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

Splunk Universal Forwarder (UF), the OpenTelemetry™ Add-on and You:

Let's Get All the Data!

PLA1117B

.conf24

splunk>

Bring on the future.

Meet The Presenters

Jonathan Fair

TS&I Observability Architect
Splunk

Jason Riley

Observability Professional Services Architect
Splunk

https://splk.it/pla1117b

Journey to the Bottom of the Ocean

Images sourced by Jonathan Fair

Nothing prepares you for a shark!

(even a mostly harmless reef shark such as this)

Images sourced by Jonathan Fair

Just keep swimming

Images sourced by Jonathan Fair

Journey to the Bottom of your Data

Splunk + Observability

Images sourced by Jonathan Fair

Why Splunk® Observability Cloud?

How you should think about it

Core Splunk + Observability Cloud

=

Splunk Observability Cloud

Splunk Observability Cloud

- Real Time Streaming Metrics
- AI driven directed troubleshooting
- No Sample tracing
- End to End Visibility
- Open Standards Open Source

What is OpenTelemetry™?

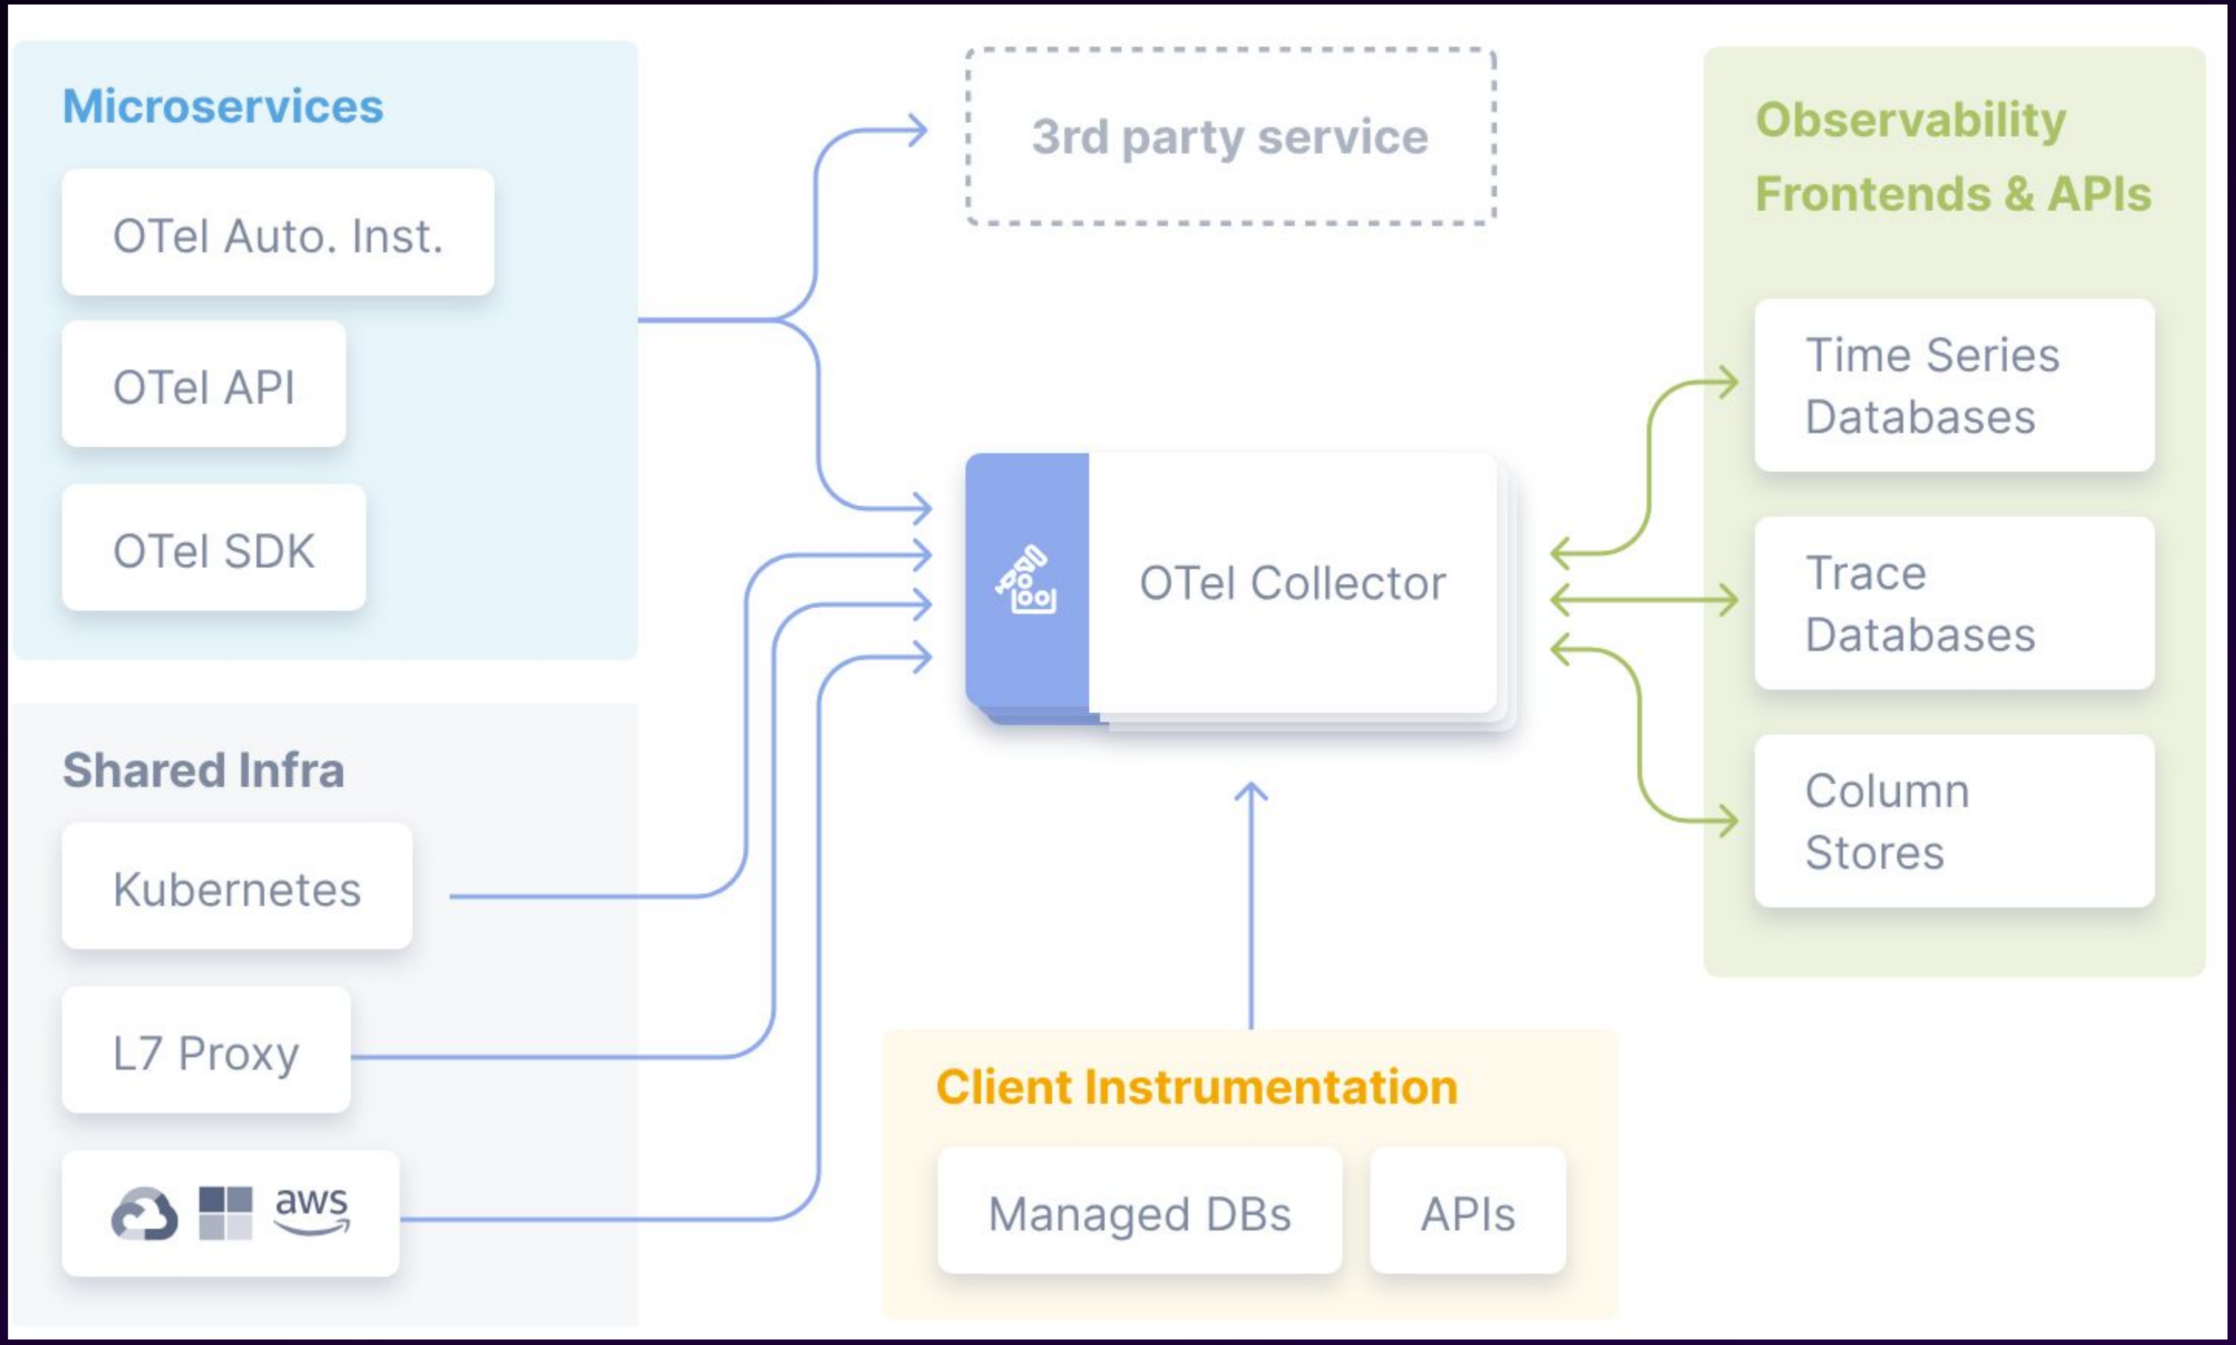

ttps://opentelemetry.io/docs/

Splunk UF Add-on for OTel

UF

OTel

Table of Contents

What Is The OTel Add-on

Basic Setup And Deployment

Lab Exercise 1

Deployment Deep Dive

Best Practices And Considerations

Lab Exercise 2

Configuration Deeper Dive

Lab Exercise 3

Recap And Q&A

What Is The Splunk Add-on for OpenTelemetry Collector?

A Splunk Technical Add-on, that packages the Splunk OTel Collector as a component of the Splunk Universal Forwarder

- Discoverable and accessible on Splunkbase
- Enables collection of metrics and traces for Splunk Observability Cloud

Why did we develop it?

- Enhance accessibility of Splunk OTel Collector
- Simplify deployment via familiar frameworks
- Ease integration for customers with existing UF deployments
- Target audience: Customers with an established fleet of UFs
  - Even better if they are managed via the Splunk Deployment Server

How To Deploy

- Follows the same installation pattern as other Splunk TAs
- Utilize the Deployment Server for streamlined distribution and management
- Possible to use custom deployment frameworks

Time to try it yourselves

Images sourced by Jonathan Fair

Lab Exercise 1:

Basic deployment

- Connect to the Lab Environment
- Review Deployment Server setup
- Deploy pre-configured Splunk TA for OTel Collector to Linux server
- Validate successful deployment and data ingestion into Splunk O11y Cloud

