Using Splunk Application Performance Monitoring

splunk>

Foundational Knowledge

- To be successful, students should have a working understanding of these topics:
  - Familiarity with navigating Splunk Observability Cloud
  - Basic knowledge of visualization and alerting on metrics in Splunk Observability Cloud

Course Objectives

- Identify key elements of Splunk APM
- Differentiate between metrics, traces and logs
- Define spans, traces, metrics, and metadata
- Navigate Splunk APM
- Apply Splunk APM features to troubleshoot an issue

Course Outline

- Module 1: Overview of Splunk APM
- Module 2: Environments, Services, and Endpoints
- Module 3: Traces, Spans, Tags, and Business Workflows
- Module 4: Profiling and Database Query Performance Monitoring
- Module 5: Troubleshooting Using Splunk APM

Module 1: Overview of Splunk Application Performance Monitoring (APM)

Module Objectives

- Describe the Splunk Observability solution
- Identify key elements of Splunk APM
- Differentiate between metrics, traces and logs

Splunk Observability Cloud

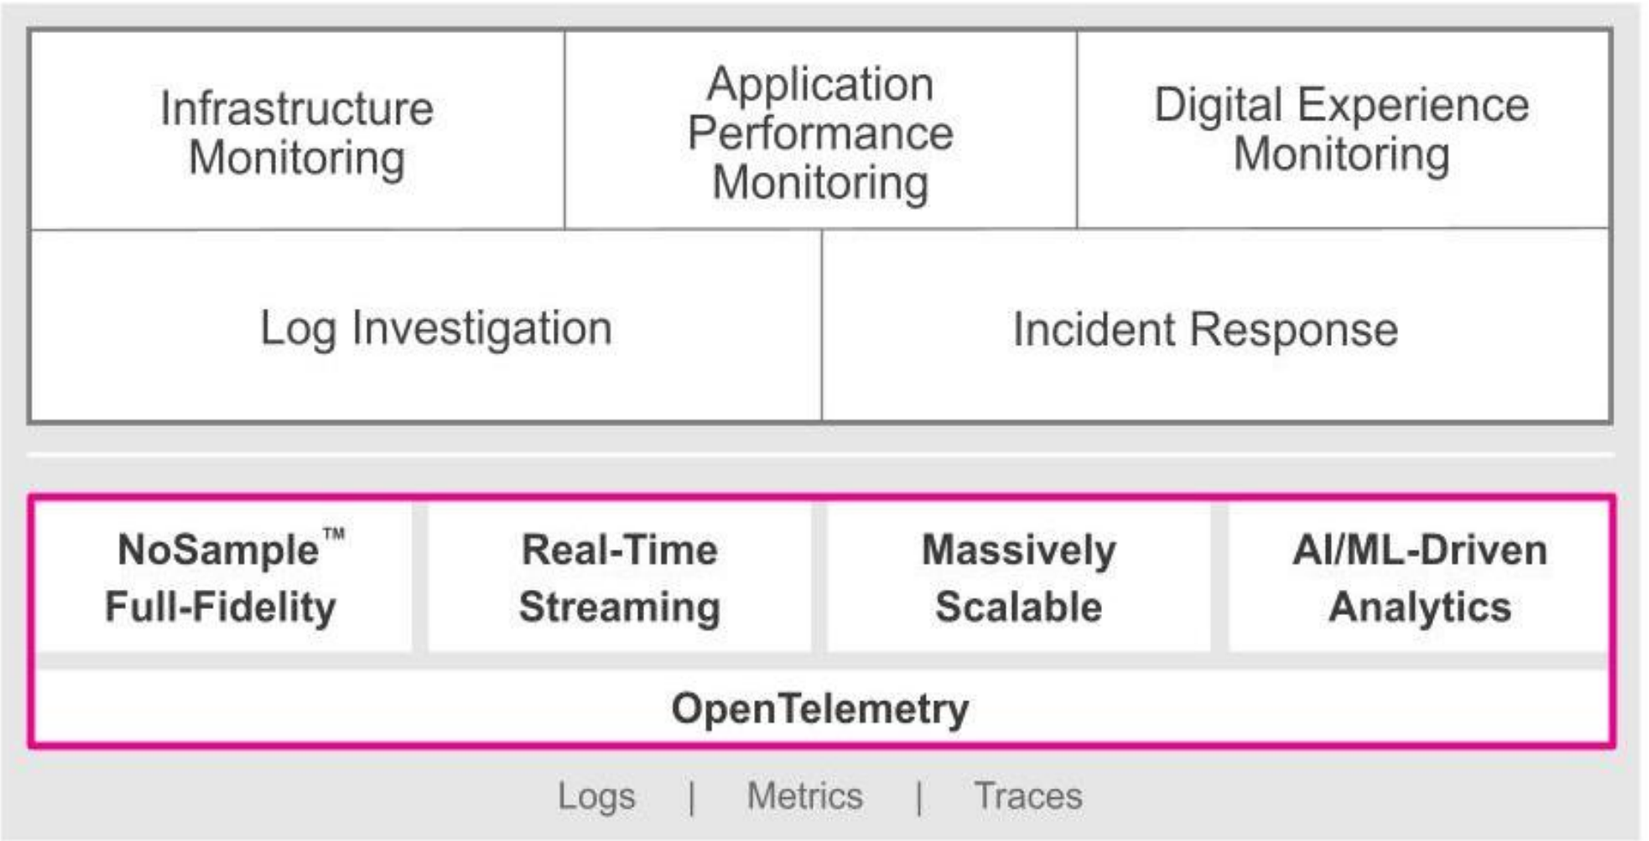

Real-Time Streaming Architecture

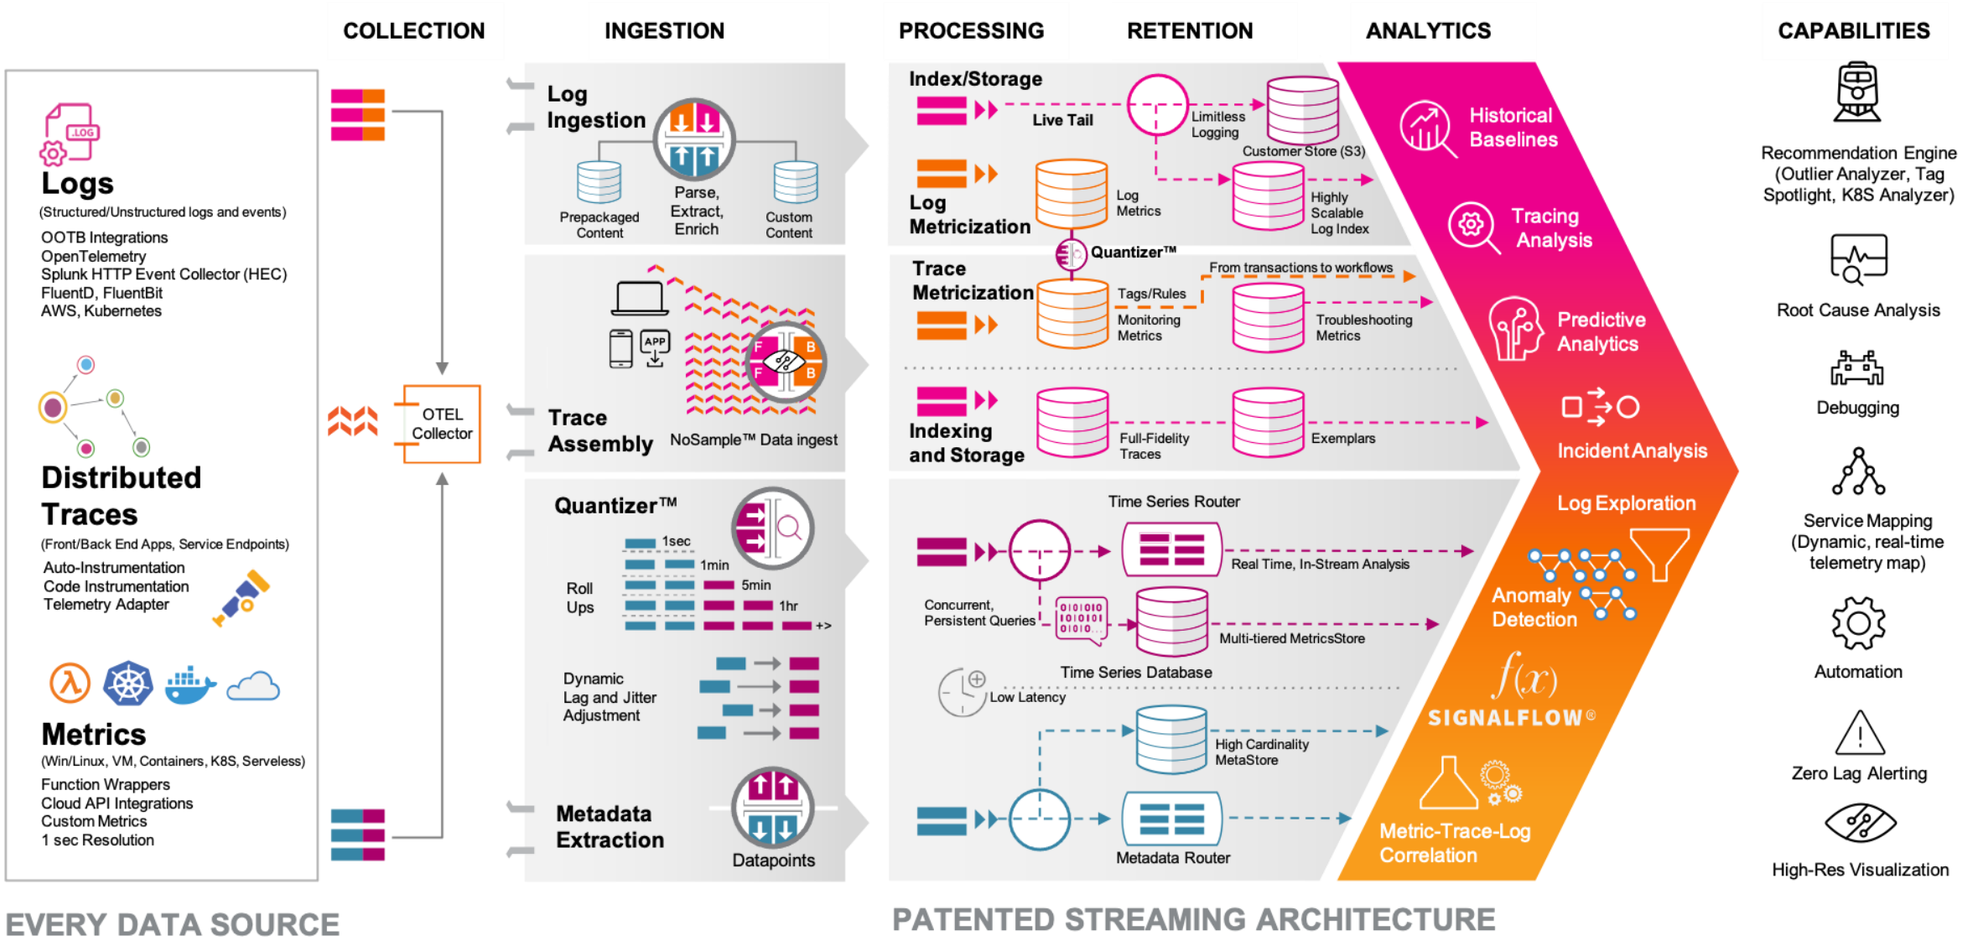

Real-Time Streaming Architecture

How We Get Data Into Splunk Observability Cloud

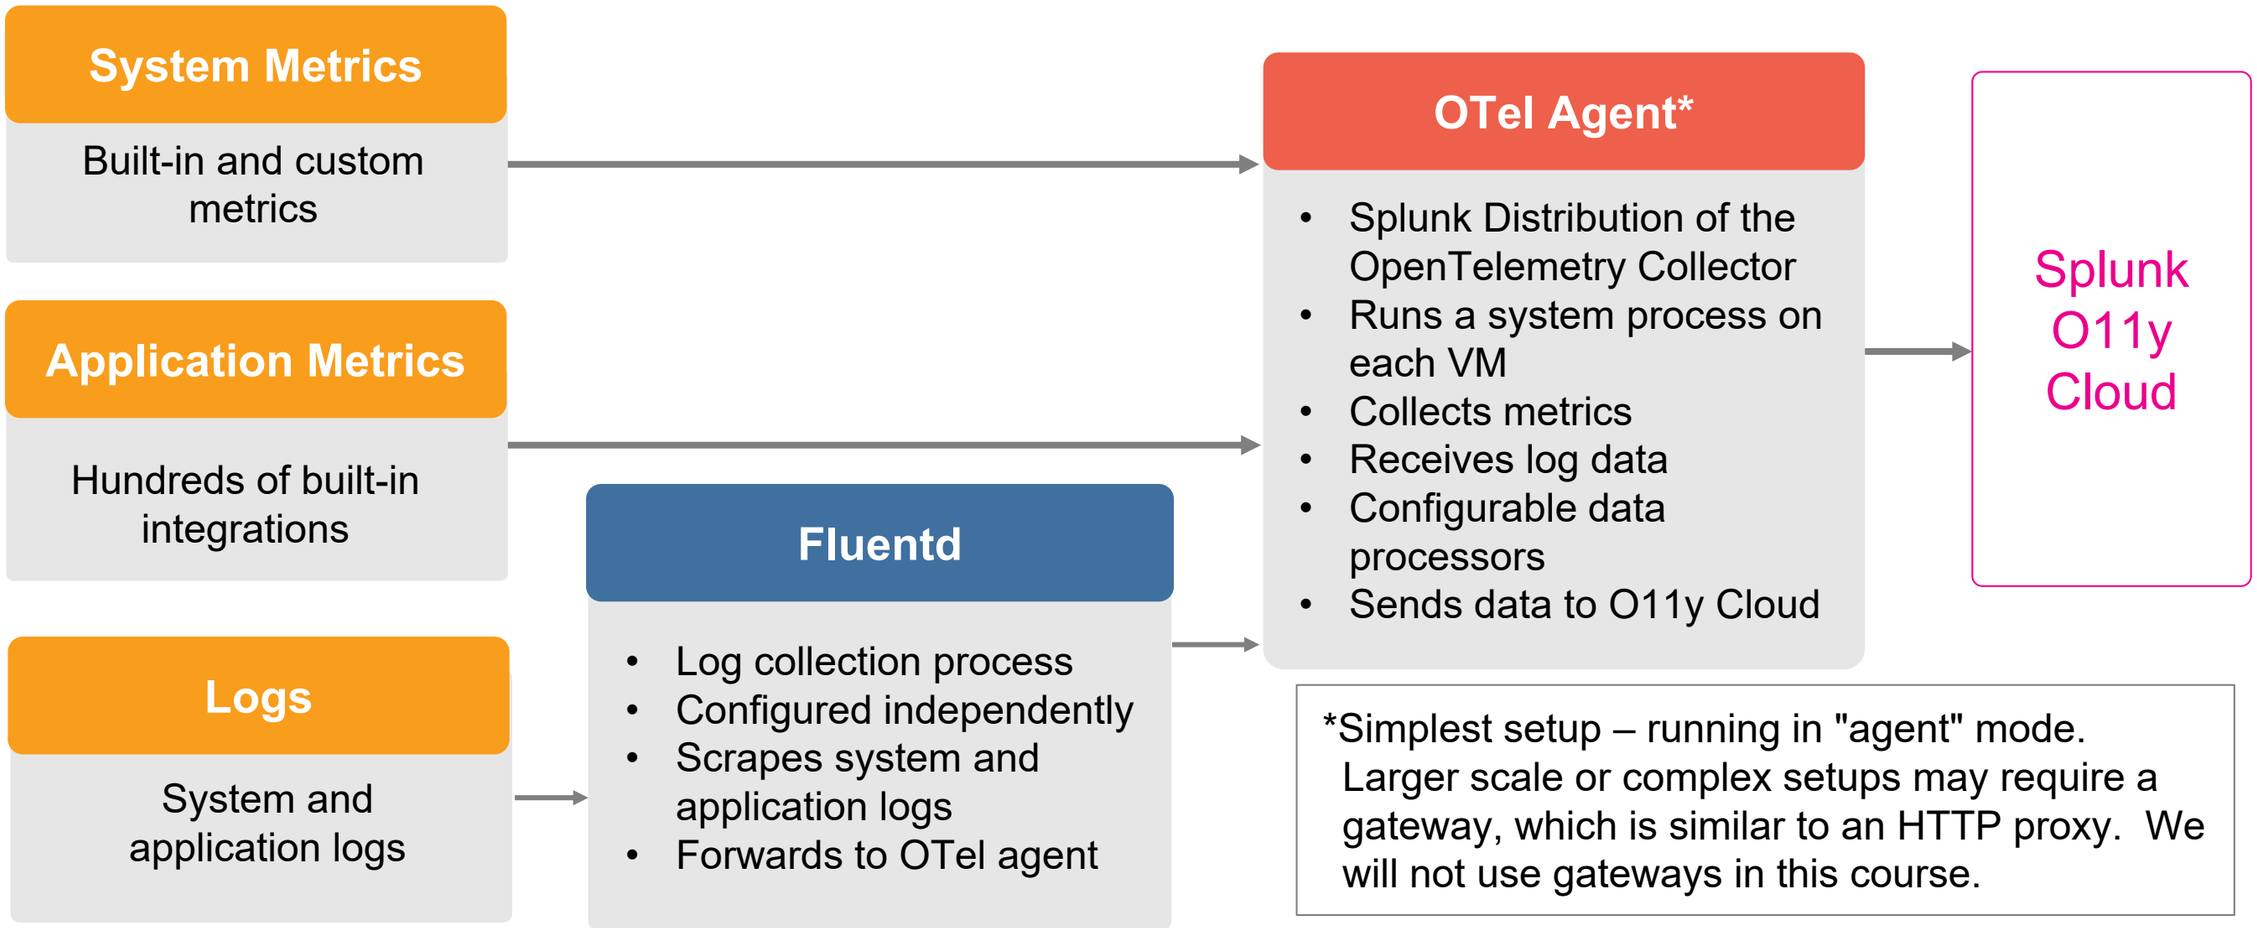

About Splunk APM

- Analyze all traces
  - Rate, Error and Duration (RED) metrics and percentiles generated for every trace and span are collected and stored for 13 months
  - Traces and spans with granular level of detail of each operation and endpoints are stored for 8 days
- Ability to break down application behavior based on span metadata (container, customer or key business logic)

About Splunk APM (cont.)

- Auto-generated dashboards (service, endpoint, business workflows)
- Real-time performance of upstream and downstream dependencies and underlying infrastructure
- Real-time alerting
- One-click troubleshooting
  - Easily identity root cause errors
  - Drill down to code-level

Who uses Splunk APM?

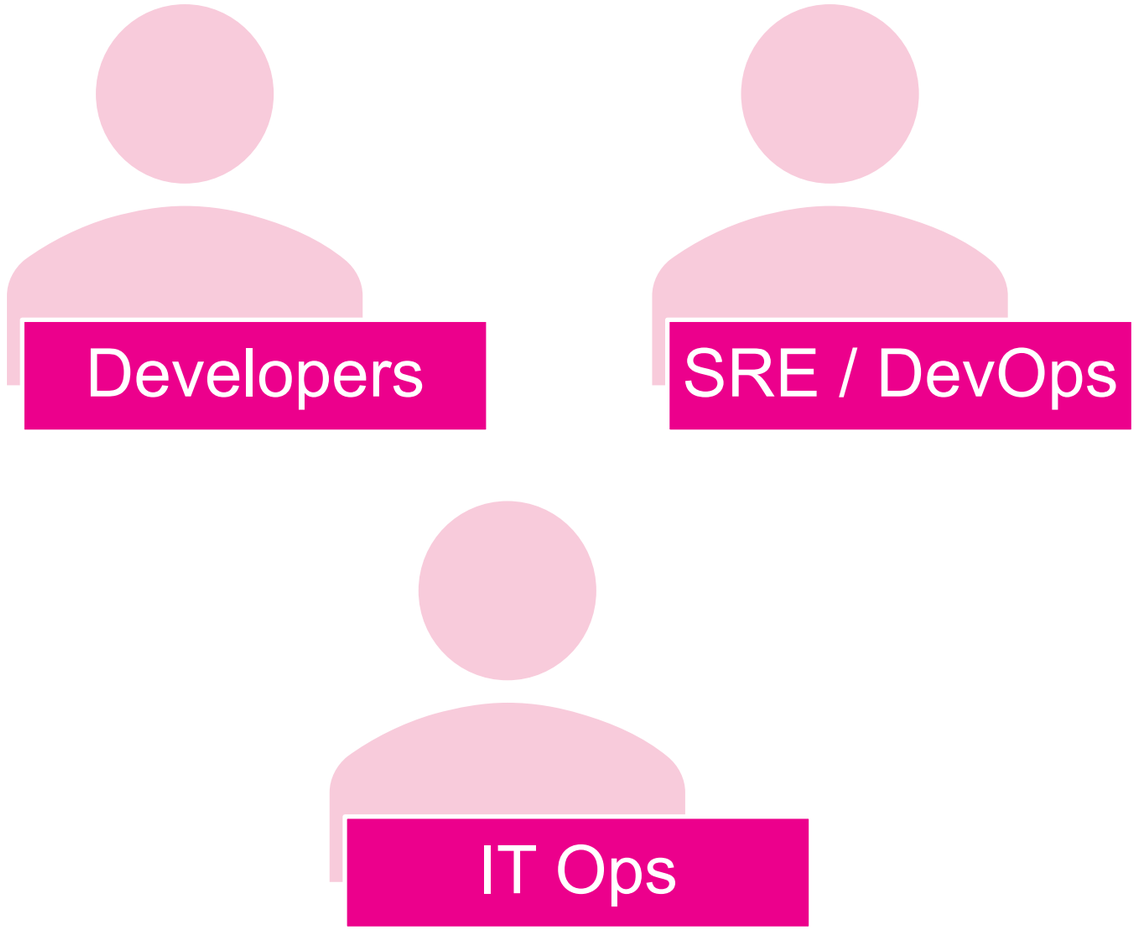

Why do People use APM?

1

Monitoring

- Alerting on SLIs as well as end-user experiences
- Monitoring during deployment / release

