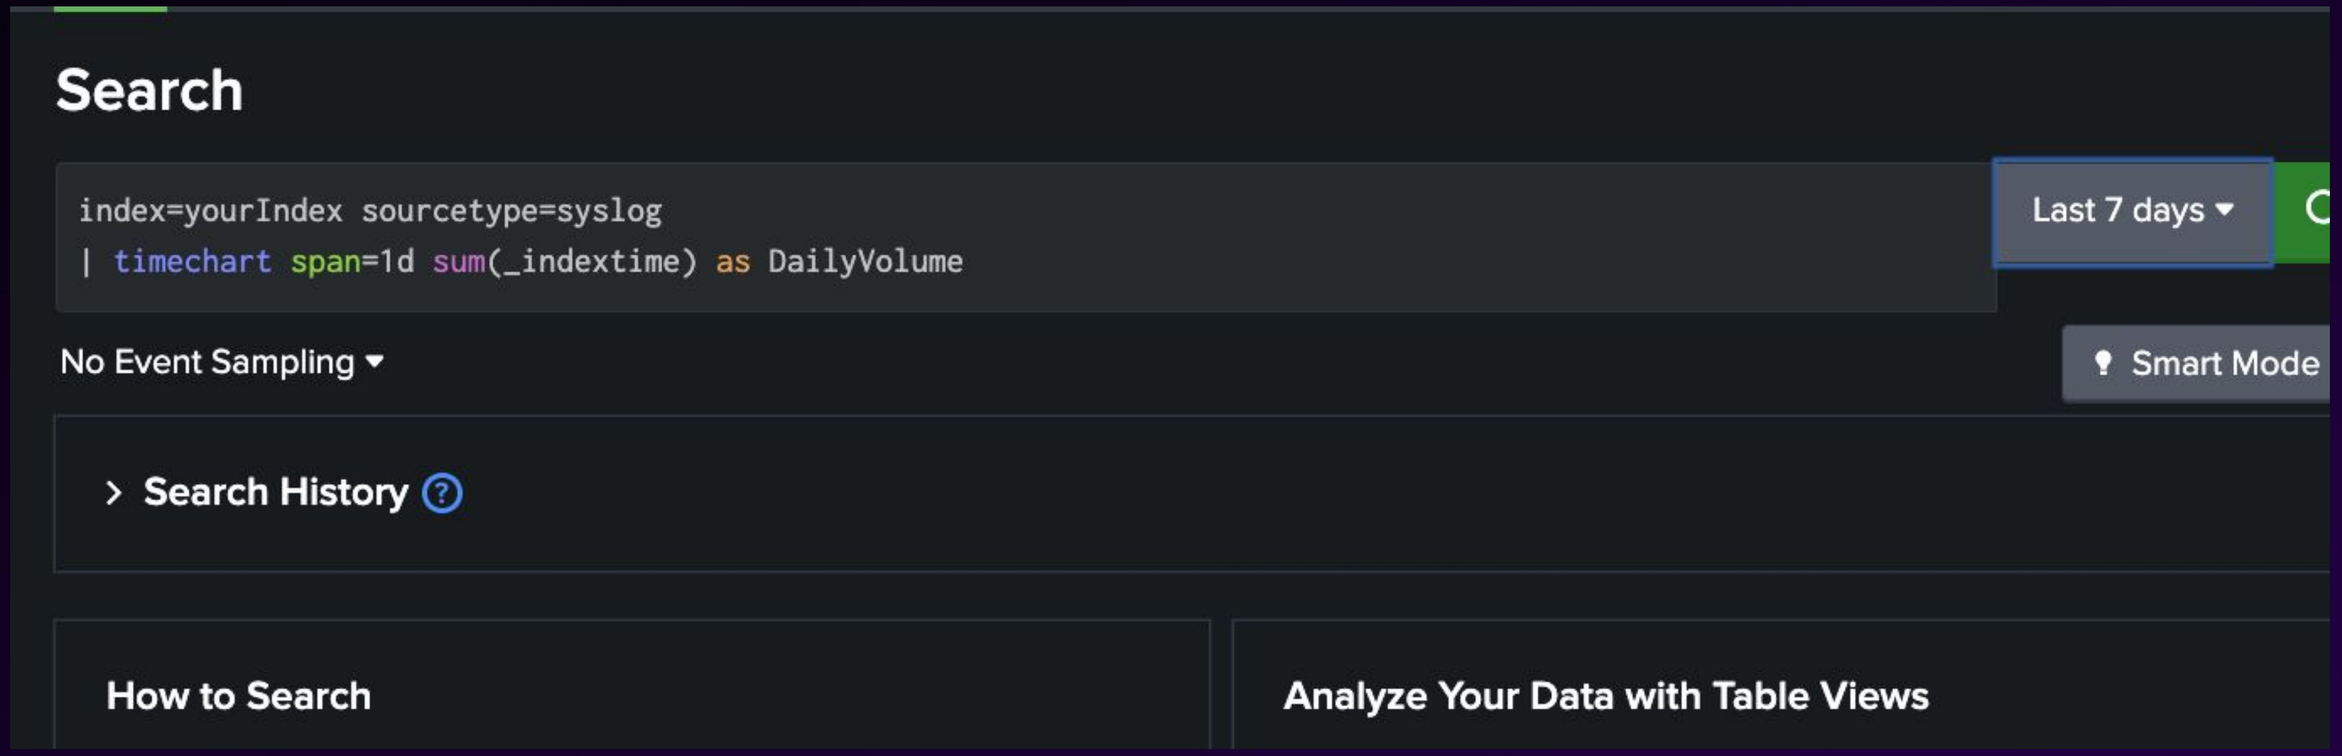

Check for sourcetype=syslog in Your Splunk

Search

```
index=yourIndex sourcetype=syslog
| timechart span=1d sum(_indextime) as DailyVolume
```

Thank you