https://splk.it/pla1117b

End of Lab 1

You should have a setup similar to this on your Splunk Deployment Server

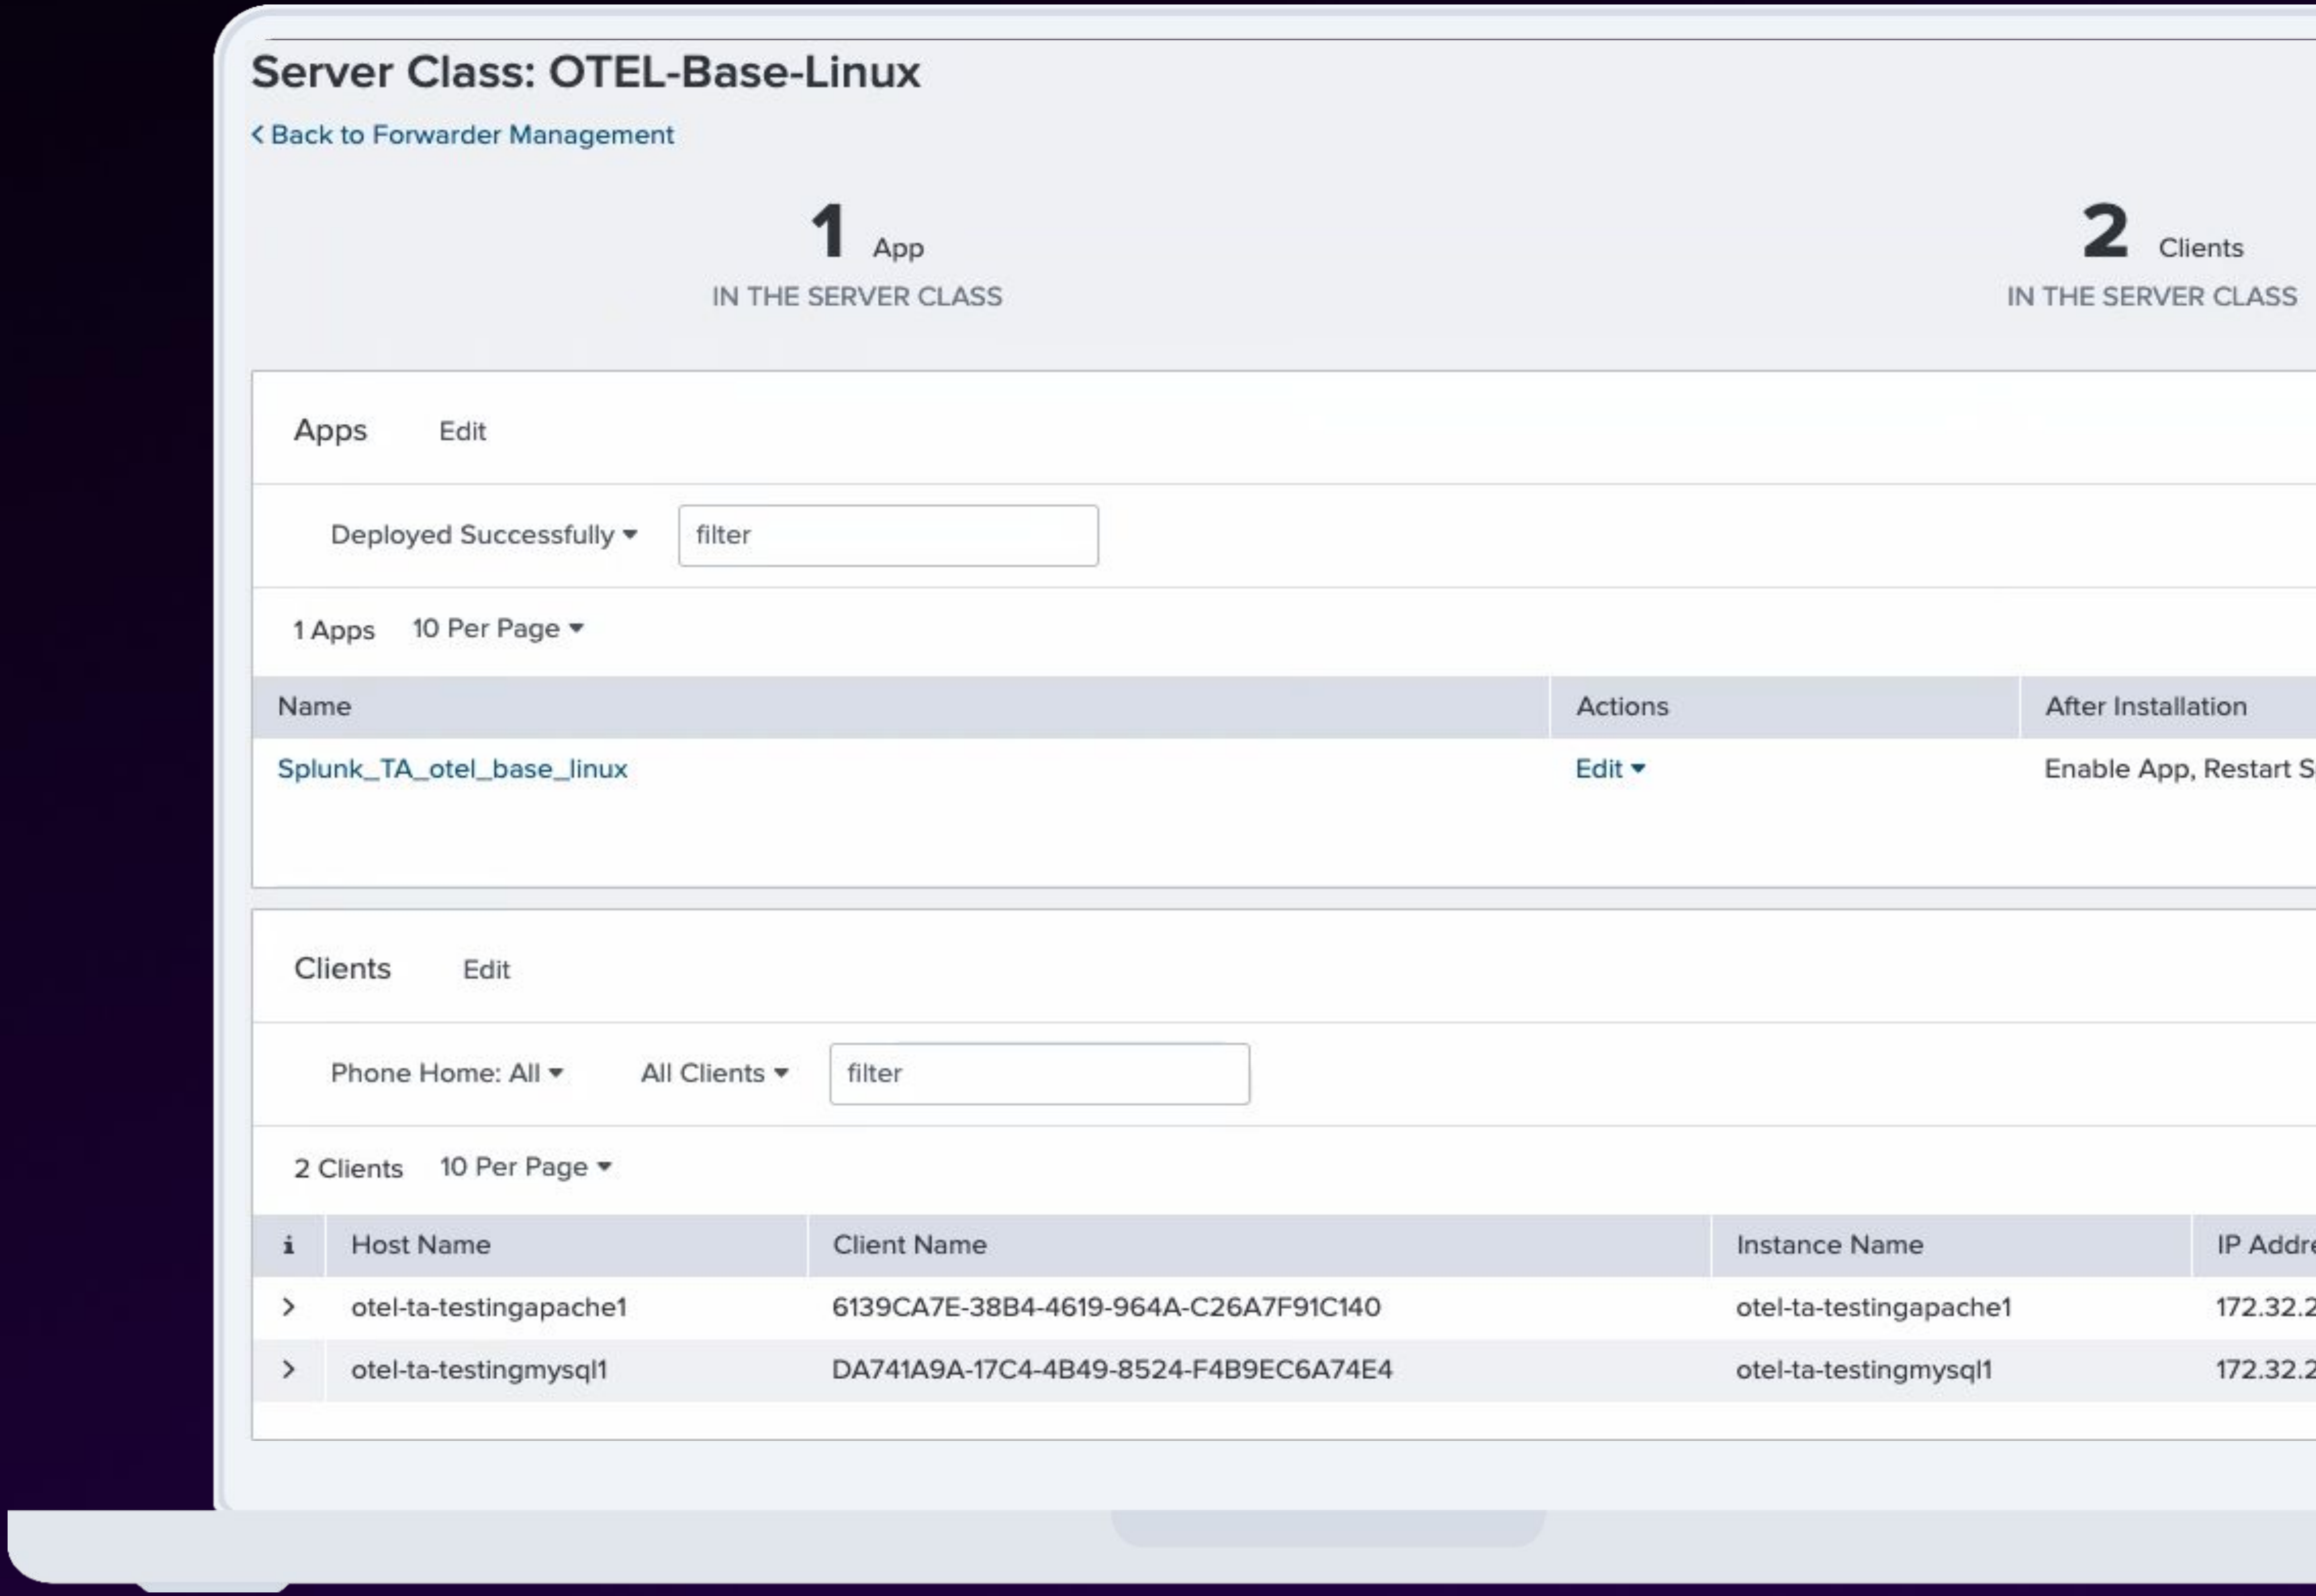

Server Class: OTEL-Base-Linux

< Back to Forwarder Management

1 App IN THE SERVER CLASS

2 Clients IN THE SERVER CLASS

Apps Edit

Deployed Successfully ▾ filter

1 Apps 10 Per Page ▾

Name	Actions	After Installation
Splunk_TA_otel_base_linux	Edit ▾	Enable App, Restart S

Clients Edit

Phone Home: All ▾ All Clients ▾ filter

2 Clients 10 Per Page ▾

i	Host Name	Client Name	Instance Name	IP Addre
>	otel-ta-testingapache1	6139CA7E-38B4-4619-964A-C26A7F91C140	otel-ta-testingapache1	172.32.2
>	otel-ta-testingmysql1	DA741A9A-17C4-4B49-8524-F4B9EC6A74E4	otel-ta-testingmysql1	172.32.2

End of Lab 1

And your Linux server data should be visible in your Splunk Observability Cloud environment.