2

Troubleshooting

- Troubleshooting application availability issues such as errors and high latency
- Root causing to service/endpoint, underlying infra, database or down to line of code
- Understanding radius of impact of an incident

3

Optimization

- Identifying and addressing bottlenecks in code and architecture

About Tracing

- Records operations while processing a user request
- Propagates the trace context from service to service - keeps context
- It is user/event-triggered
- Contains parent-child relationship → topology information
- Data from each operation is converted into **metrics** enabling you to look at historical data, patterns

About Tracing (cont.)

Tracing is like using a microscope to see the details of your application's services and operations

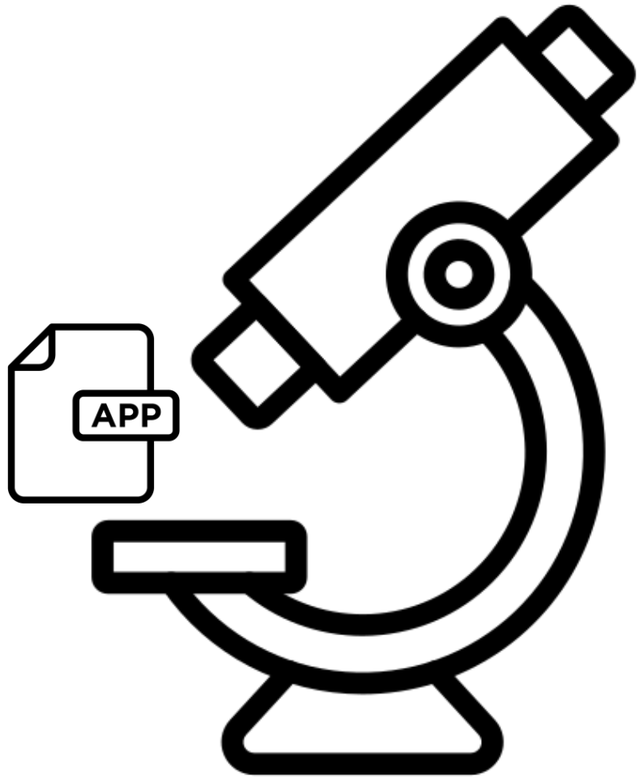

SPAN	
Span ID	c825c8db904f7676
Parent ID	12d8815cc8aaf113
OPERATION	
loadCart	
TAGS	
container_id	"ssidhucheckout-162837"
container_image	"quay.io/image/checkout"
container_spec_name	"checkout"
environment	"ssidhu"
http.method	"POST"
http.status_code	"200"
http.url	"/checkout/42"
kubernetes_cluster	"ssidhu"
kubernetes_node	"ssidhunode1"
kubernetes_node_id	"ssidhunode1-sto7w34t"
kubernetes_pod_name	"ssidhucheckout-1"
kubernetes_pod_uid	"ssidhucheckout-1"
namespace	"commerce"
span.kind	"client"
tenant	"Gold"

Traces and Spans

- Splunk APM ingests trace spans
- A trace is comprised of spans and is captured by instrumenting code
- A span represents an endpoint, operation or section of code from an application
- Spans have a reference to their parent span, turning the trace into a tree

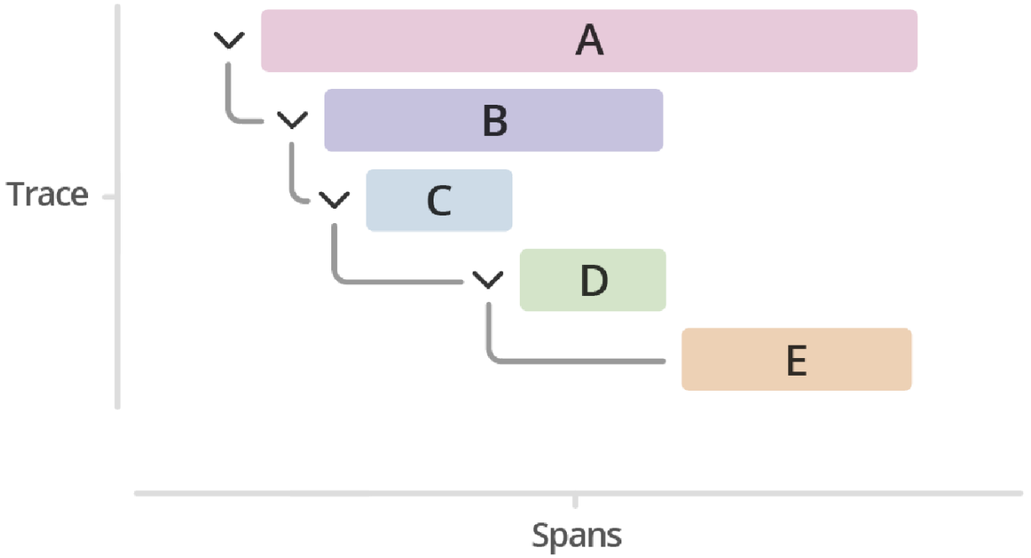

Traces and Spans (cont.)

Data Driven Approach to Observability

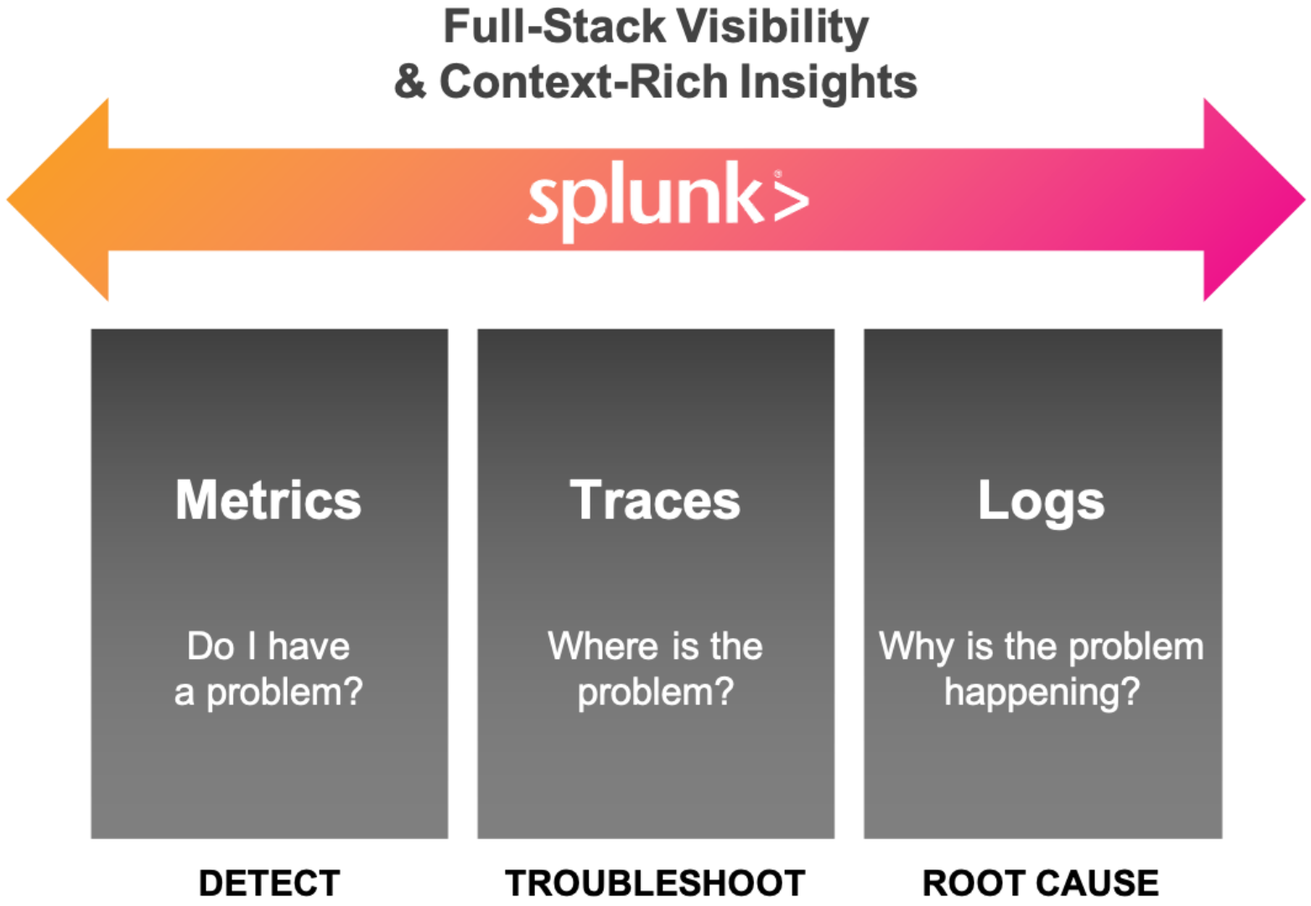

Live Demo Scenario

- Imagine you are the service owner of an online boutique retailer for the **ad** service
- Using the online boutique, customers can browse items in the store
- Customers have been complaining that pages with ads are slow to load
- You represent an SRE, DevOps engineer or a developer who will be using Splunk APM to understand what is going on

Live Demo Scenario (cont.)