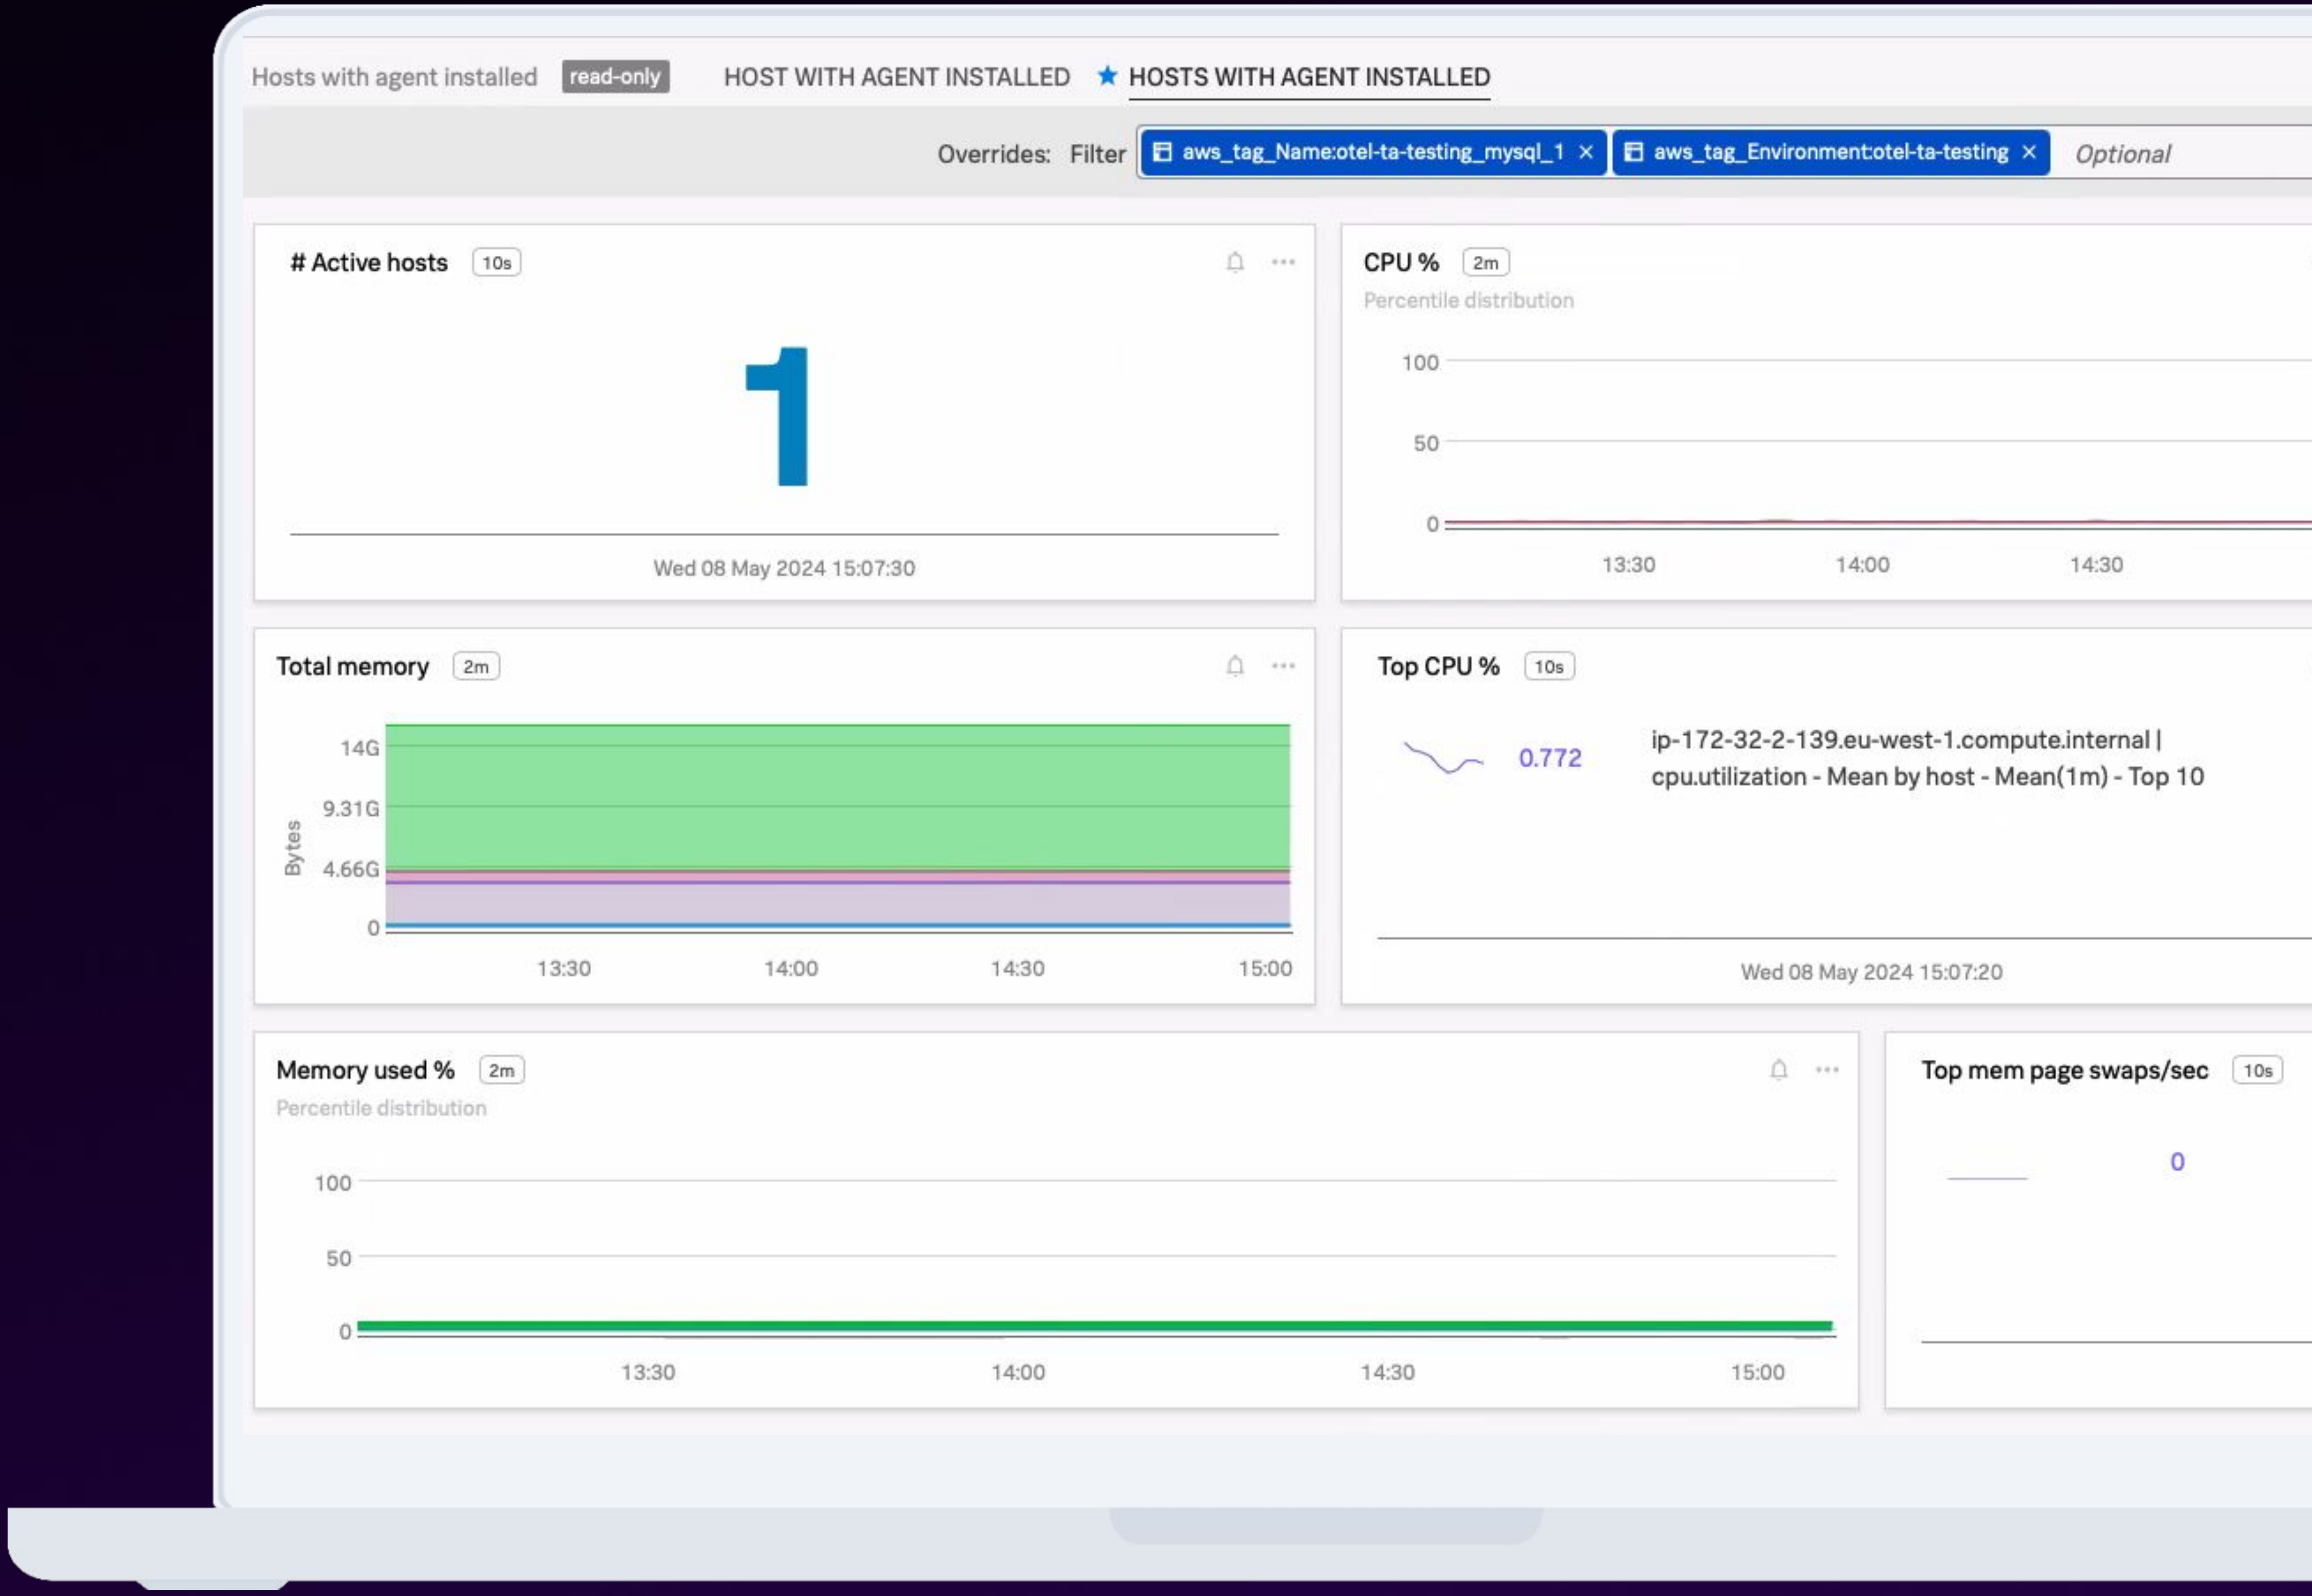

Overrides: Filter aws_tag_Name:otel-ta-testing_mysql_1 × aws_tag_Environment:otel-ta-testing × Optional

Deployment Basics

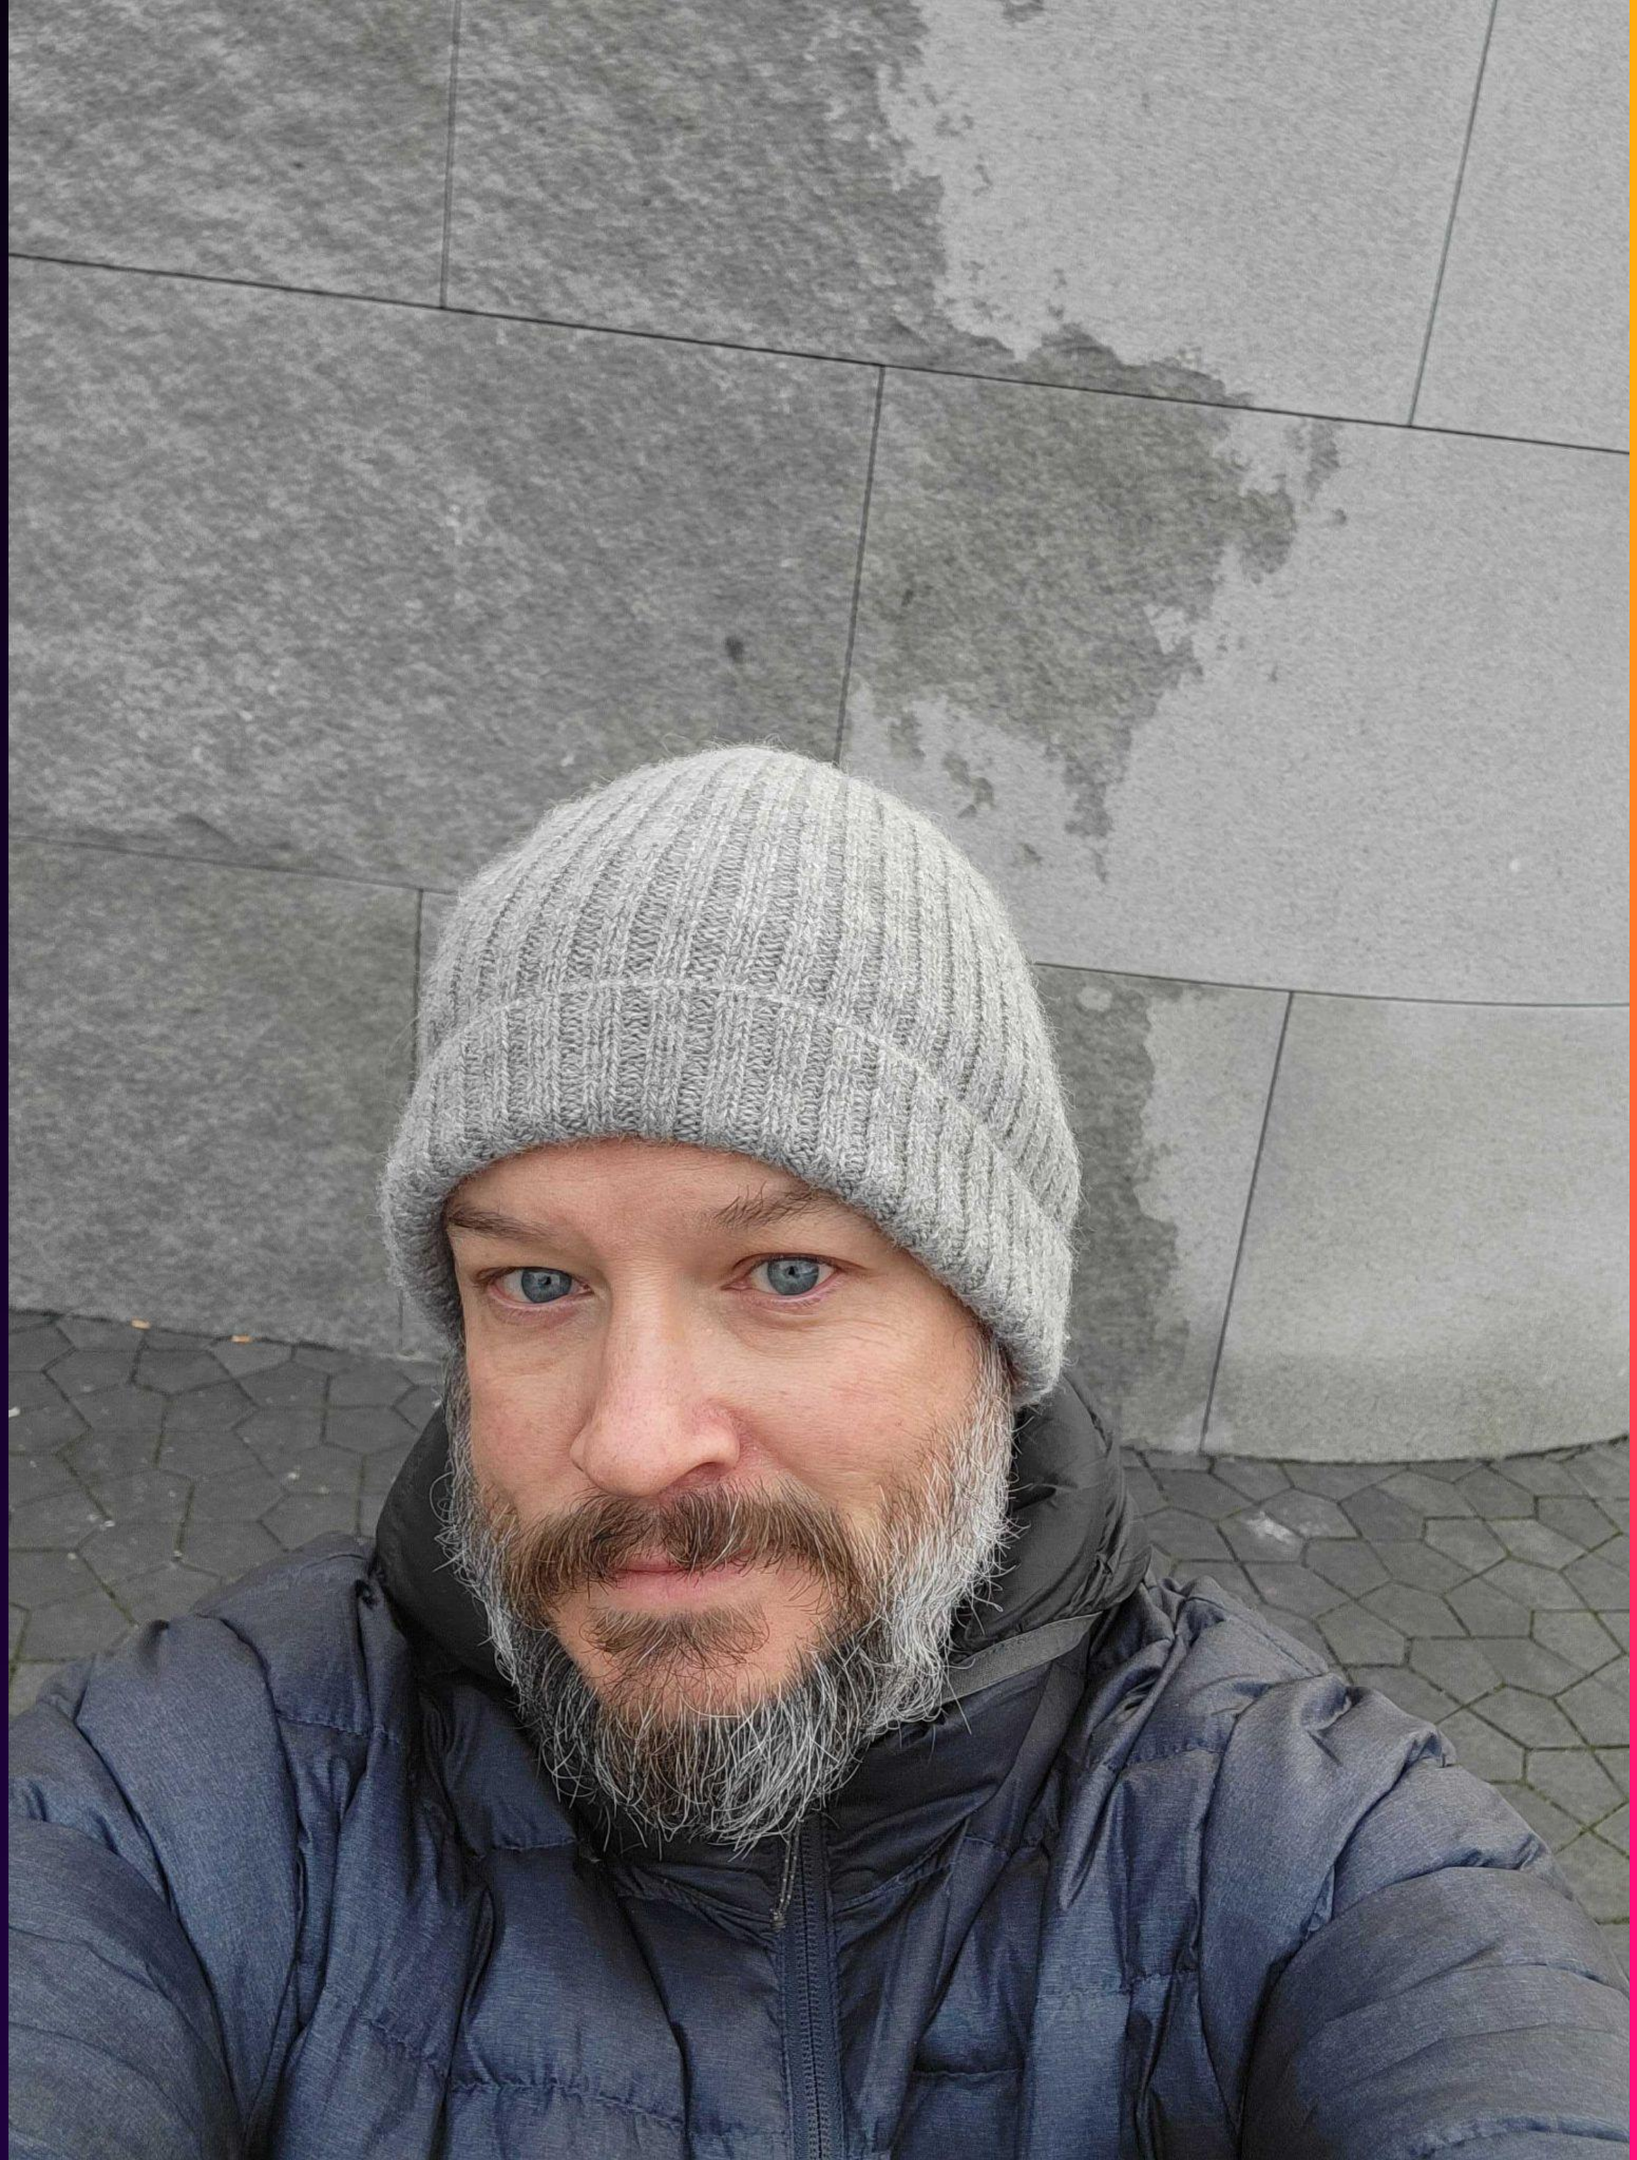

Images sourced by Jason Riley

Deployment Server 101

- Centrally manages configuration packages as apps for clients
- Includes Forwarder Management as a graphical interface
- Capability to restart remote Splunk instances
- Requires an Enterprise license and dedicated server

Deployment Server Components

- Deployment Apps:
  - Packaged config files like `inputs.conf`, residing in `SPLUNK_HOME/etc/deployment-apps/`
- Deployment Clients:
  - Splunk instances phoning home to DS, initiate connections
- Server Classes:
  - Group deployment clients, define app deployment via serverclass.conf

How To Deploy

1. Server Classes and Add-on Packages
2. Configure instances as deployment clients - clients phone home to DS
3. Client downloads subscribed apps as directed by Server Classes on DS
4. Client uses app configurations

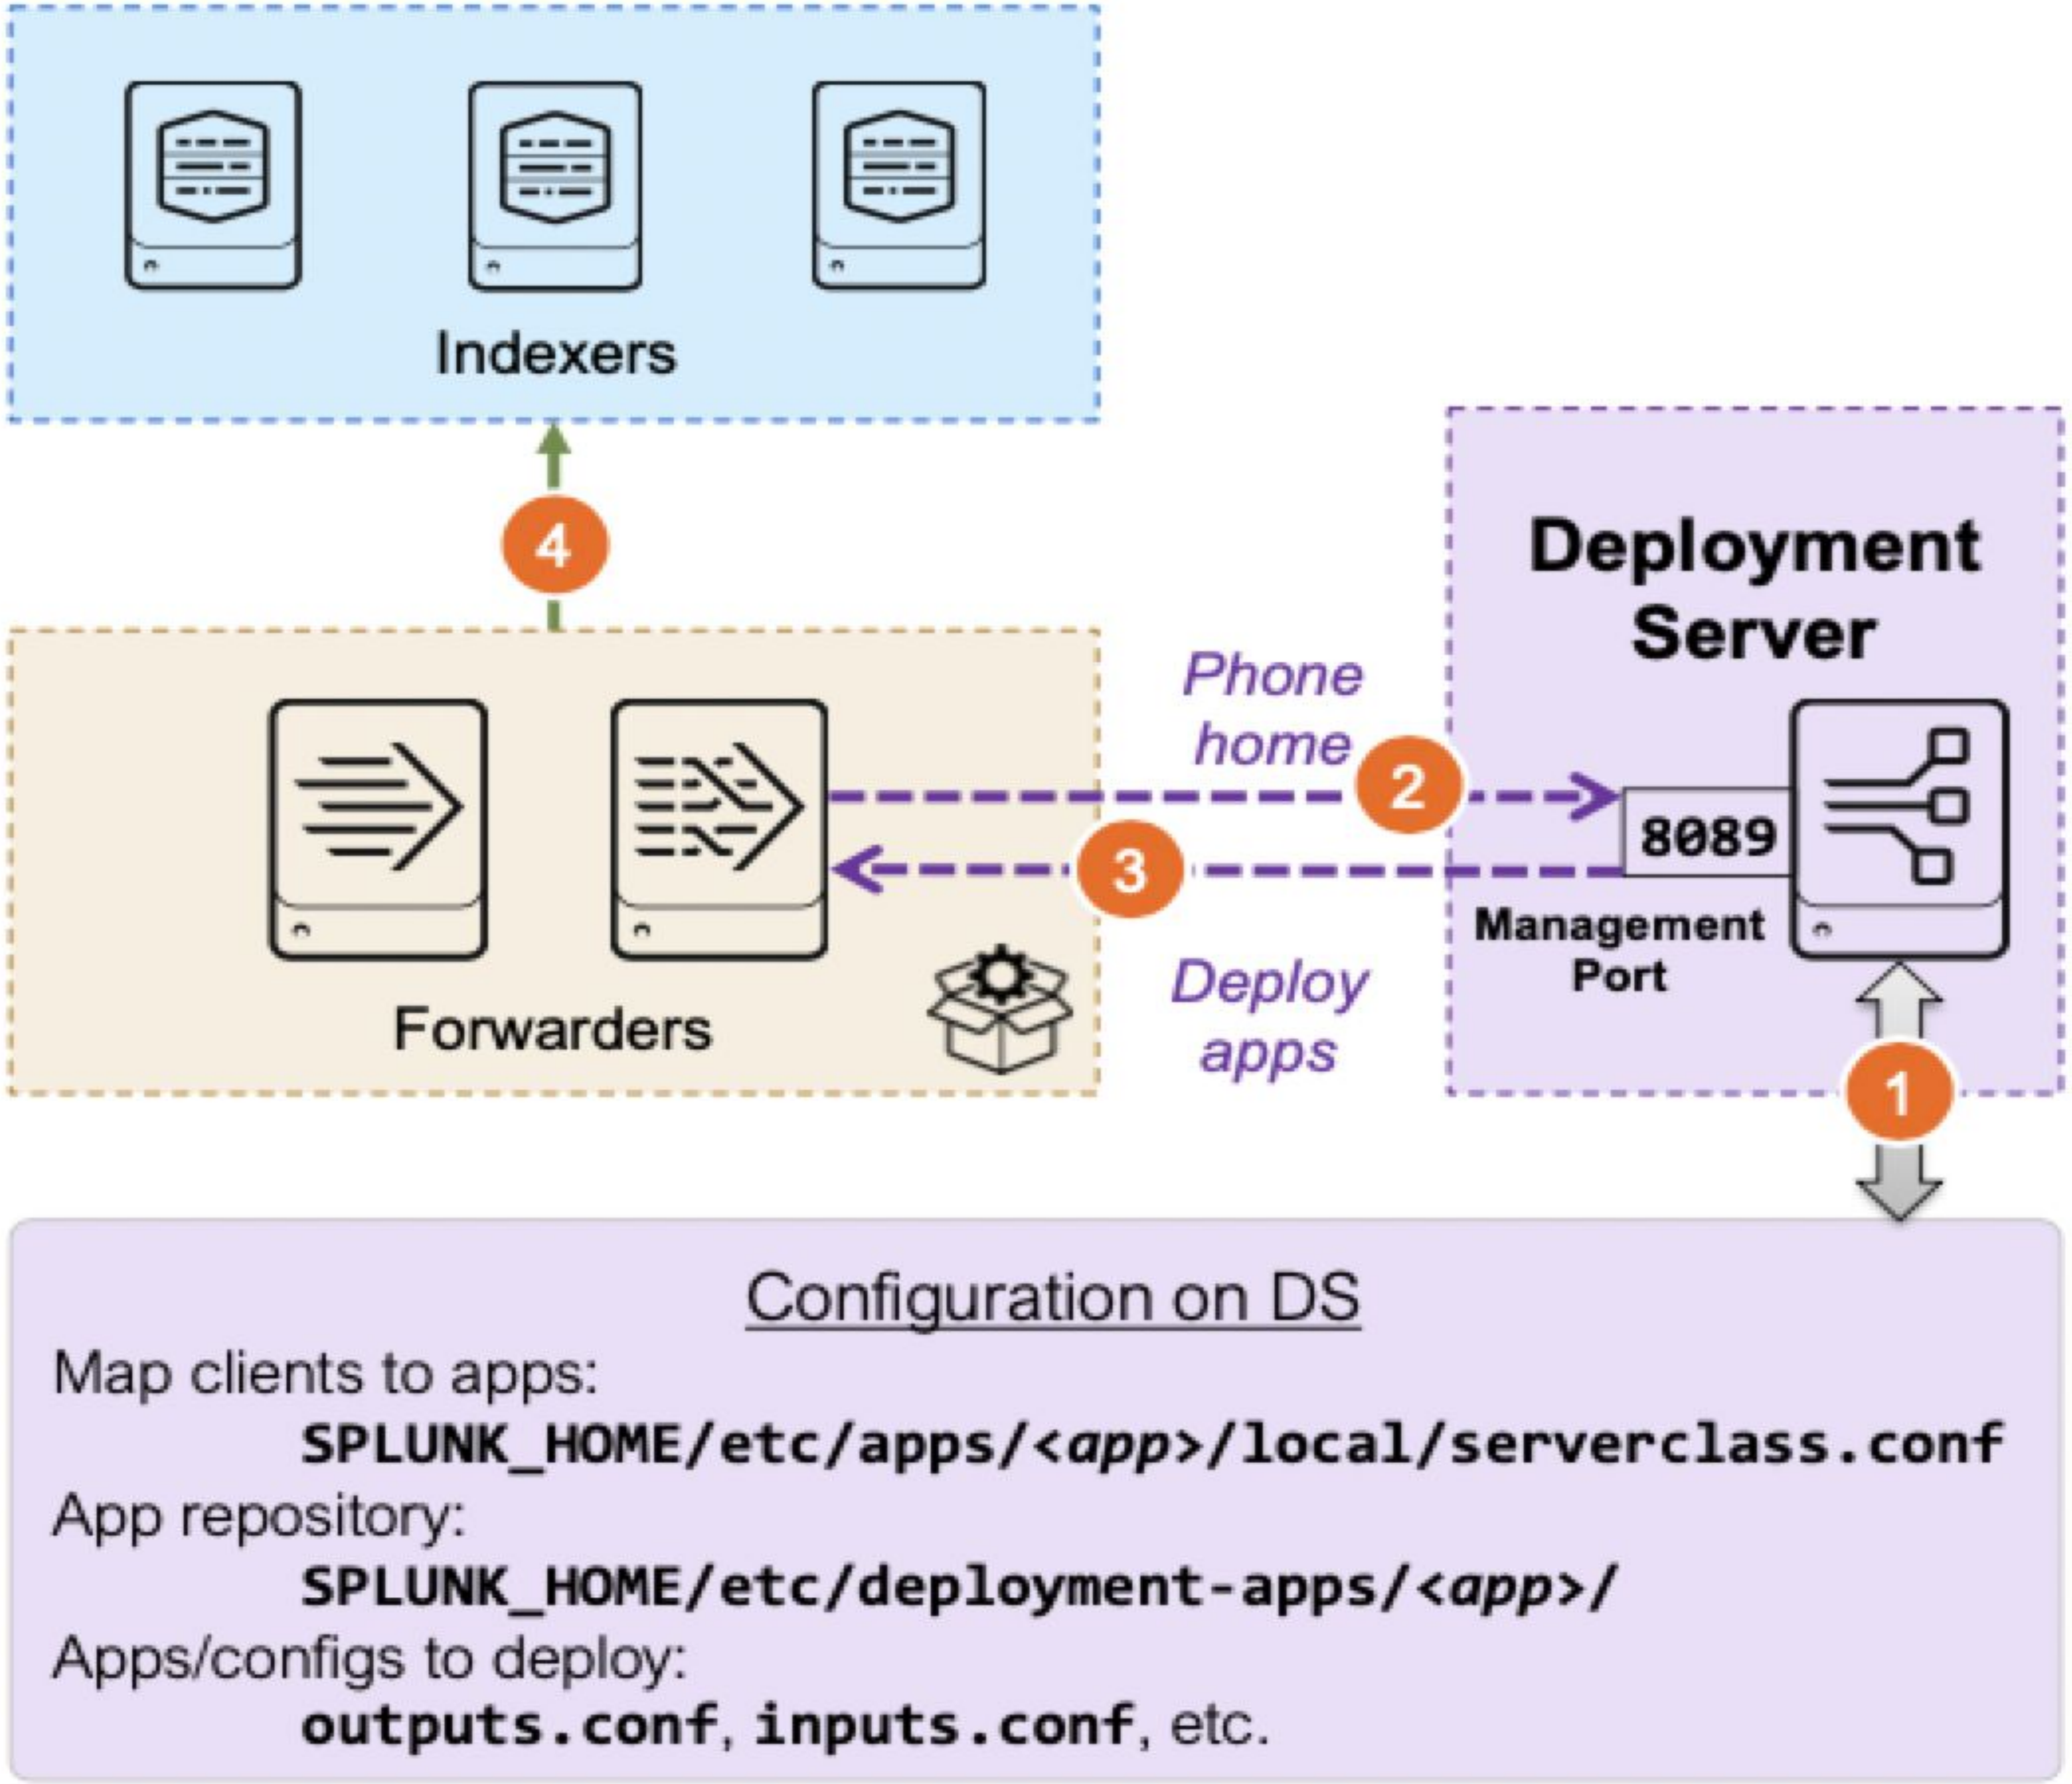

Configuration on DS

Map clients to apps:
`SPLUNK_HOME/etc/apps/<app>/local/serverclass.conf`

App repository:
`SPLUNK_HOME/etc/deployment-apps/<app>/`

Apps/configs to deploy:
`outputs.conf`, **`inputs.conf`**, etc.

Things to Consider

- Differences in deployment method compared to standalone OTEL collector
- Lack of automatic discovery and configuration support
- Log collection disabled in Collector TA
- Large package size due to binaries
- OTEL TA runs as a child process of UF, adhering to the same user privileges

Additional Considerations

- Performance on par with standalone OTEL collector post-installation
  - See sizing recommendations
- Core version compatibility: Tested on UF versions 8.x and 9.x
- Scalability limits of the Deployment Server apply

Configuration Basics

Photo by Jonathan Fair

Images sourced by Jonathan Fair

Basic Configuration of the Splunk Add-on for OTel Collector

- New attributes
  - inputs.conf: Sets the Splunk Observability Cloud realm (default us0), metric and trace endpoints
  - Access Token: Required for Splunk Observability Cloud, specified in a separate file referenced in inputs.conf
  - Packaged with the default configuration of the Splunk distribution of the Collector

Add-on Folder Structure

```
README
   inputs.conf.spec
configs
   ta-agent-config.yaml
default
   access_token
   app.conf
   inputs.conf
linux_x86_64/bin
windowsx86_64/bin
```

Is it that simple then?