- Distributed services environment
- Multiple services written in different languages
- Search for an item
- Click on item from search results
- Services include:
  - Recommendation
  - Ad
  - Shipping
  - Payment

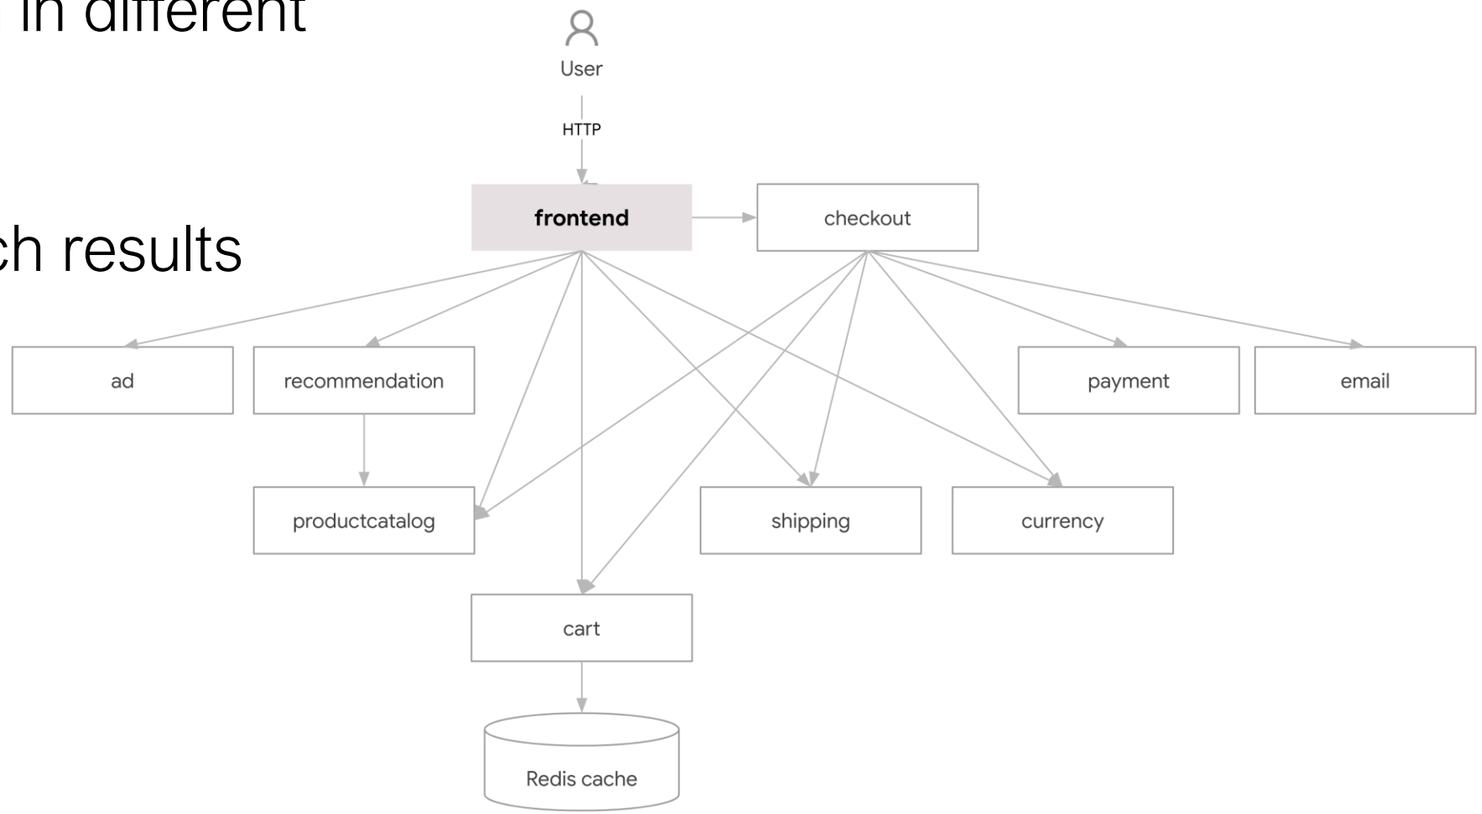

Lab Data Scenario

- You are the service owner of an online boutique retailer for the **checkout** service
- Using the online boutique, customers can add items to the cart but when they place an order or checkout, the transaction takes much longer than expected
- You represent an SRE, DevOps engineer or a developer who will be using Splunk APM to understand what is going on

Module 1 Lab Exercise 1

Time: 15 minutes

Description: Identify metrics that could be used for the online boutique and explore the APM landing page

Tasks:

- Review metrics, traces, and APM definitions
- Create metrics for the online boutique
- Explore the APM landing page

Module 2: Environments, Services, and Endpoints

Module 2 Objectives

- Define environments, services, endpoints, and operations
- Explore the following to troubleshoot latency and errors in the ad and checkout services:
  - APM Overview Page
  - Service Map

Environment

- Distinct deployment in Splunk APM that does not interact directly with other deployments of the same application
- A logical application environment or deployment
  - Dev, QA, Prod environments
- Specified by the sf_environment span tag
- Must be present on every span! (important for metrics)
- If not specified, assume environment=unknown

Service

- A small, flexible autonomous unit of software
- Interacts with other services via RPCs (HTTP or otherwise) or Pub/Sub systems to make a complete application
- Has one or more service instances deployed and serving requests
- Or may be “realized” by cloud functions
- Service names are stable, usually low 10s or 100s of services in an environment
- Specified through the tracer’s configuration

Inferred Service

- A remote service that is not instrumented in Splunk APM
- Usually, 3rd party services like databases and queues
- Span metadata helps determine the type and infer a name
- No Monitoring MetricSets available for inferred services (no dashboards)

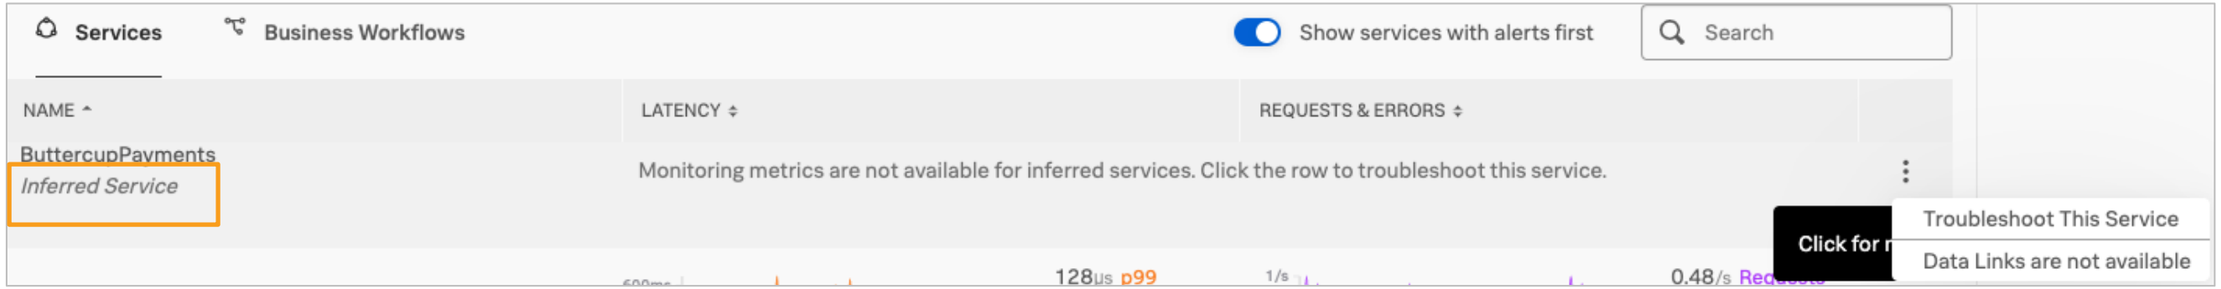

Endpoints and Operations

Term	Description
Endpoint	The first point into a service
Operation	A span within a single service

Getting Started

- Start here: **Menu > Splunk APM**
- Select your environment
- Find answers to these questions:
  - Which services have alerts* firing?
  - Which services have the highest latency?
  - Which services have the highest error rates?
  - Has anything changed with that service over time? (Service Dashboard)
  - Explore further
- Each span must have an environment tag. If no value set for environment, service appears under Unknown or All

Splunk APM – Landing Page

- Worst performers at a glance
  - Top services by error rate
  - Top services by latency (P90)
- For each service, provides:
  - Request rate and error rate
  - P99, P90 and P50 of latency
  - Link to built-in APM dashboard
  - Link to troubleshooting
- To investigate further, open the APM Services dashboard for a service

Splunk APM – Service Dashboard

- To access the Services dashboard from the Splunk APM > Overview page
  - Select the actions menu from a service in the **Services** menu
  - Select View Dashboard: APM Service
- Services Dashboard displays Rate/Error/Duration (RED) metrics for the selected service
- It also provides information about the endpoints such as errors or requests per endpoint
- Services Dashboard also includes infrastructure metrics related to the service

Splunk APM – Service Dashboard (cont.)

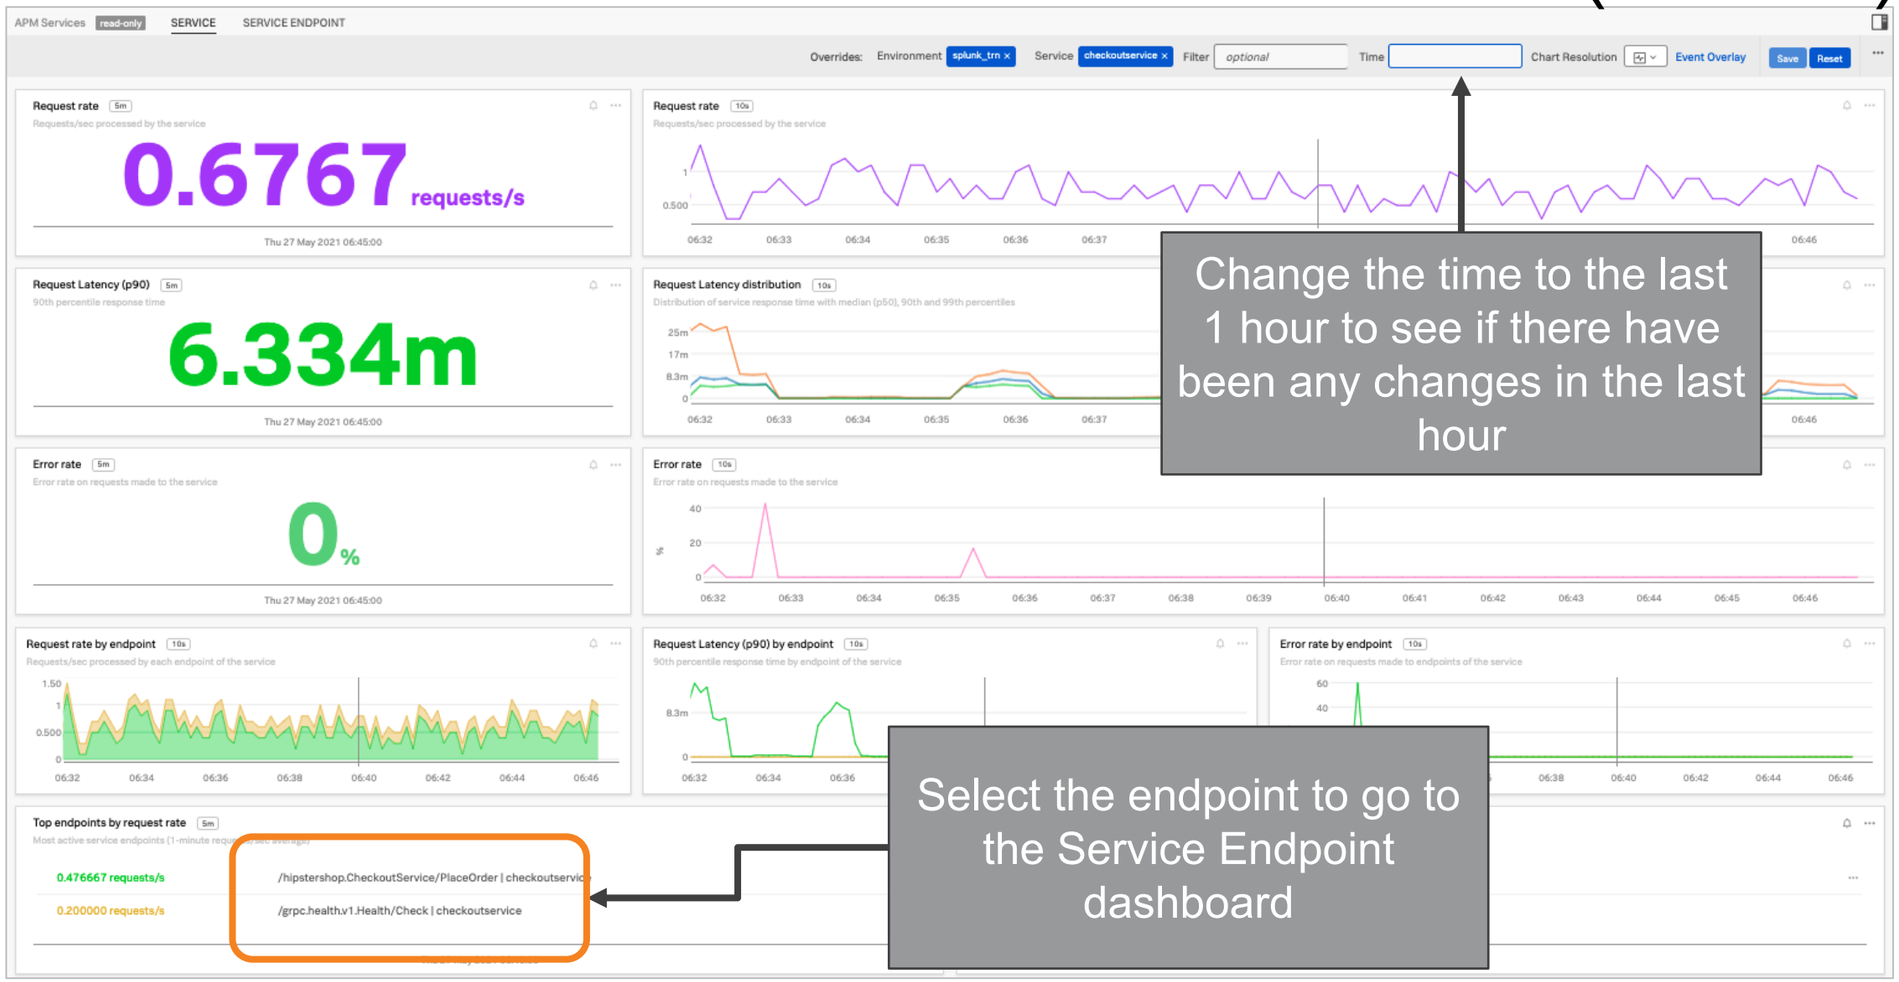

Splunk APM – Service Endpoints Dashboard

- This is similar to the Service dashboard
- Filter on endpoints in addition to services
- You can navigate to the Service Endpoint dashboard by clicking an endpoint from the Top Endpoints by request rate (or error rate) chart in the Service dashboard
- Keeps the context when you go from Service to Endpoint

Explore View

This page provides fine-grained and causality details related to service performance degradation

How to get to the Explore View?