- Short answer: No - for several reasons
  - Multiple operating systems
  - Multiple data sources to ingest
  - Permission considerations
  - Multiple access tokens

But there is a smooth way to do it

- Creating a tiered app structure
- Base Apps for Windows and Linux
- Include targeted binaries and common configs
  - Reduce the storage and network footprint of the TA deployment
- Server role specific apps with tailored configurations:
  - `agent_config.yaml` and `access_token`

Back to the Lab

Images sourced by Jason Riley

Lab Exercise 2:

Configuration

- Creating a new version of the Splunk TA for OTel Collector for Windows servers
- Modifying the Linux deployment of the Splunk TA for OTel
- Configuring and deploying the Splunk TA for OTel Collector to a Windows instance
- Validating successful deployment and data ingestion into Splunk O11y Cloud

End of Lab 2

You should now have an additional Windows server class on your Splunk Deployment Server.

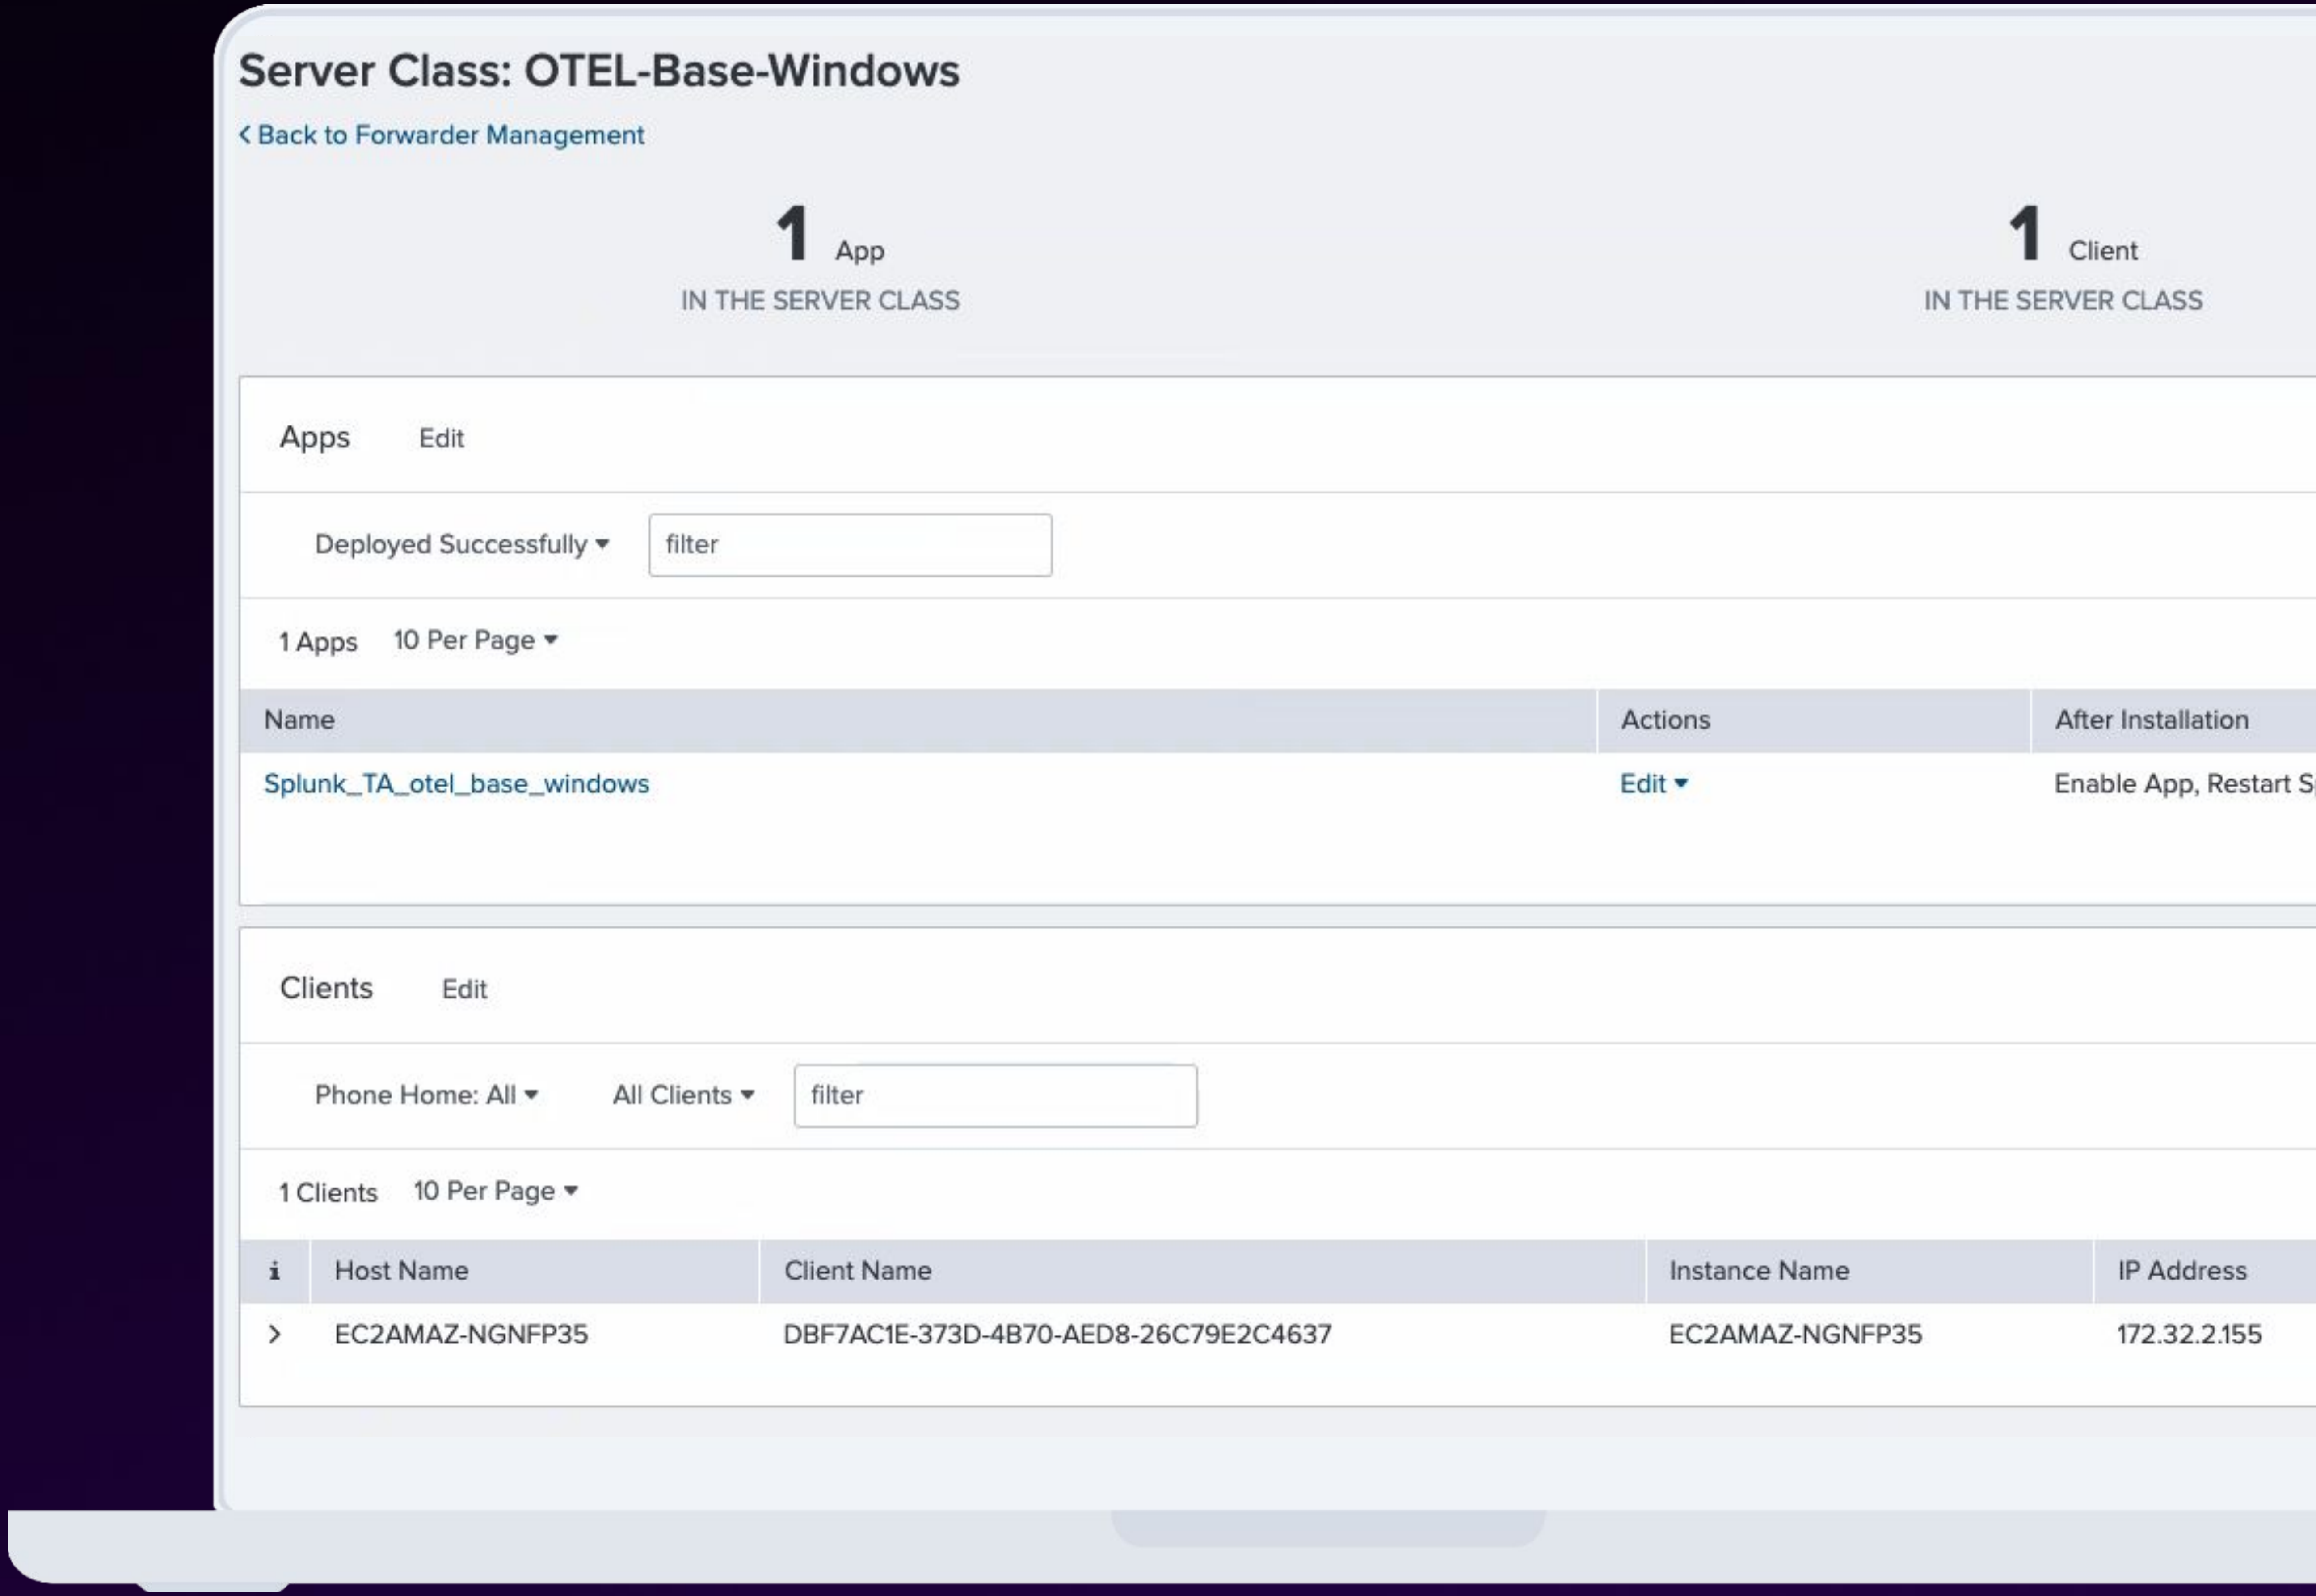

Server Class: OTEL-Base-Windows

‹ Back to Forwarder Management

1 App IN THE SERVER CLASS

1 Client IN THE SERVER CLASS

Apps Edit

Deployed Successfully ▾ filter

1 Apps 10 Per Page ▾

Name	Actions	After Installation
Splunk_TA_otel_base_windows	Edit ▾	Enable App, Restart S

Clients Edit

Phone Home: All ▾ All Clients ▾ filter

1 Clients 10 Per Page ▾

i	Host Name	Client Name	Instance Name	IP Address
›	EC2AMAZ-NGNFP35	DBF7AC1E-373D-4B70-AED8-26C79E2C4637	EC2AMAZ-NGNFP35	172.32.2.155

End of Lab 2

And your Windows server data should be visible in your Splunk Observability Cloud environment.

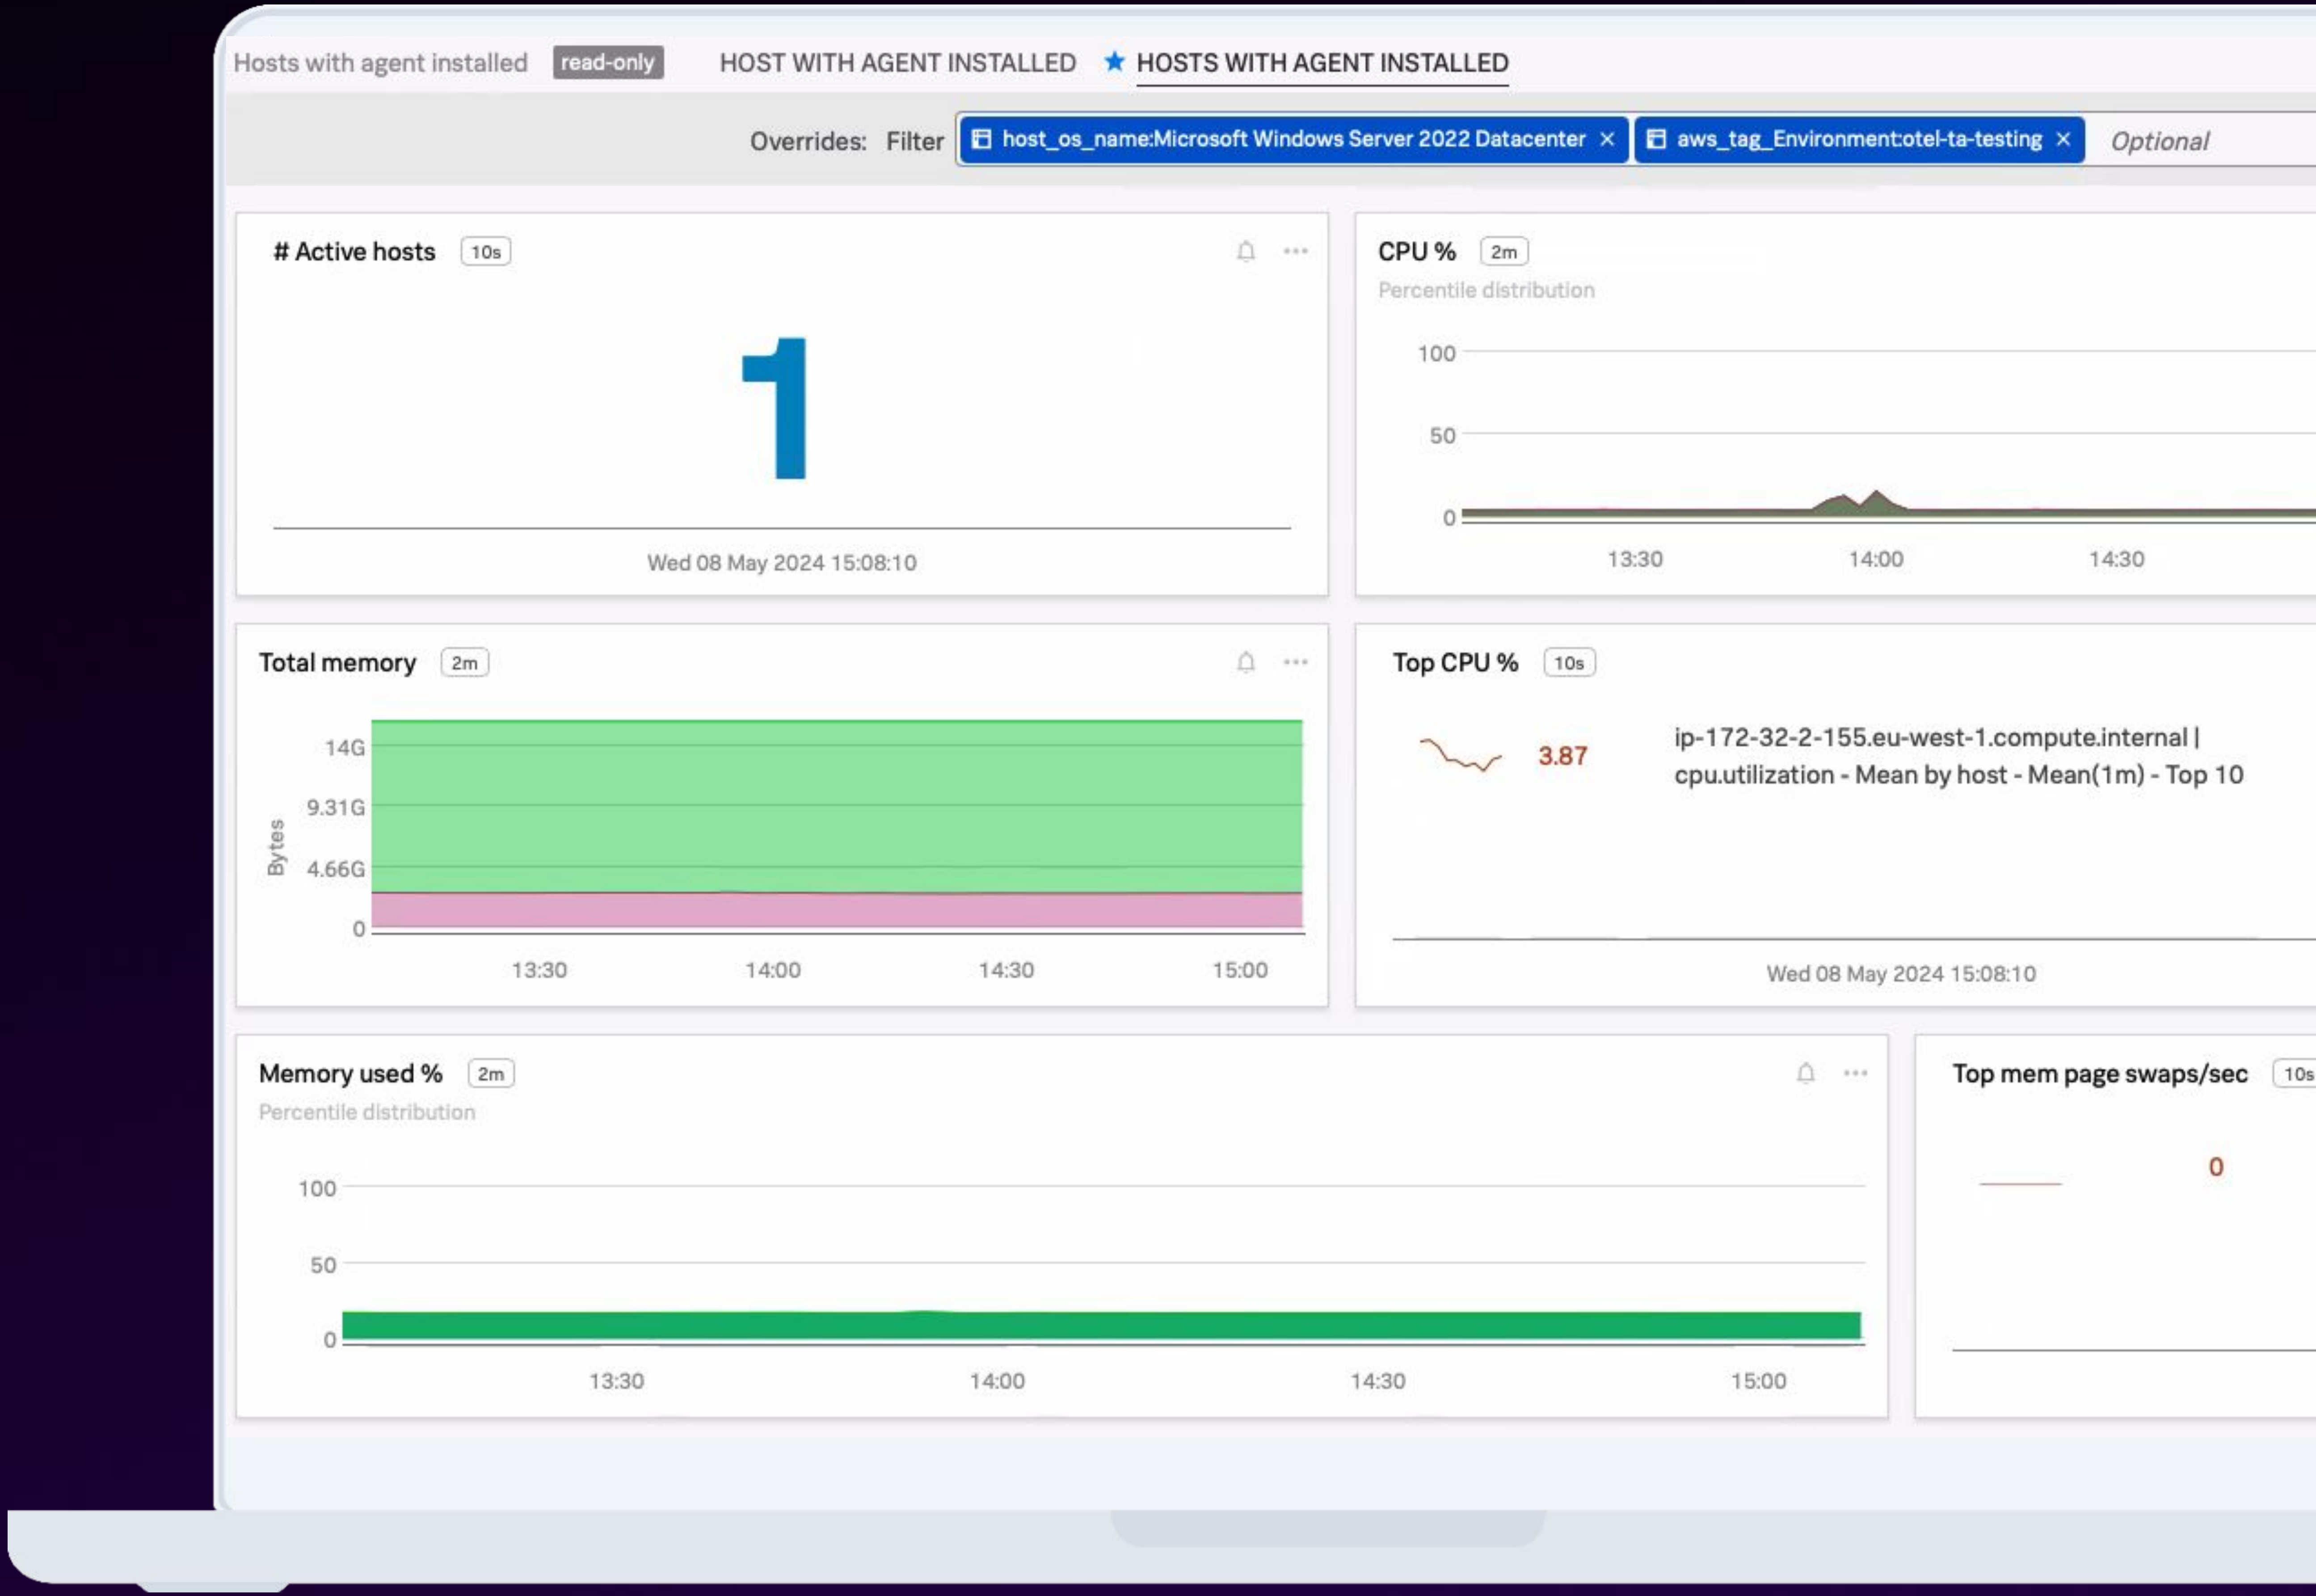

Overrides: Filter host_os_name:Microsoft Windows Server 2022 Datacenter × aws_tag_Environment:otel-ta-testing × Optional

ip-172-32-2-155.eu-west-1.compute.internal | cpu.utilization - Mean by host - Mean(1m) - Top 10

Let's Dig Deeper

Images sourced by Jonathan Fair

App Tiering Continued

- App configurations merge on the DS
  - Lexicographical precedence: Numbers, Capitals, Lowercase
- Naming of folders critical
- Local folder overrides Default
- Let’s see how that looks