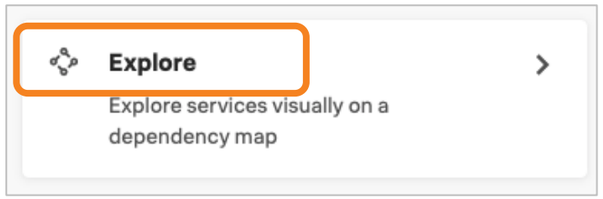

Explore option in the APM Overview Page

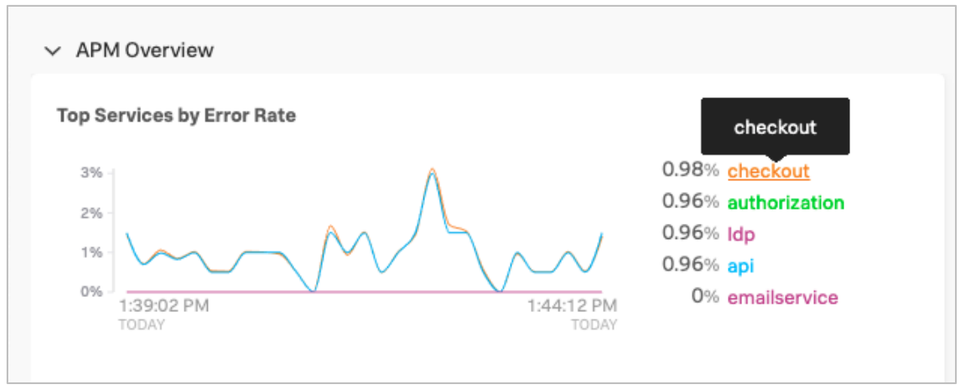

Clicking a Service name in the APM Overview page

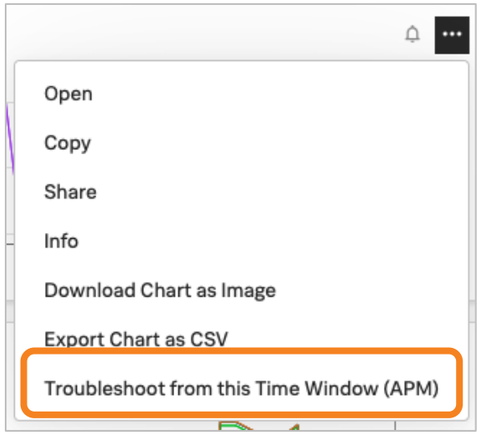

Chart Actions menu on charts in APM built-in dashboards

Explore Services

- Overall Service map
  - Real-time view of your application's architecture
  - Very quickly determine which services have errors and which service interactions have high latency
  - Gives you possible cause of errors from downstream services
- Filters
  - You can filter by environment, workflows, service and tags
  - You can also filter to a service by clicking the service in the service map
  - Charts on the right are filtered for the selected service

Explore Services (cont.)

- Find service dependencies
  - Inbound and outbound services
  - RED metrics for each dependent service
- View traces
- Open Tag Spotlight

Service Maps

- A service is intended to represent a component of your distributed application's architecture
- Nodes represent logical services that make up your architecture and their dependencies

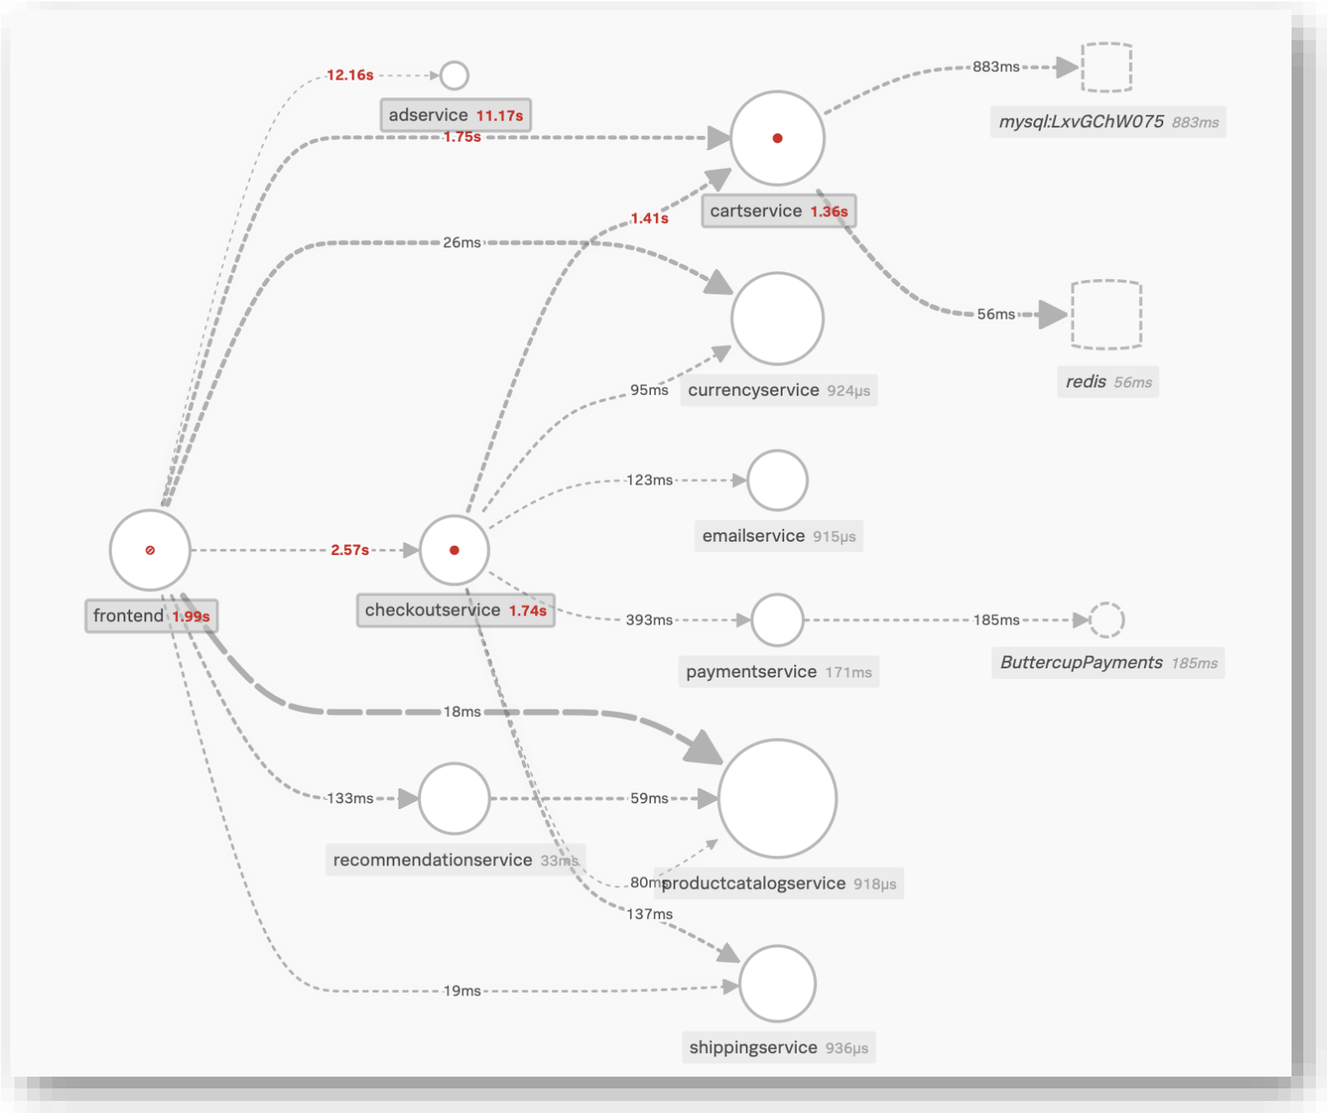

Service Maps (cont.)

- Using the service map, you can do the following:
  - Select any service in the service map to see charts for that service
  - Use the Breakdown selector to break the service's activity down by any indexed span tag
  - Select any chart in this view to show example traces that match the parameters of the chart
  - Select **Endpoint Performance** to filter, sort, and compare endpoints

Investigating Services Scenario

- Customers are complaining about slowness on item pages with ads
- Use the APM overview page to see that there is latency in the **ad** service
- Use the Service map to investigate more details about the latency issue

Investigating Services Steps

- Navigate to the **APM Overview** page
- Filter for the appropriate environment
- Look under the **Services** tab and select a service
- Select the service to view the service map
- Navigate back to the **APM Overview** page
- Select the **Explore** option to see the high level service map
- Service map highlights errors and latency with red

Module 2 Lab Exercise 2

Time: 20 minutes

Description: In this exercise you will navigate the service views and investigate the service dashboards

Tasks:

- Identify Services with High Latency and Errors
- View Service and Endpoint Dashboards
- Explore the service

Take a break

Module 3: Traces, Spans,Tags, and Business Workflows

Module 3 Objectives

- Define spans, traces, metrics, metadata in more detail
- Explore the following to troubleshoot latency and errors in the ad and checkout services:
  - Trace Analyzer and Search
  - Span Waterfall View
  - Tag Spotlight
- Explore Tags and MetricSets

What Is Tracing

A way to record all operations involved in the handling of a request or transaction through the entire application stack and backend infrastructure

What is Instrumentation

Process through which application code is extended to capture and report trace spans for operations of interest in the path of handling a particular request or transaction

Instrumentation Agnostic

- Splunk APM is instrumentation agnostic
- Supports instrumentation libraries such as OpenTelemetry, Jaeger and Zipkin
- As long as spans are sent in a recognized format, it doesn't matter how you instrument your code
- Default instrumentation for languages such as Java, Ruby, Python, Node, Go, .NET, PHP

Auto Instrumentation

- Run-time automated process automatically identifies which frameworks and libraries are in use
  - Instruments those libraries to capture tracing instrumentation (without code change)

Splunk APM: Full-Fidelity Architecture

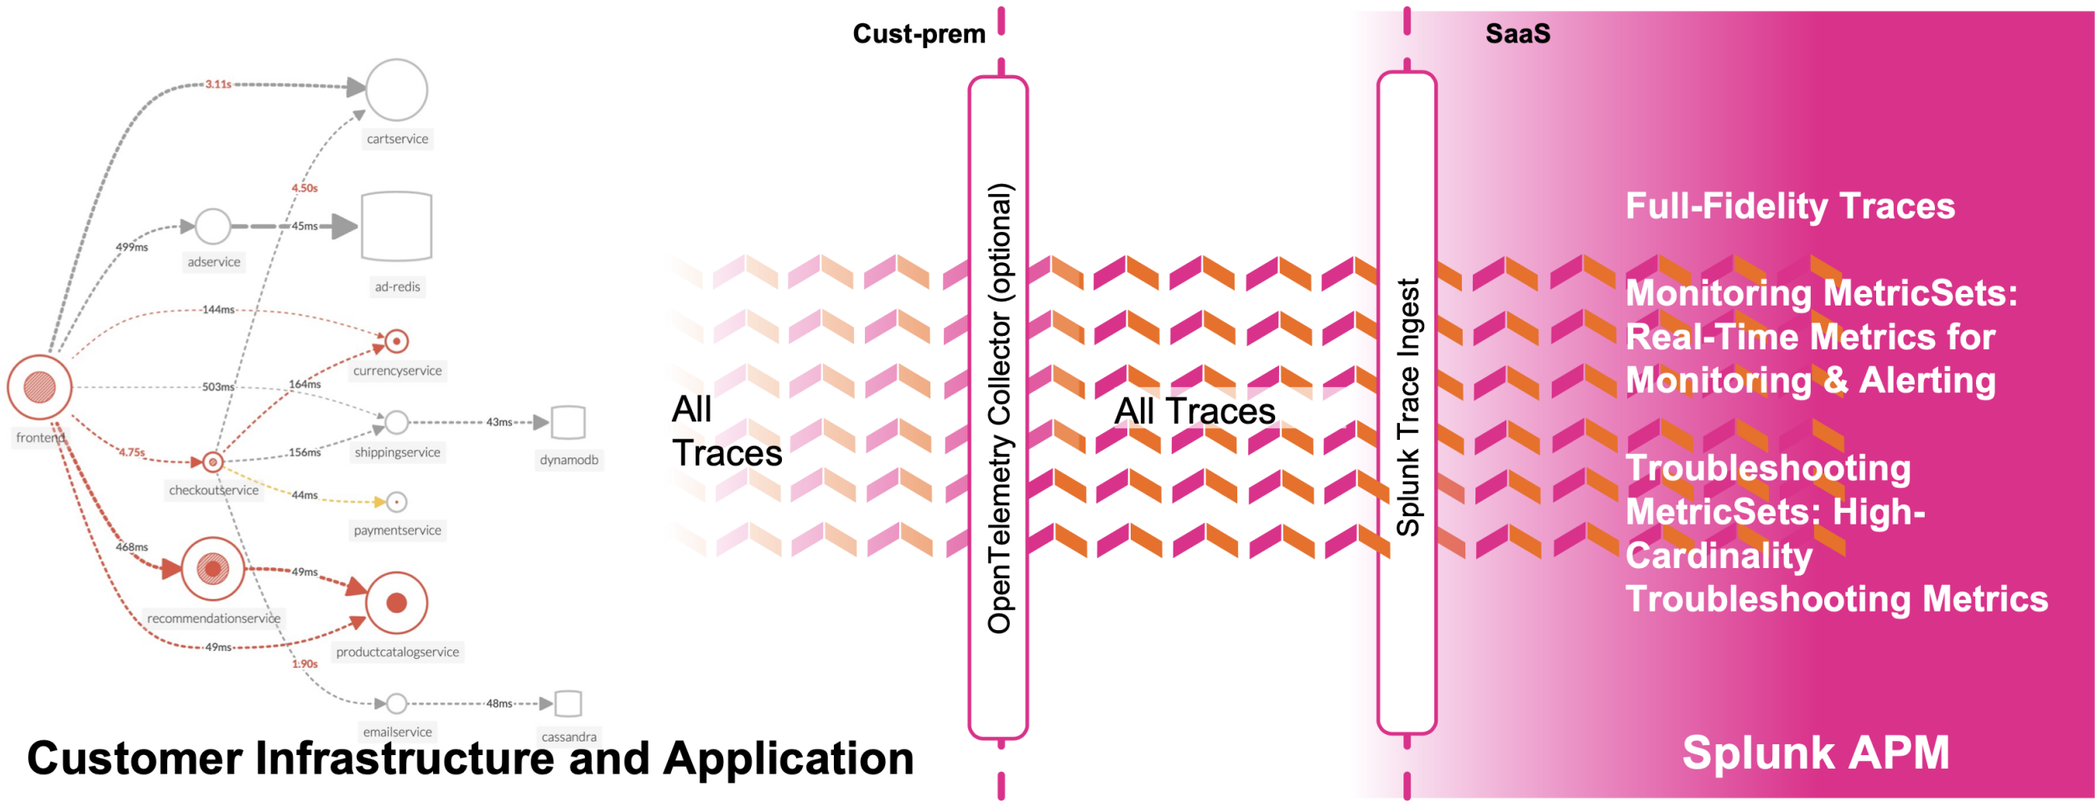

Traces

- Identified by a unique `traceId`
- A collection of spans (represents unique transaction handled by app and constituent services)