App Tiering Structure

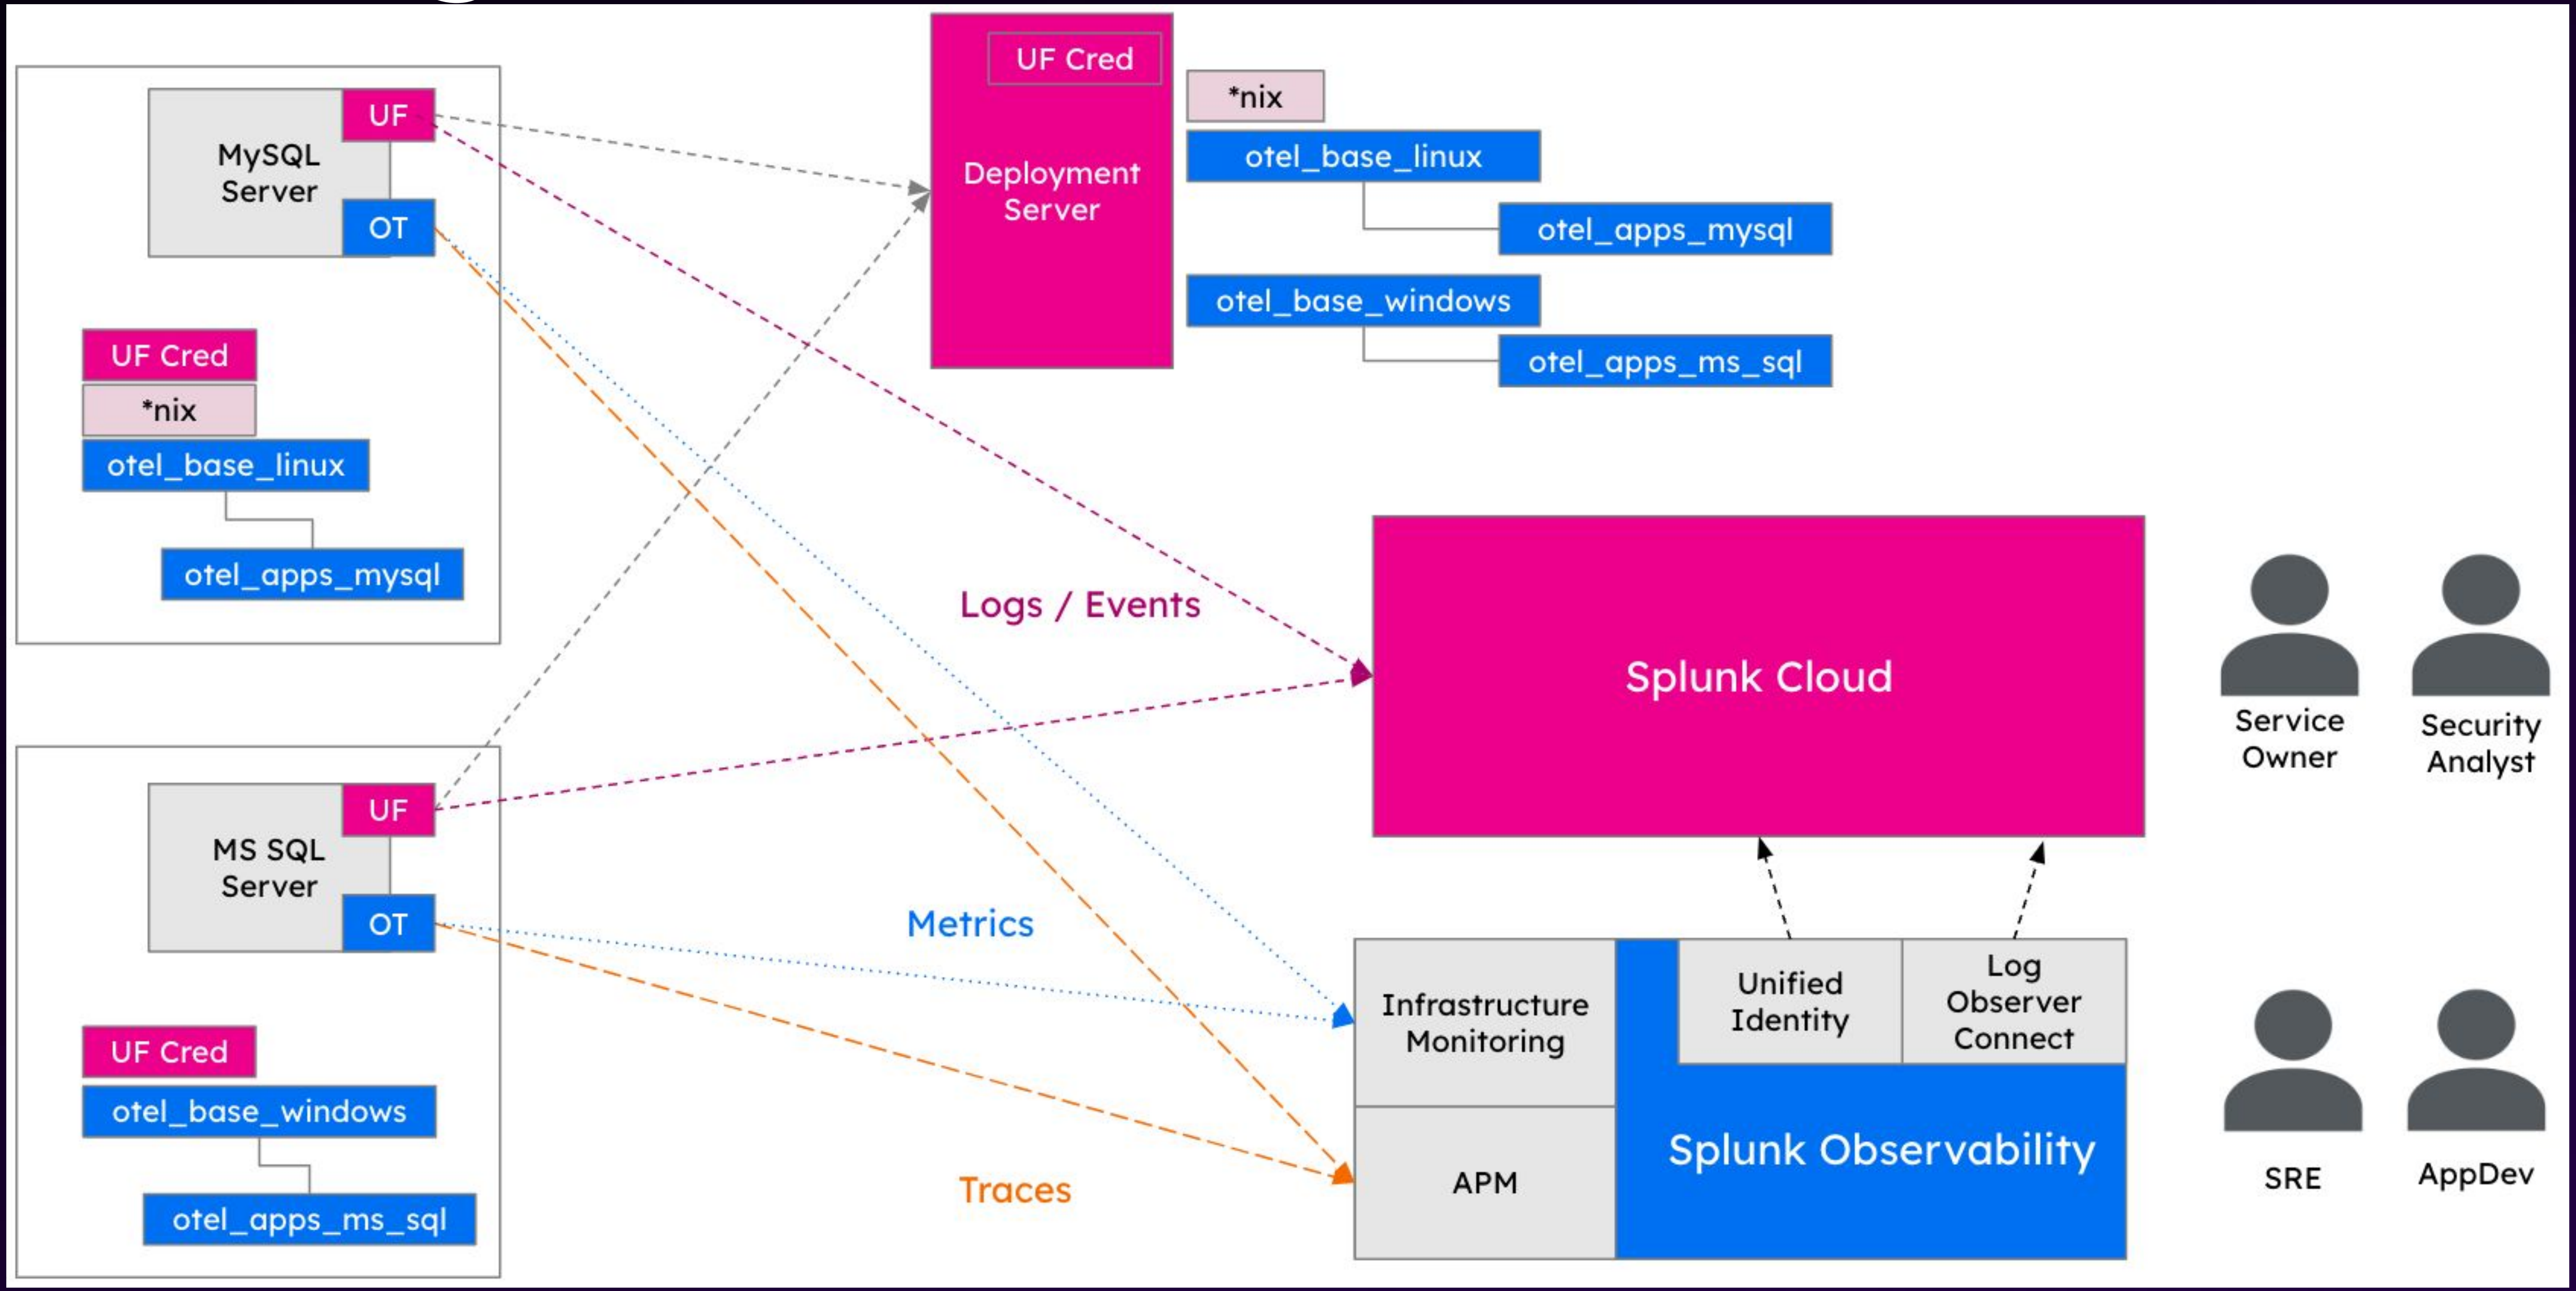

Presenter created diagram.

Troubleshooting

- Check the following logs in `$SPLUNK_HOME/var/log/splunk/`
  - `otel.log` (or custom log file name)
  - `Splunk_TA_otel.log`
- Make sure that there are no competing collector agents running on the server

- Use Splunk Search to easily troubleshoot your logs!!!

Let’s Try Something More Complex

Images sourced by Jason Riley

Lab Exercise 3:

Advanced Use Case

- Modify the Linux deployment of the Splunk TA for OTel
- Add custom configuration for new data source onboarding
- Validate successful deployment and data ingestion into Splunk O11y Cloud

End of Lab 3

For this final lab you should have this additional server class configured on the Splunk Deployment Server

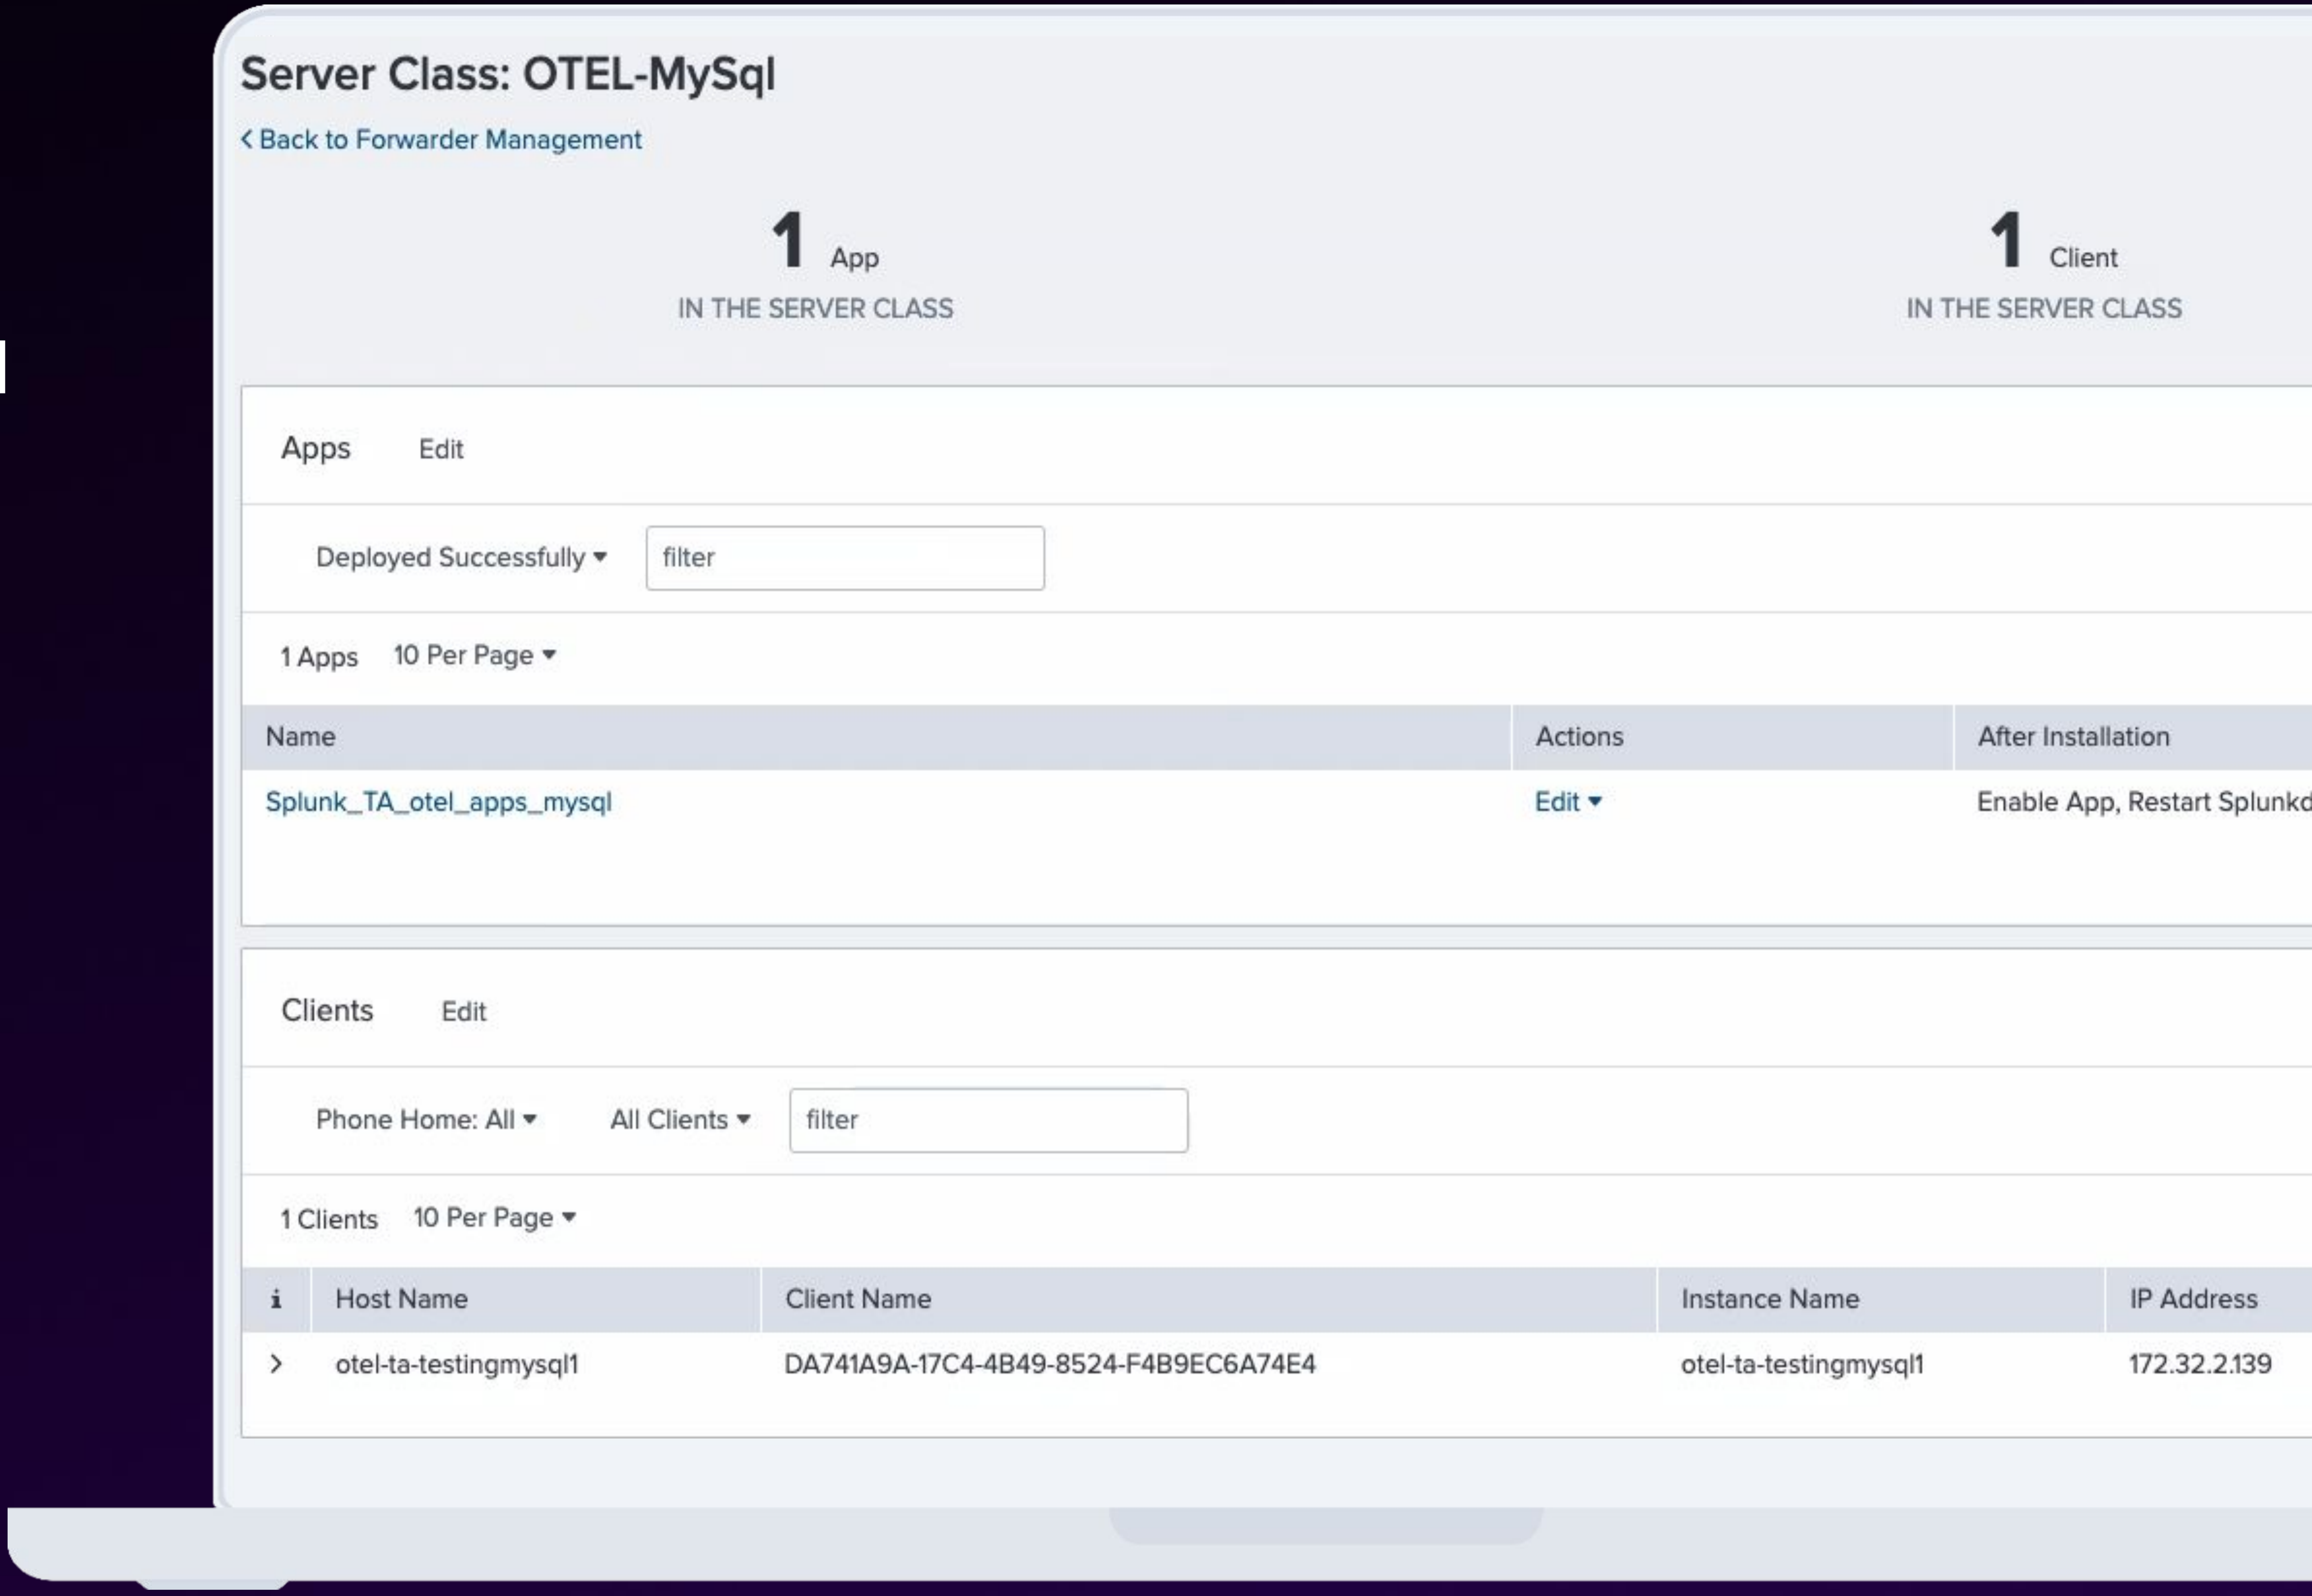

Server Class: OTEL-MySql

< Back to Forwarder Management

1 App IN THE SERVER CLASS

1 Client IN THE SERVER CLASS

Apps Edit

Deployed Successfully ▾ filter

1 Apps 10 Per Page ▾

Name	Actions	After Installation
Splunk_TA_otel_apps_mysql	Edit ▾	Enable App, Restart Splunkd