Trace Search: Classic

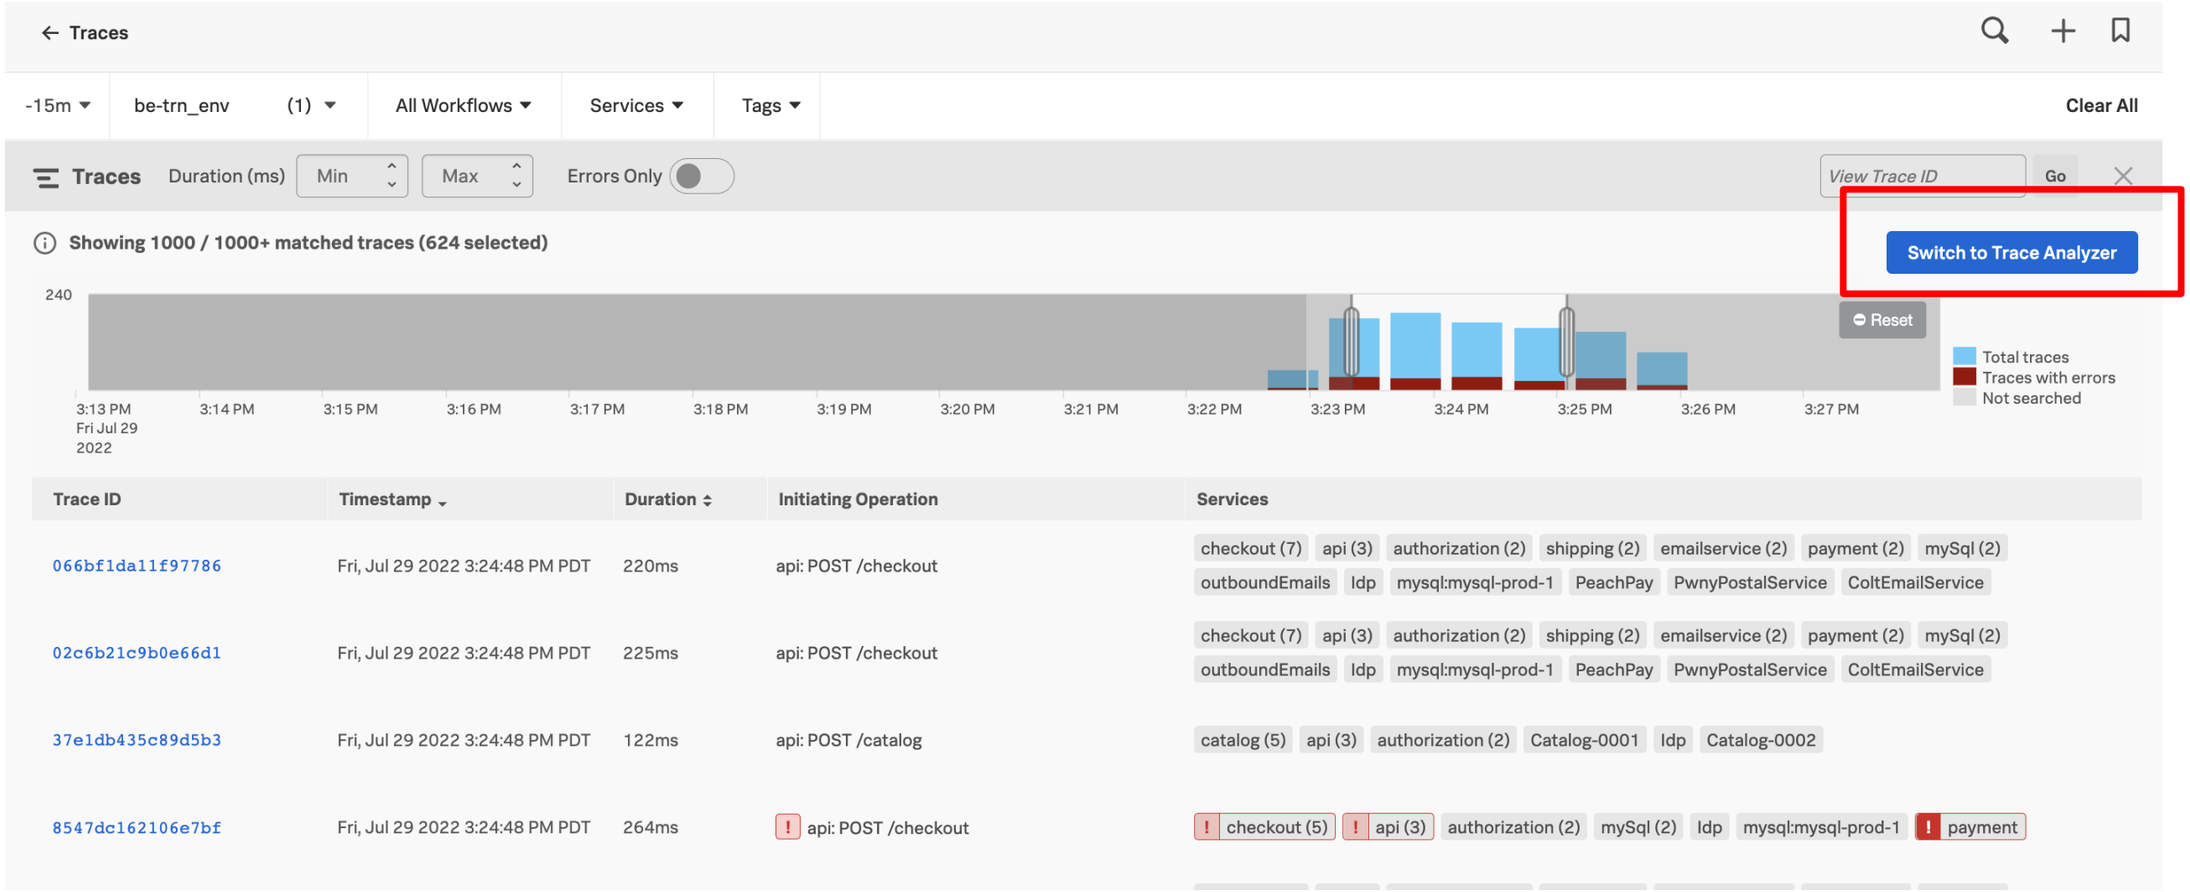

Trace Analyzer

- Similar look and feel to classic view with improved functionality
- Lets you search traces and identify patterns in the full-fidelity trace data without prior knowledge of which tags are relevant

Trace Analyzer (cont.)

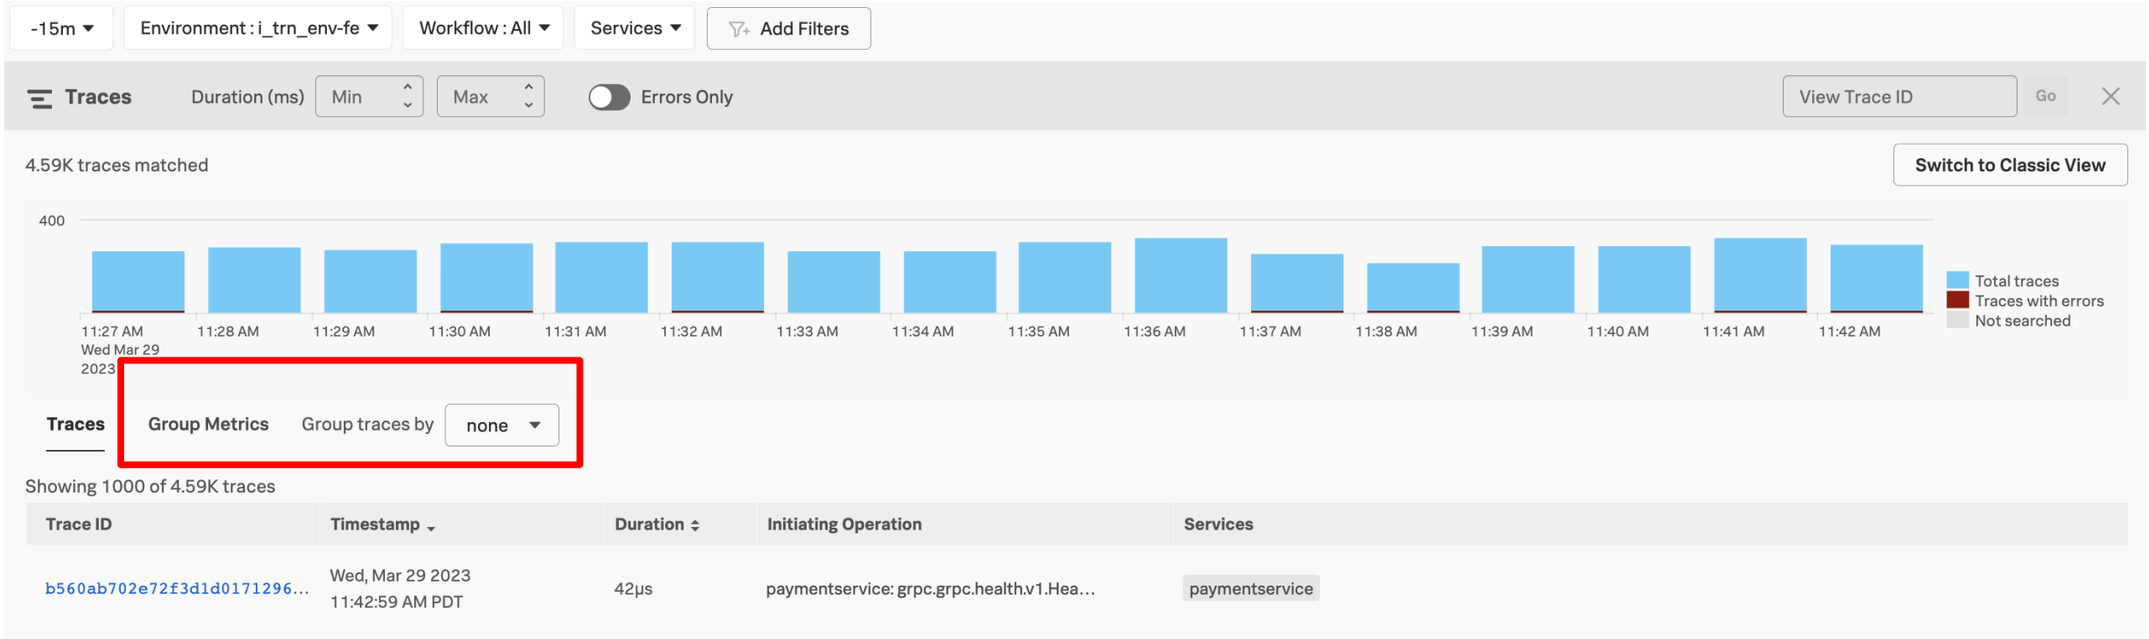

Spans

- Each span has a name representing:
  - Service where the operation took place
  - The operation captured by this span
- Each span contains the following information:
  - Service name
  - Operation name
  - Start time of operation (microsecond precision)
  - Logical name of service on which the operation took place
  - IP address of service instance on which the operation took place

Spans (cont.)

- Spans can have parent/child relationships
- Spans form a tree representing all operations involved in the handling of the request, their sequencing and relationships
- Spans also carry metadata as key/value pairs
  - Example: http.method, http.url

Waterfall View

- Call graph
- Displays all spans in the trace
- Total time for each operation

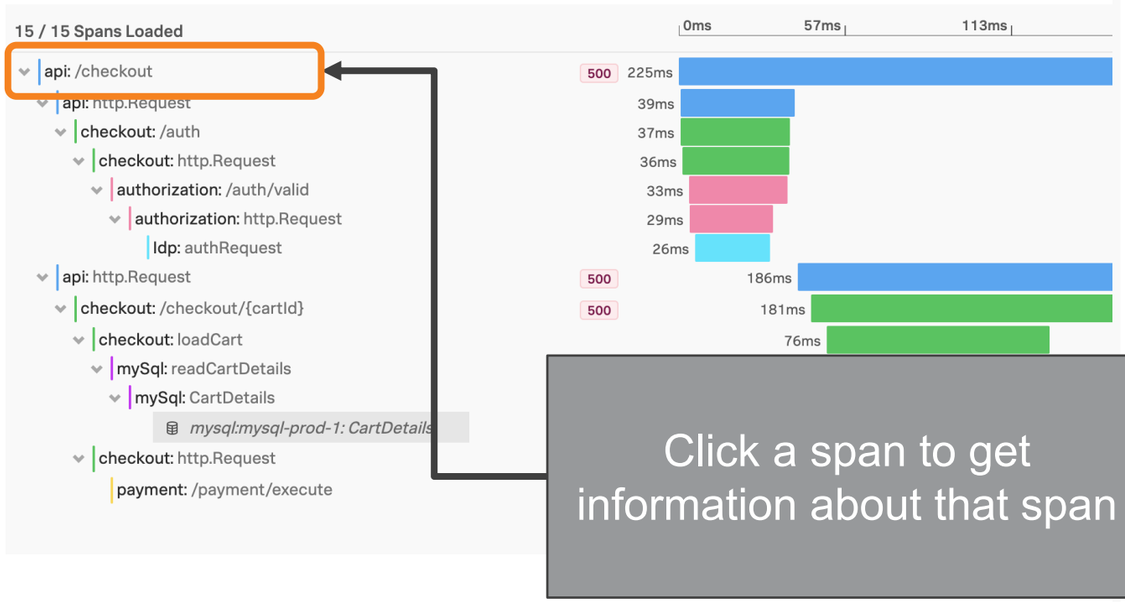

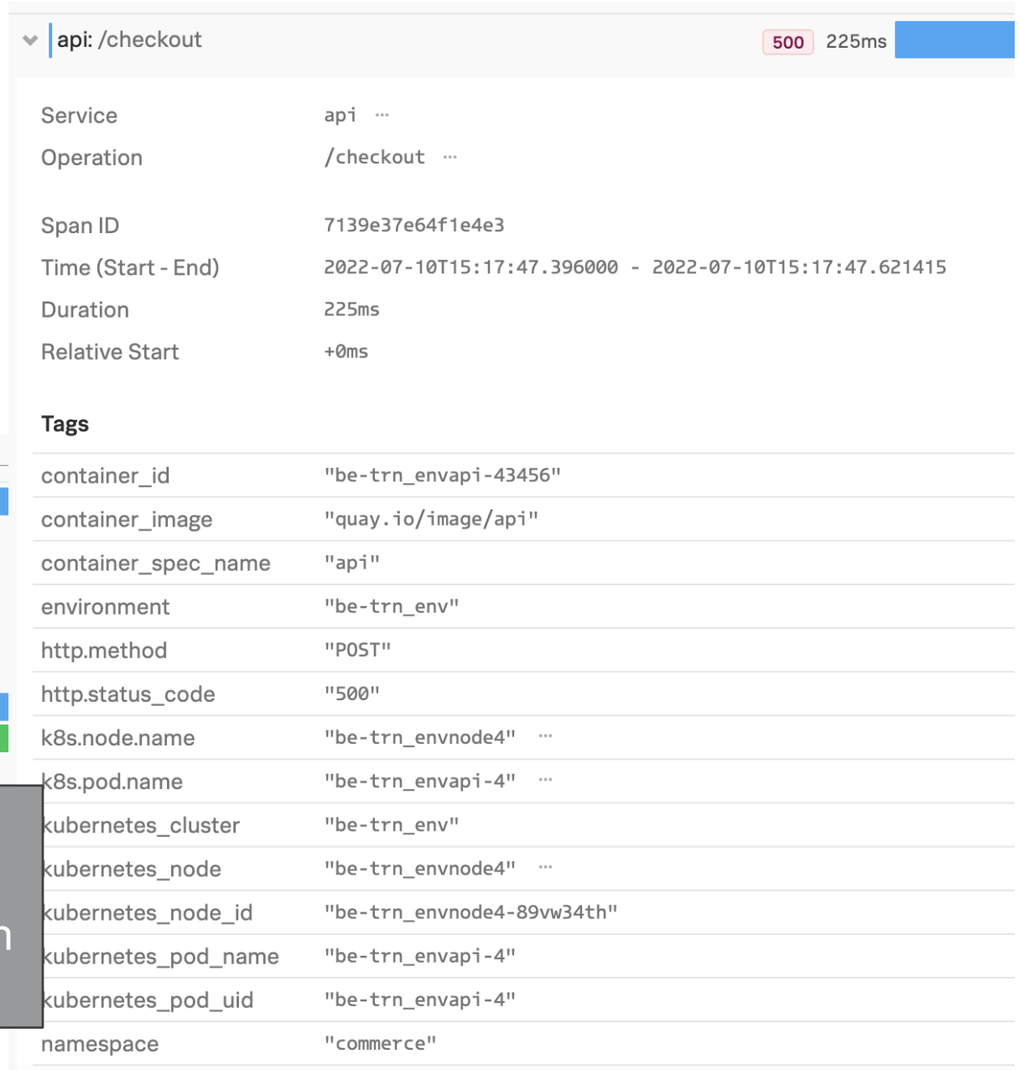

Span Performance

Tells you where time is spent within an individual trace

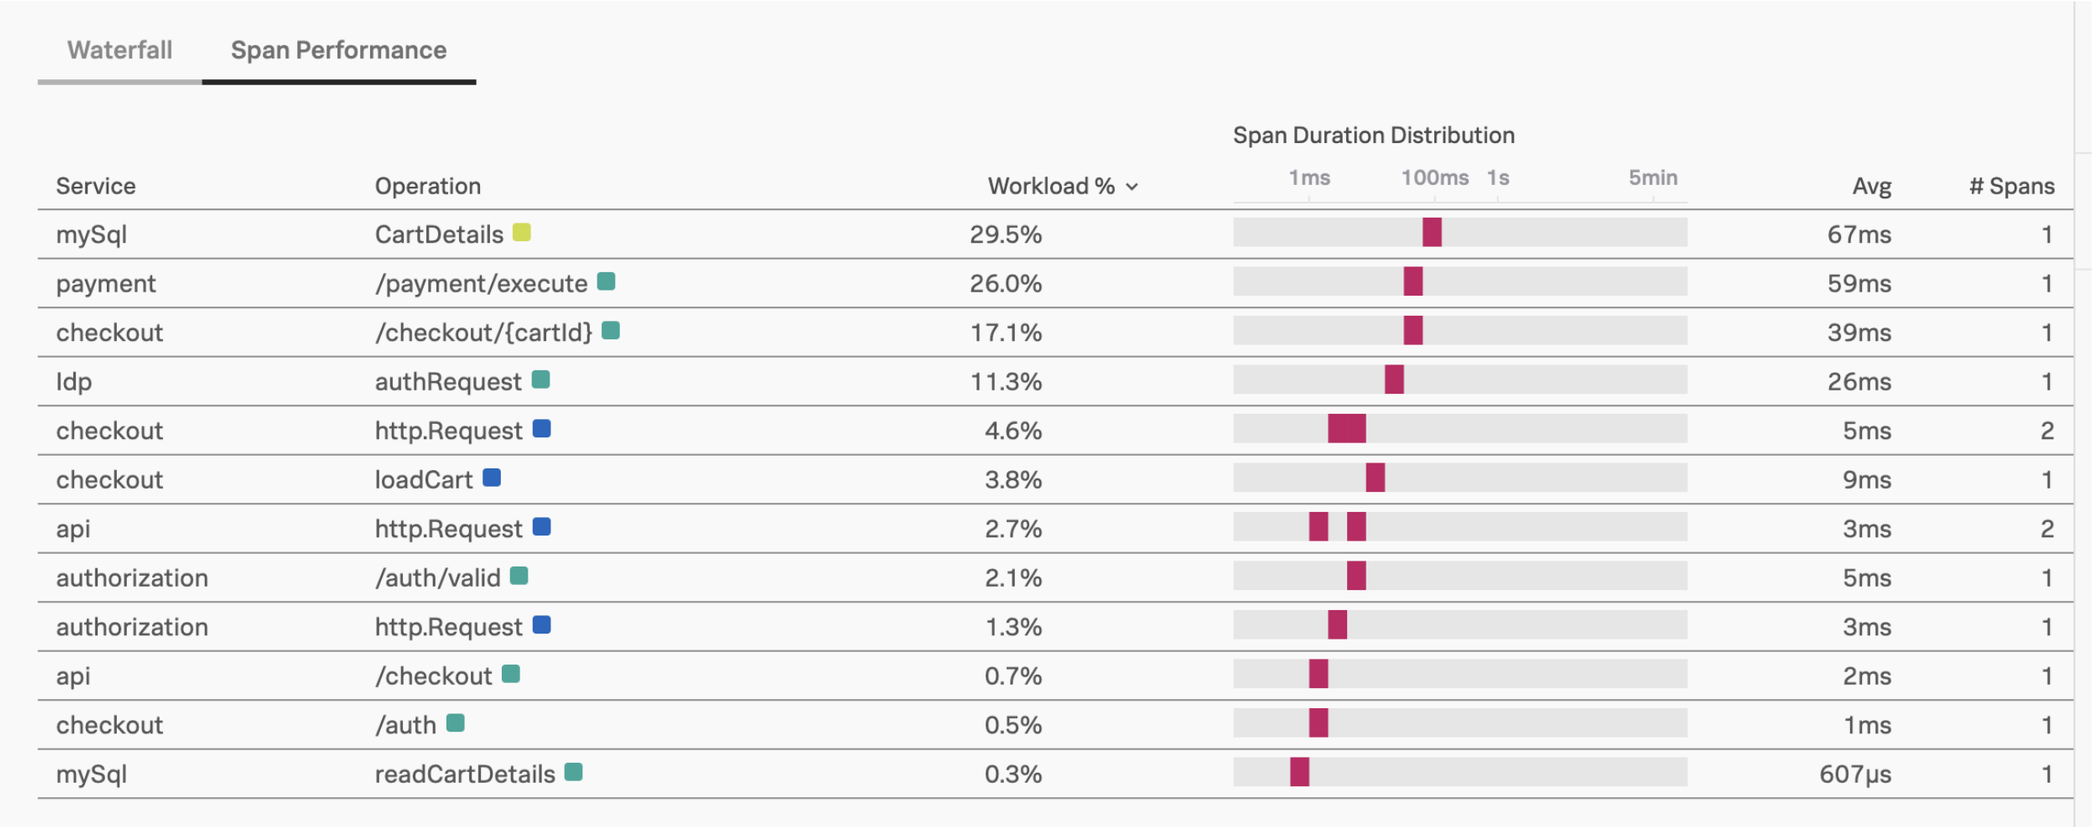

Service	Operation	Workload %	Avg	# Spans
mySql	CartDetails	29.5%	67ms	1
payment	/payment/execute	26.0%	59ms	1
checkout	/checkout/{cartId}	17.1%	39ms	1
ldp	authRequest	11.3%	26ms	1
checkout	http.Request	4.6%	5ms	2
checkout	loadCart	3.8%	9ms	1
api	http.Request	2.7%	3ms	2
authorization	/auth/valid	2.1%	5ms	1
authorization	http.Request	1.3%	3ms	1
api	/checkout	0.7%	2ms	1
checkout	/auth	0.5%	1ms	1
mySql	readCartDetails	0.3%	607µs	1

Span Metadata

- Additional metadata that provides information and context about operations they represent
- Can be used to query and filter traces that contain desired metadata
- Provides extra information about each operation when inspecting spans for troubleshooting
- Defined as span tags

Span Tags

- Span tags are sets of key-value pairs
- Tag keys must be unique within a given span
- Both keys and values are text strings (mostly free-form)
- Standard tags are included by default
  - http.method, http.url, db.statement
- Developers can add custom tags through code instrumentation

Span Tags Added by OTel Collector

- Trace spans emitted by the application are sent to the Splunk OTel collector
- The OTel Collector automatically adds tags to every span captured on a host
- These tags identify on which piece of infrastructure each trace span was executed
- These tags include:
  - host tag
  - AWSUniquieID or gcp_id
  - container_id (Docker) or kubernetes_pod_uid (Kubernetes cluster)…

Indexed Span Tags

- You can index span tags in Splunk APM
- You can use indexed tags to drill down on the performance of your services
- Indexed span tags are available in different Splunk APM pages that can be used to “filter” and in “breakdowns”
- Indexing span tags generates Troubleshooting MetricSets (discussed later)

Indexing Tags

- Only Admins can index span tags
- **Settings > APM MetricSets > New MetricSet**
- Cardinality contribution analysis calculates potential cardinality of indexing this tag which
  - Helps account for any limits

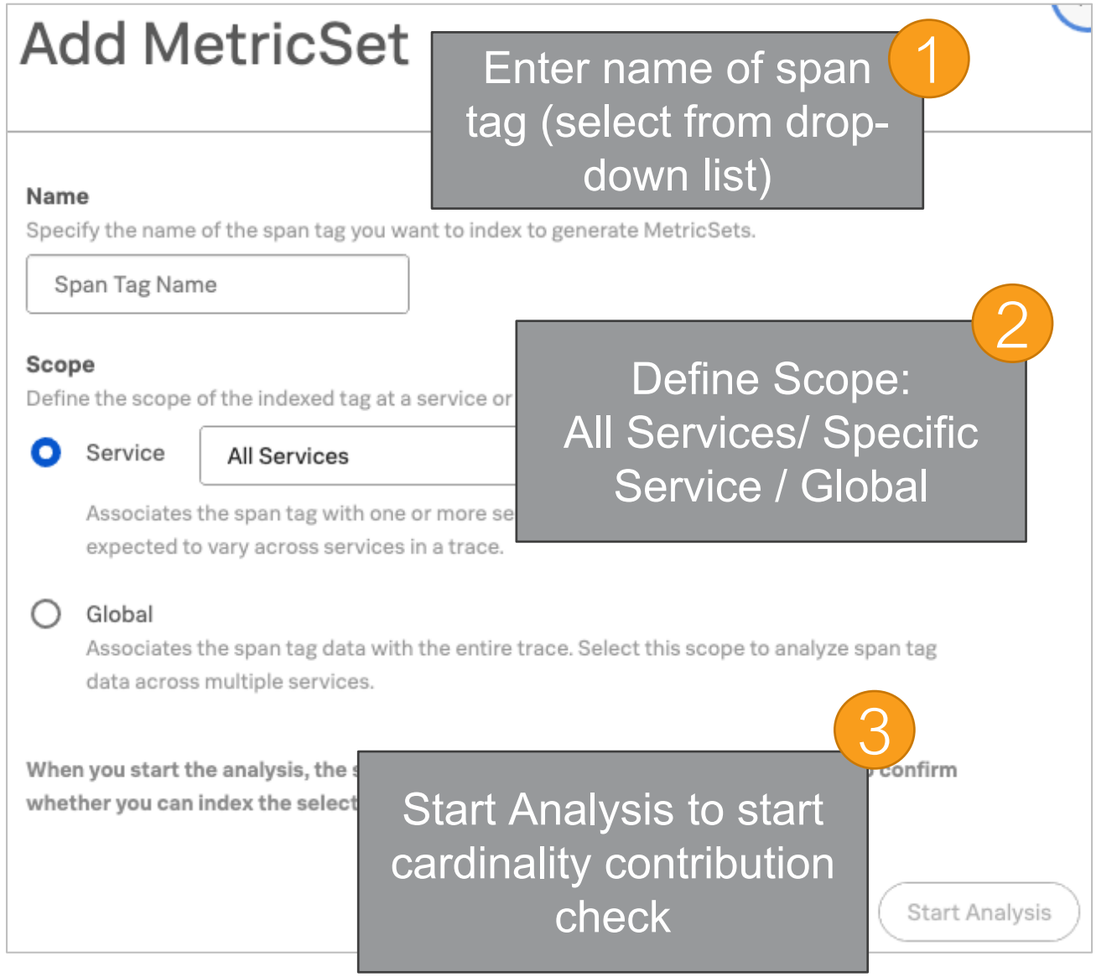

Managing Indexed Span Tags

- Custom MetricSets have one of the the following statuses:
  - Active, Paused, Stopped
- Once span tags have been indexed, you can:
  - Edit the MetricSet configuration
  - Pause or stop the MetricSet generation
  - Delete the MetricSet configuration
- Only a Splunk Observability Cloud Admin can perform these actions

Limits on Indexing Span Tags

- There is a limit to the number of span tags that you can index
- It depends on the number of Troubleshooting MetricSets that are generated from indexed span tags
- To determine total potential number of Troubleshooting MetricSets for each service:

  (# of tag values indexed tag1)*(# of tag values indexed tag 2)* ….
- Total potential number of Troubleshooting MetricSets = sum of Troubleshooting MetricSets across all services

Indexed Tags

- Use span tags to break down your SLO by multiple variables
  - Examples: devices, regions, version, customer level
- Defined as dimensions that help you narrow down causality
- Index only those span tags you want to drill down into to gain insights about performance or to address a specific incident

Indexed Tags (cont.)

- Avoid indexing span tags with a high level of cardinality
  - For example, indexing query_id generates MetricSets for every unique query which is not that useful
- Also avoid indexing span tags that represent ephemeral resources like container_id

Types of Span Tags

Type	Example	Description
Global	Tenant class, application ID	A tag with only a single value across a trace. This associates the span tag value with an entire trace.
All services	Version, host, status code	A tag that exists for most or all services. This associates the span tag value with services in a trace. The value could change across services in a trace.
Specific service(s)	Job ID, database name	A tag that exists for only a single service in a trace. This associates the span tag value with a single service in a trace.

Deciding What to Index

Answer these questions to help you decide which tags to index.

- Are there any attributes that you look at when an incident occurs?
- Do you run multiple versions or builds of the code at the same time?
- Do you deploy services in multiple regions or domains?
- Do you monitor multiple products?