Clients Edit

Phone Home: All ▾ All Clients ▾ filter

1 Clients 10 Per Page ▾

i	Host Name	Client Name	Instance Name	IP Address
>	otel-ta-testingmysql1	DA741A9A-17C4-4B49-8524-F4B9EC6A74E4	otel-ta-testingmysql1	172.32.2.139

End of Lab 3

And your MySQL server data should now be available to monitor in your Splunk Observability Cloud environment

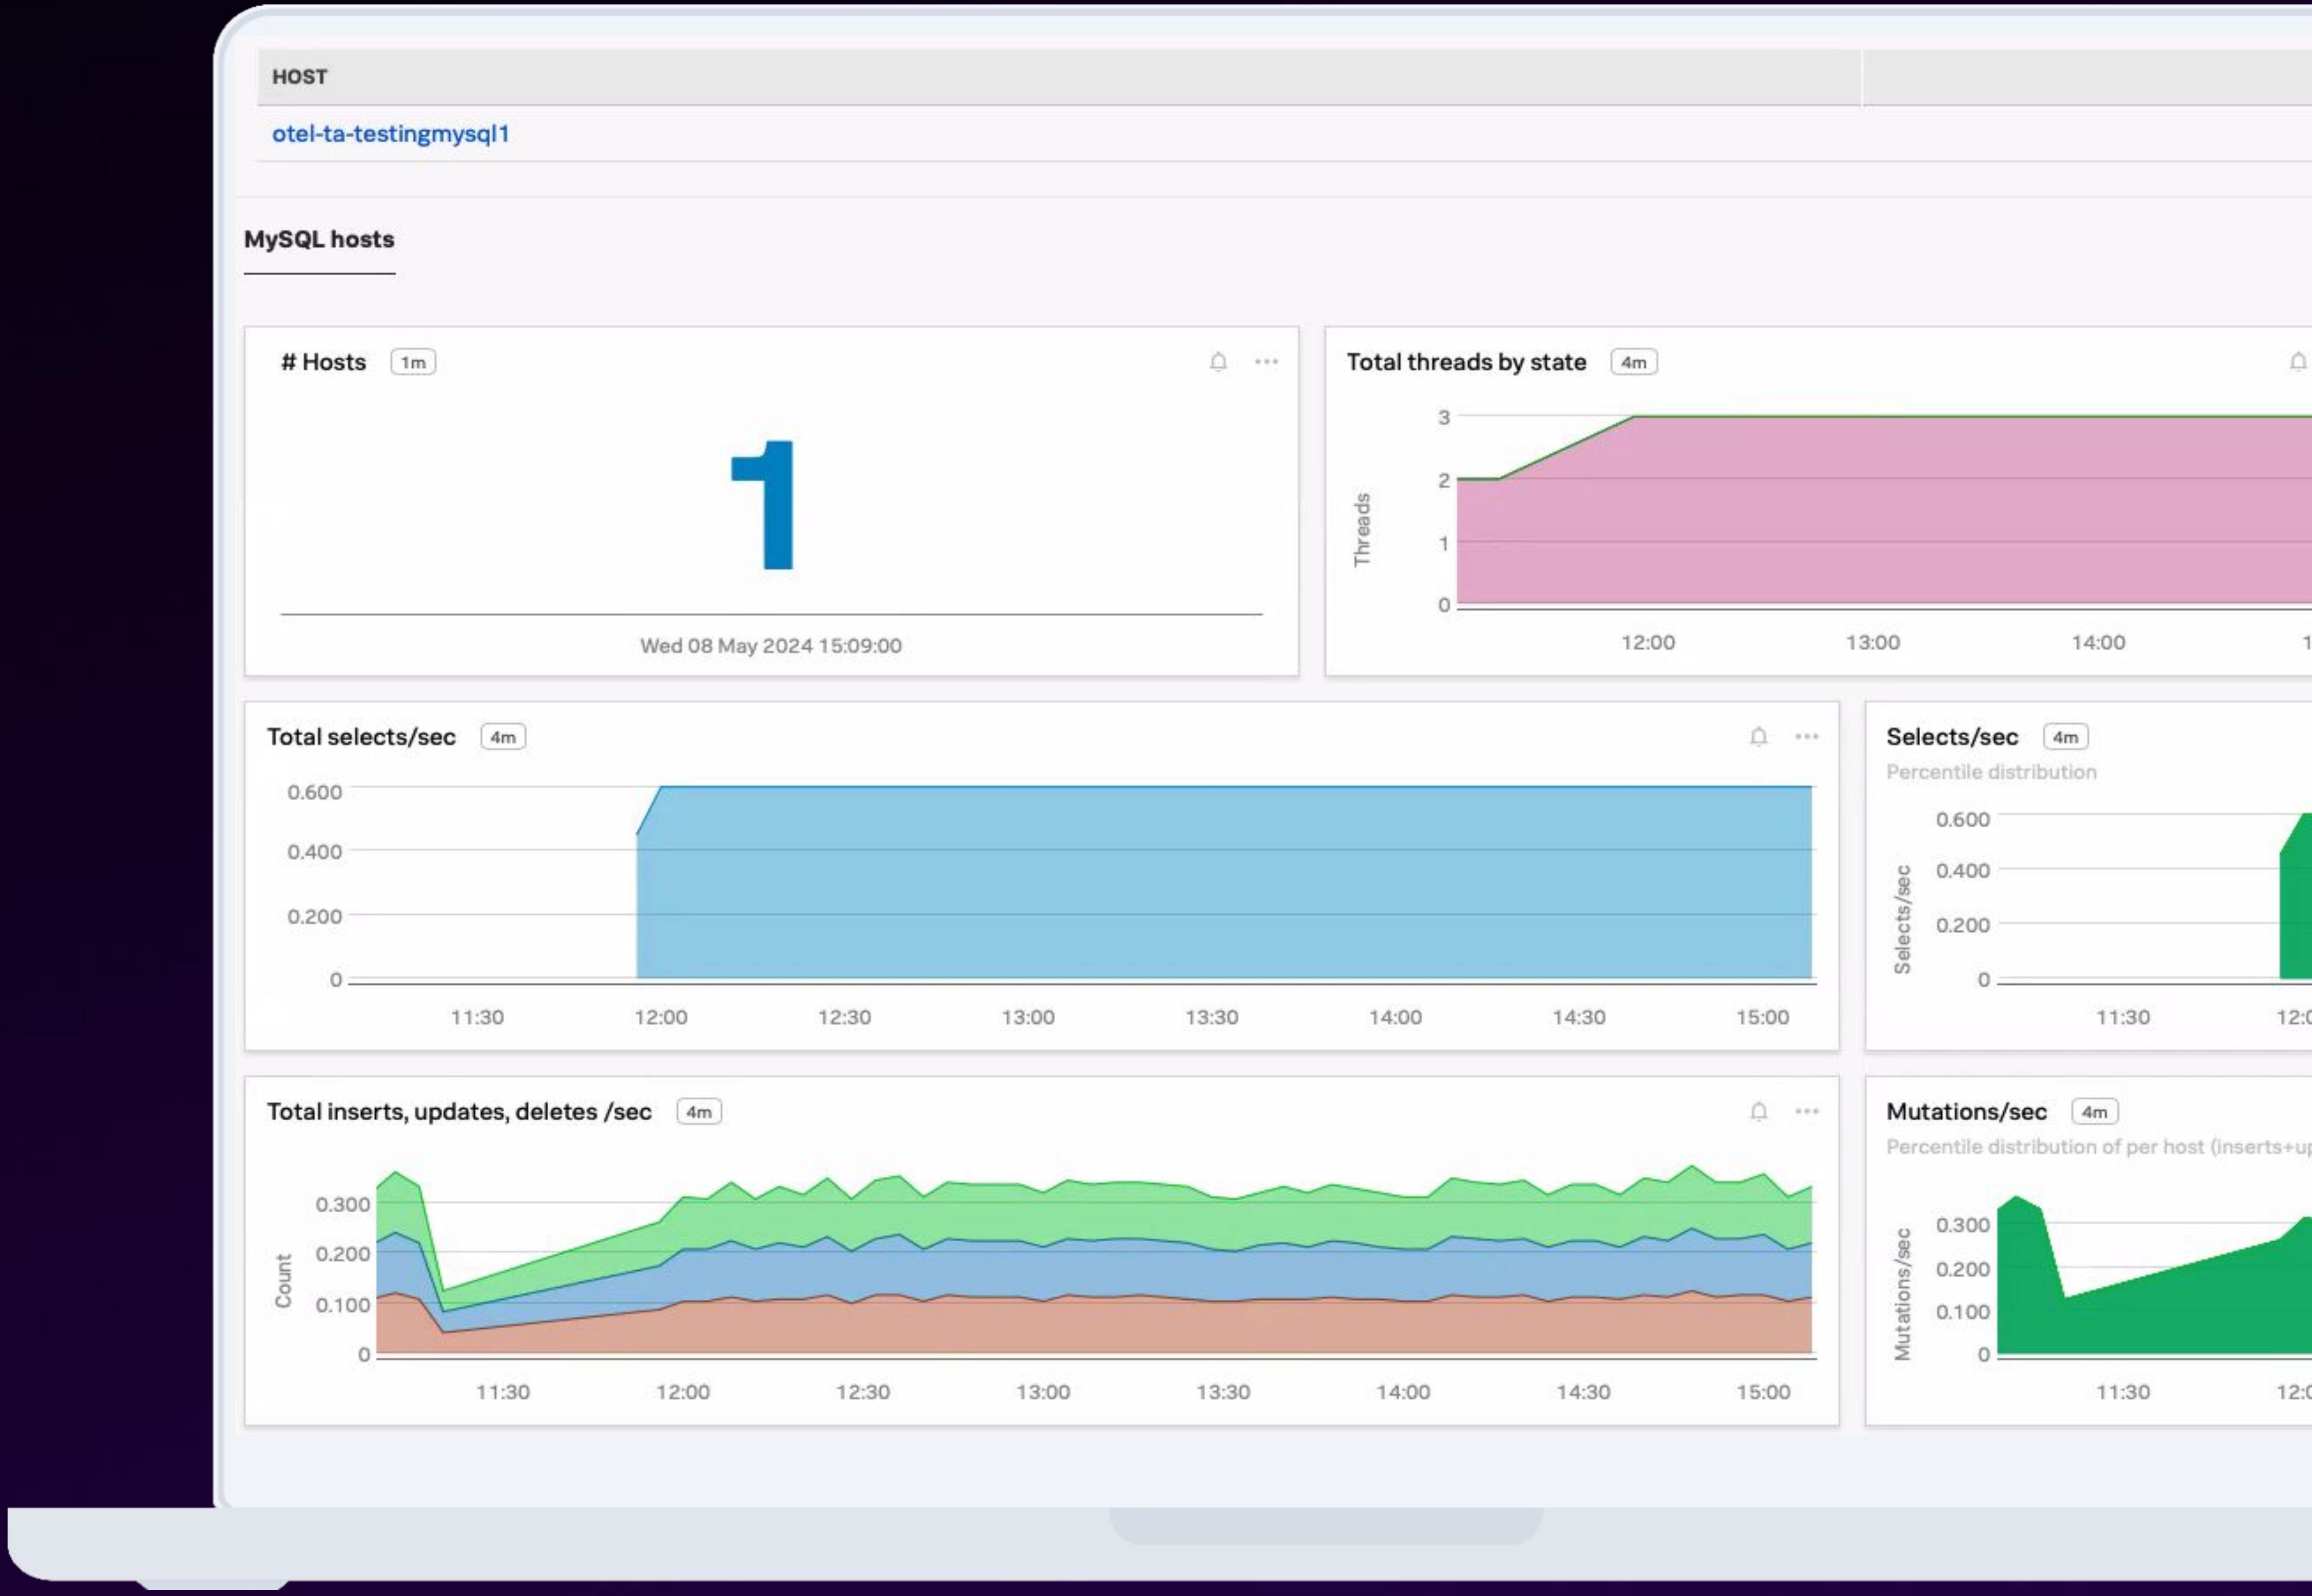

Recap

Images sourced by Jason Riley

So what did we learn?

- OTEL Collector available as Splunk Add-on
- Familiar way to deploy
- Easily get metrics and traces
- Quickly gain full Observability of your environment

Key Takeaways

Splunk Add-on for the OTel Collector

- Available on Splunkbase
- Packages up the Splunk OTel Collector as a UF component
- Enables collection of metrics and traces
- Great for Splunk Cloud or Splunk® Enterprise customers with existing fleet of UFs

Ease of deployment and management

- Can be deployed through the familiar Splunk Deployment Server
- Flexibility of choosing target hosts for deployment
- Compatible with other 3rd party configuration frameworks

Considerations and best practices

- Same permissions as the UF - can cause issues with certain data sources
- As of now, lack of zero configuration support
- One-size-fits all deployment approach won’t work
- Tiered app structure is the way to go

Any questions?

So you wanna get started?

- Splunk Add-on for the OTel Collector - Splunkbase Download
- Scope out the Observability Demo Area @ the Pavilion!!!

Special Thanks

Georgios Glymidakis

Original Co-Presenter who couldn’t be here today

.conf24
splunk>

Special Thanks

Couldn’t have gotten this together without the support of these people

- Aunsh Chaudhari
- Chris Crocco
- Daniel Pan
- Geoff Higginbottom

Thank you

Scope out the Observability Demo Area
at the source=*Pavilion

.conf24
splunk>