Examples of Indexed Tags

- Infrastructure dependencies
  - Availability zone, host, region, Kubernetes cluster, node, pod
- Application-specific
  - Controller in Ruby, mvc controller in .net, job.id, normalized queries for database calls
- Business logic
  - Tenant class, device type, code version

Unindexed Tags

- You can use unindexed tags to filter as well
  - However, no visual analysis will be available
  - Instead, you are taken directly to the list of traces

Tag Spotlight

- Use to analyze performance of services
- For every service, provides a RED metrics time-series chart
- Displays the total number of requests, errors, root-cause errors, and latency
- Displays the total number of requests, errors, root-cause errors, and latency for every value of an indexed span tag
- See all indexed tags in one place
- Analyze all tags on Service B, limited to requests from Service A

Using Tag Spotlight

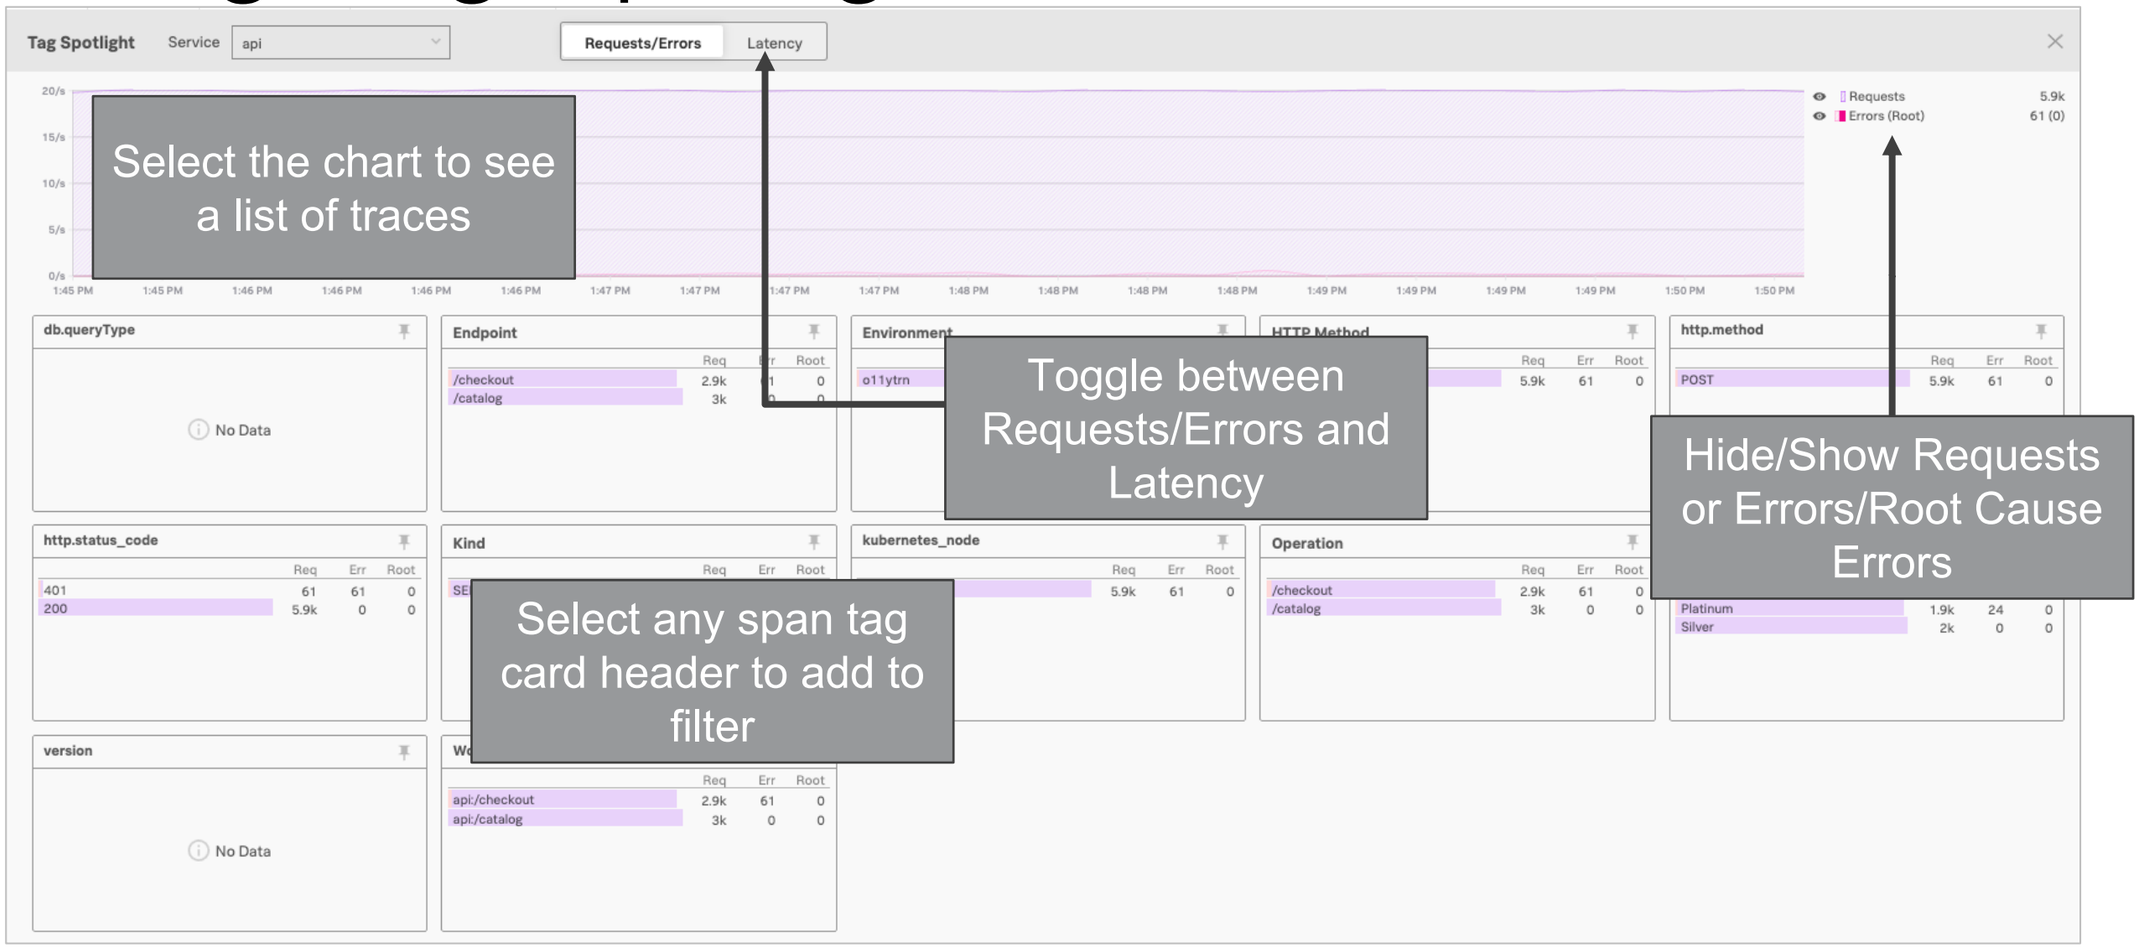

Using Tags to Filter and Apply Breakdowns

- In the Explore page, you can filter on tags
  - Click Tags
  - Select the tag on which you want to filter and the value
  - The filter is applied to the service graph and charts on the right
- You can apply breakdowns to see how each segment within a tag is doing
  - For example, how a tenant type is impacted
- Breakdown provides 100% accurate SLIs for a given tag value within the service

MetricSets

- In addition to traces, RED (rate, error, duration) metrics are also saved
- Monitoring MetricSets
  - Used for real-time monitoring and alerting
- Troubleshooting MetricSets
  - Used in the Splunk APM User Interface for filtering service-graphs and breaking down SLIs to enable historical comparison for spans and workflows

Monitoring MetricSets

- Used for monitoring and alerting
- Created out-of-the-box for combinations of
  - Service
  - Endpoint
  - Workflow

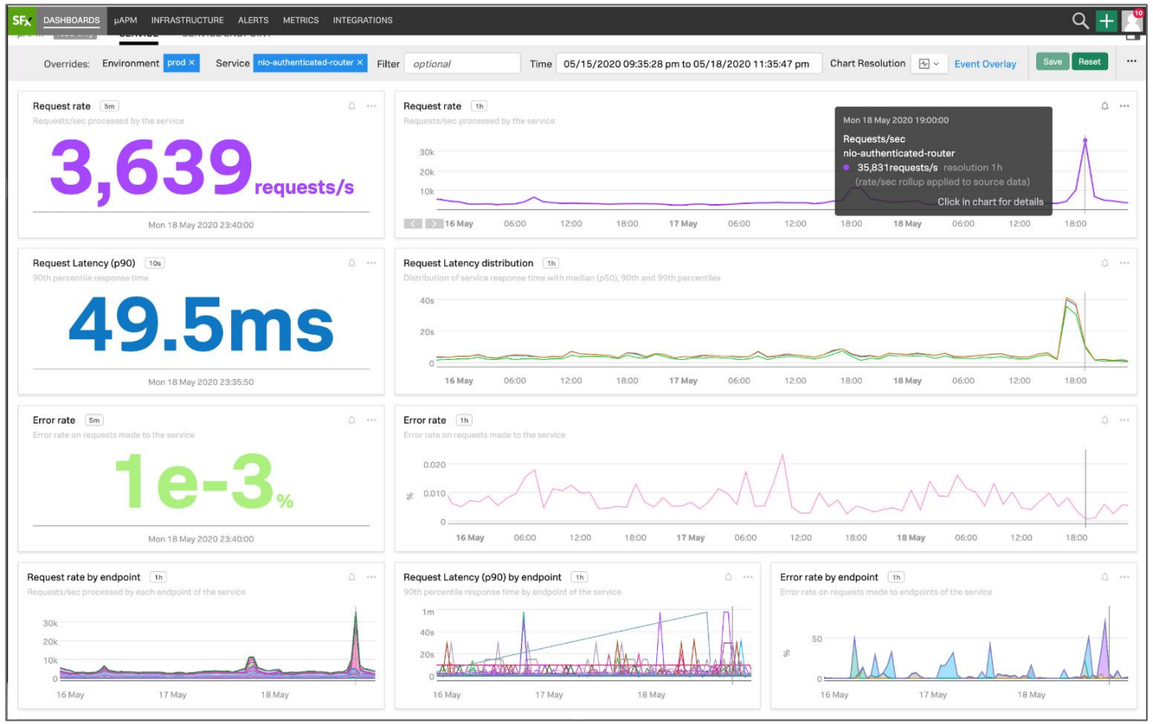

Monitoring MetricSets (cont.)

- Each MetricSet contains the following metrics:
  - Request-Rate
  - Error-Rate
  - Latency: Min, Max, P50, P90, P99
- Stored for 13 months by default

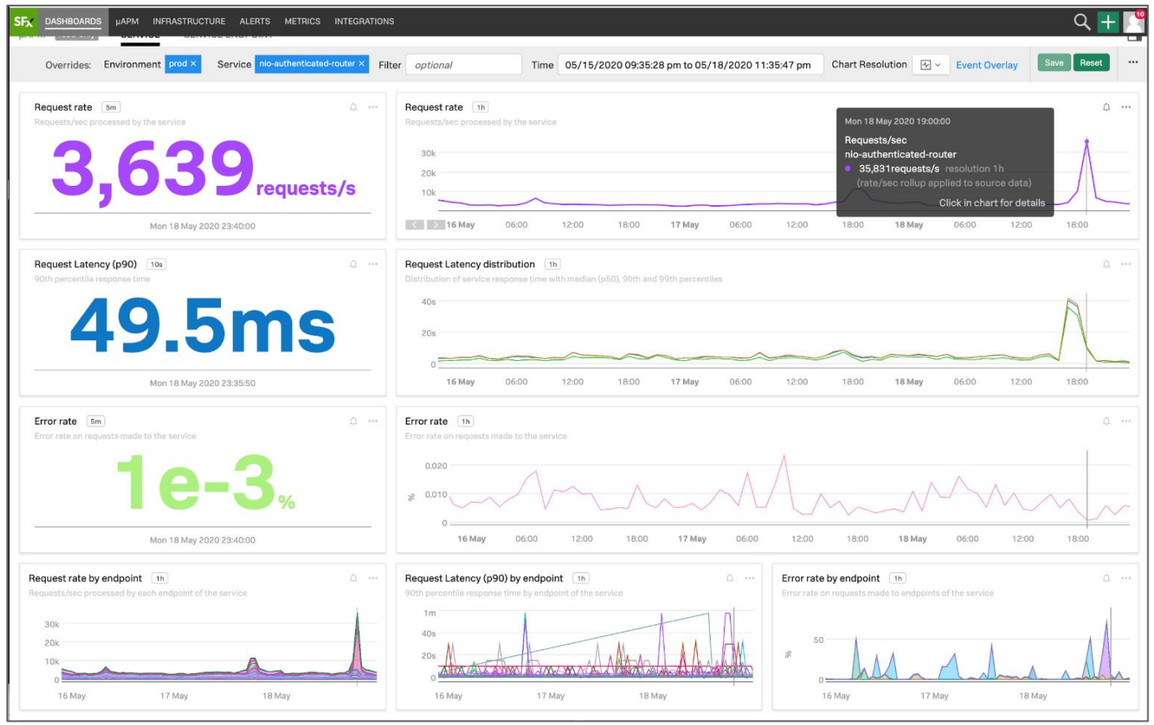

Monitoring MetricSets (cont.)

- Used to populate Splunk APM built-in dashboards and APM Overview page
- If there is no span.kind of type SERVER or CLIENT, no dashboard created/populated
- span.kind describes the relationship between the Span, its parents, and its children in a Trace

Business Workflow (cont.)

- A Business Workflow is the start-to-finish journey of the collection of traces associated with a given activity or transaction
- Provides a way to easily group relevant traces based on:
  - Their initiating operation
  - Any other tag/endpoint on a downstream service

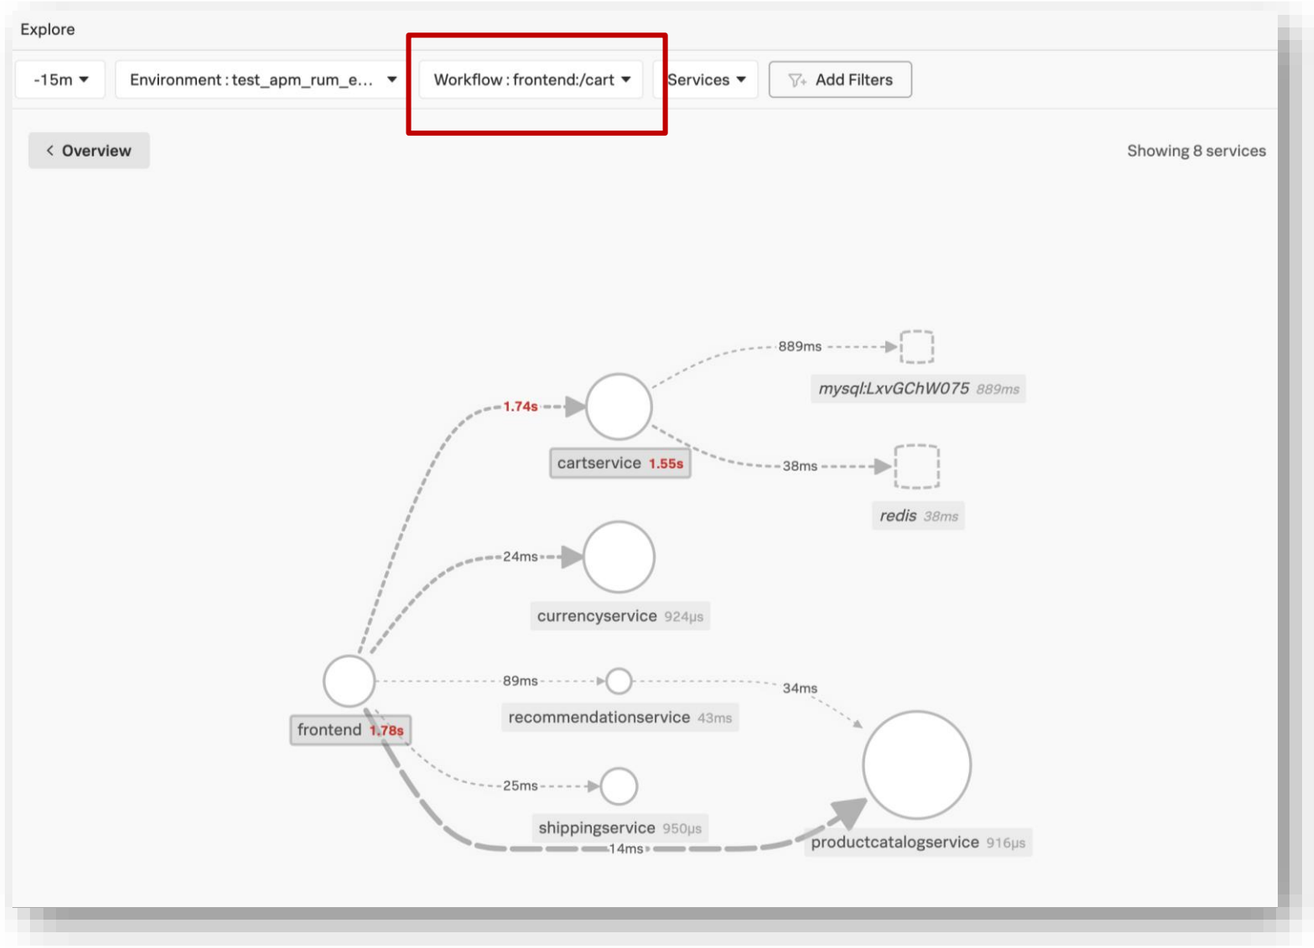

Business Workflow

- Enables monitoring of end-to-end KPIs and identifying root causes and bottlenecks
- Only a Splunk Observability Cloud Admin can add business workflows (Setting up workflows is out of the scope of this class)

Business Workflows Use Cases

- Business/engineering executive
  - Monitor end-to-end RED metrics for most critical transactions
  - Want to ensure you are meeting associated KPIs
  - “What is the p90 of the logins duration?”
  - “What percentage of checkout requests return an error?”
- SRE/Service Owner
  - Troubleshoot an issue
  - Narrow down context to only services involved in the transaction of interest
  - “Which downstream service is responsible for increased response of the API class?”
  - “Are the errors in the Login service only happening for API calls from the homepage?”

Business Workflows Page

- Overview page shows the Top Business Workflows by error rate and duration (p90)
- Get more information about workflows from the Business Workflows tab
  - APM Overview > Business Workflows
  - View workflows sorted by decreasing/increasing latency or decreasing/increasing requests and errors
  - To see the behavior for a workflow over a past time period, open the dashboard

Business Workflow Overview Dashboard

- Overview dashboard gives you the top workflows
  - By requests
  - By duration
  - By error rate
- You can filter the dashboard to view workflows for a specific environment

Workflow Details Dashboard

- Filter the Workflow Details dashboard for a particular workflow
- View details such as
  - Total requests
  - p90 of workflow duration
  - Workflow error rate

Traces, Spans, and Tags Scenario

- Customers are complaining about slowness on item pages with ads
- Use the trace and span views to investigate where the latency issue is occurring
- Use the tag spotlight and breakdowns to explore the issue in more detail

Exploring Traces and Spans Steps

- Open the service map for the service with high latency (ad service)
- Select Traces
  - Filter for error traces
  - Switch to Trace Analyzer view
- Select an error trace and open the span waterfall view
- View the span details

Exploring Tag Spotlight and Breakdowns Steps

- Open the service map for the service with high latency (ad service)
- Select different tags from the breakdown menu on the service map
- Navigate to the Tag Spotlight page
- Investigate the tag details on the tag spotlight page

Module 3 Lab Exercise 3

Time: 20 minutes

Description: Use trace and spans views combined with breakdowns, tag spotlight, and business workflows to investigate further

Tasks:

- View Traces and Spans
- Apply breakdowns
- Analyze in the Tag Spotlight page
- Identify business workflow with high latency
- View traces for a business workflow

Take a Break

Module 4: AlwaysOn Profiling, DB Query

Module Objectives

- Define CPU and memory profiling
- Explore the following to troubleshoot latency and errors in the ad and checkout services:
  - CPU AlwaysOn Profiling
  - Memory AlwaysOn Profiling
  - Database Query Performance

Splunk AlwaysOn Profiling

- Continuously collects data from applications and services
- Periodically collects CPU snapshots in the form of stack traces from Java, Node.js, and .NET
- Stack traces are linked to spans and traces in Splunk APM
- Enables visibility into code behavior in production
- Helps identify performance and resource bottlenecks

Splunk AlwaysOn Profiling in APM

- As part of Splunk APM, AlwaysOn Profiling provides context to spans
- You can explore stack traces directly from APM
- You can visualize the performance of each component using the flame graph

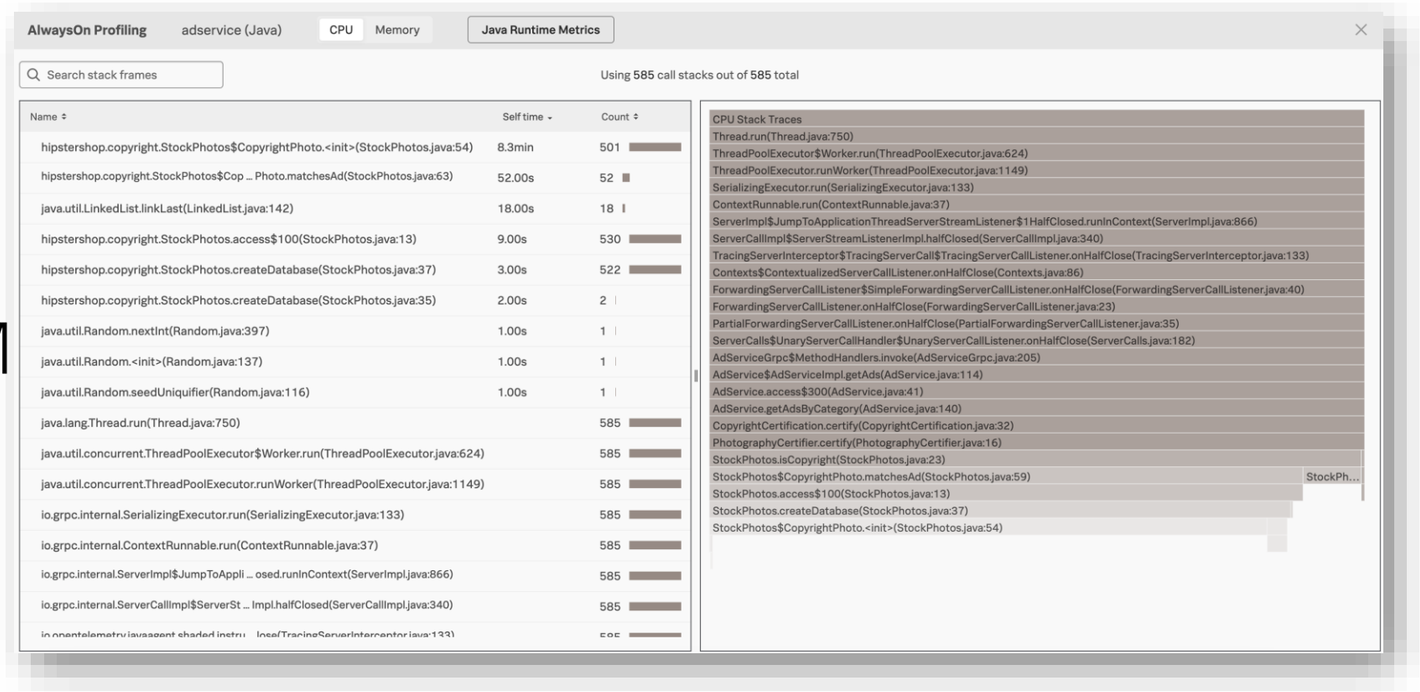

Issues That AlwaysOn Profiling Can Identify

- Slow or inefficient database queries
- Thread locks
- Thread pool starvation
- File system bottlenecks
- Slow calls to external services

Memory Profiling

- Memory profiling adds memory allocation data to stack traces and exposes runtime memory metrics
- Exposes memory metrics for your application, which you can use to build charts and dashboards

Memory Profiling (cont.)

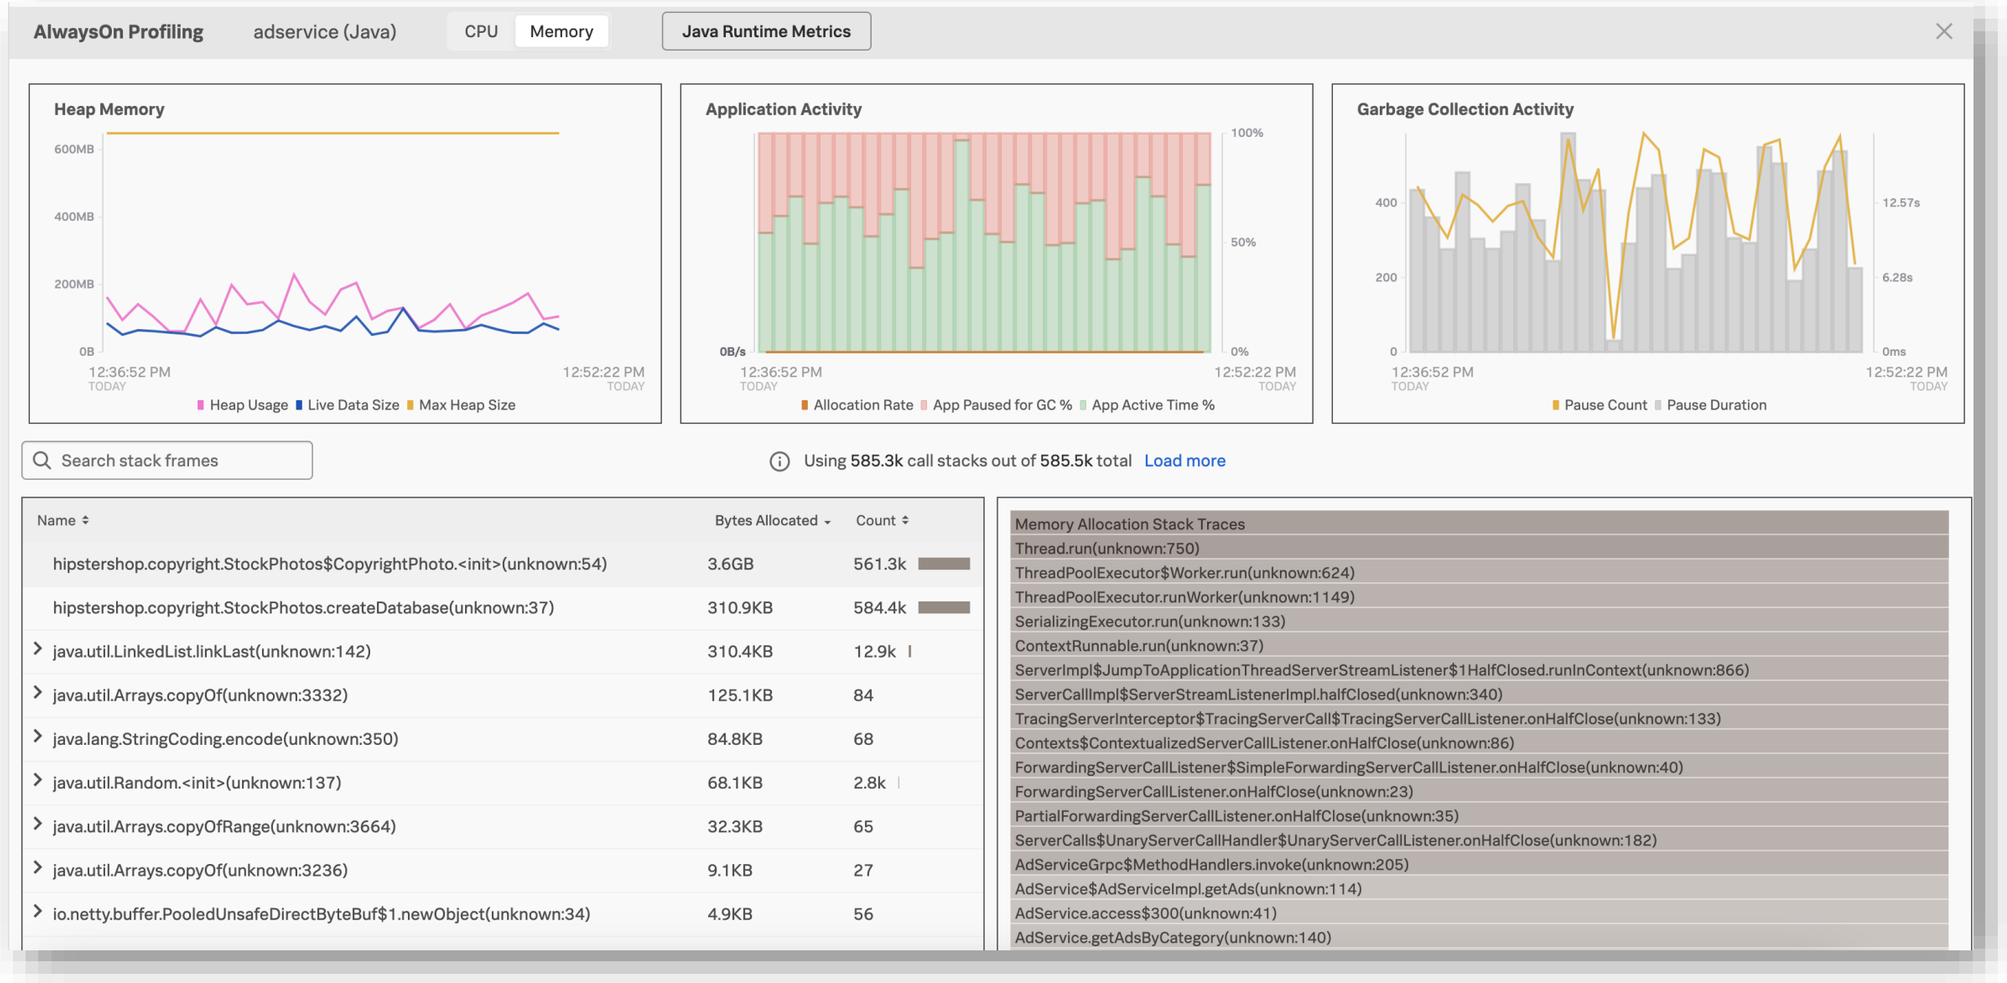

Database Query Performance

- With Database Query Performance you can
  - Pinpoint the database that is causing slowness for your application with no DB instrumentation required
  - See slow and frequently executed aggregate queries, with historical trends
  - Slice and dice in Tag Spotlight to find out what services, endpoints etc. have been affected

Database Query Performance (cont.)

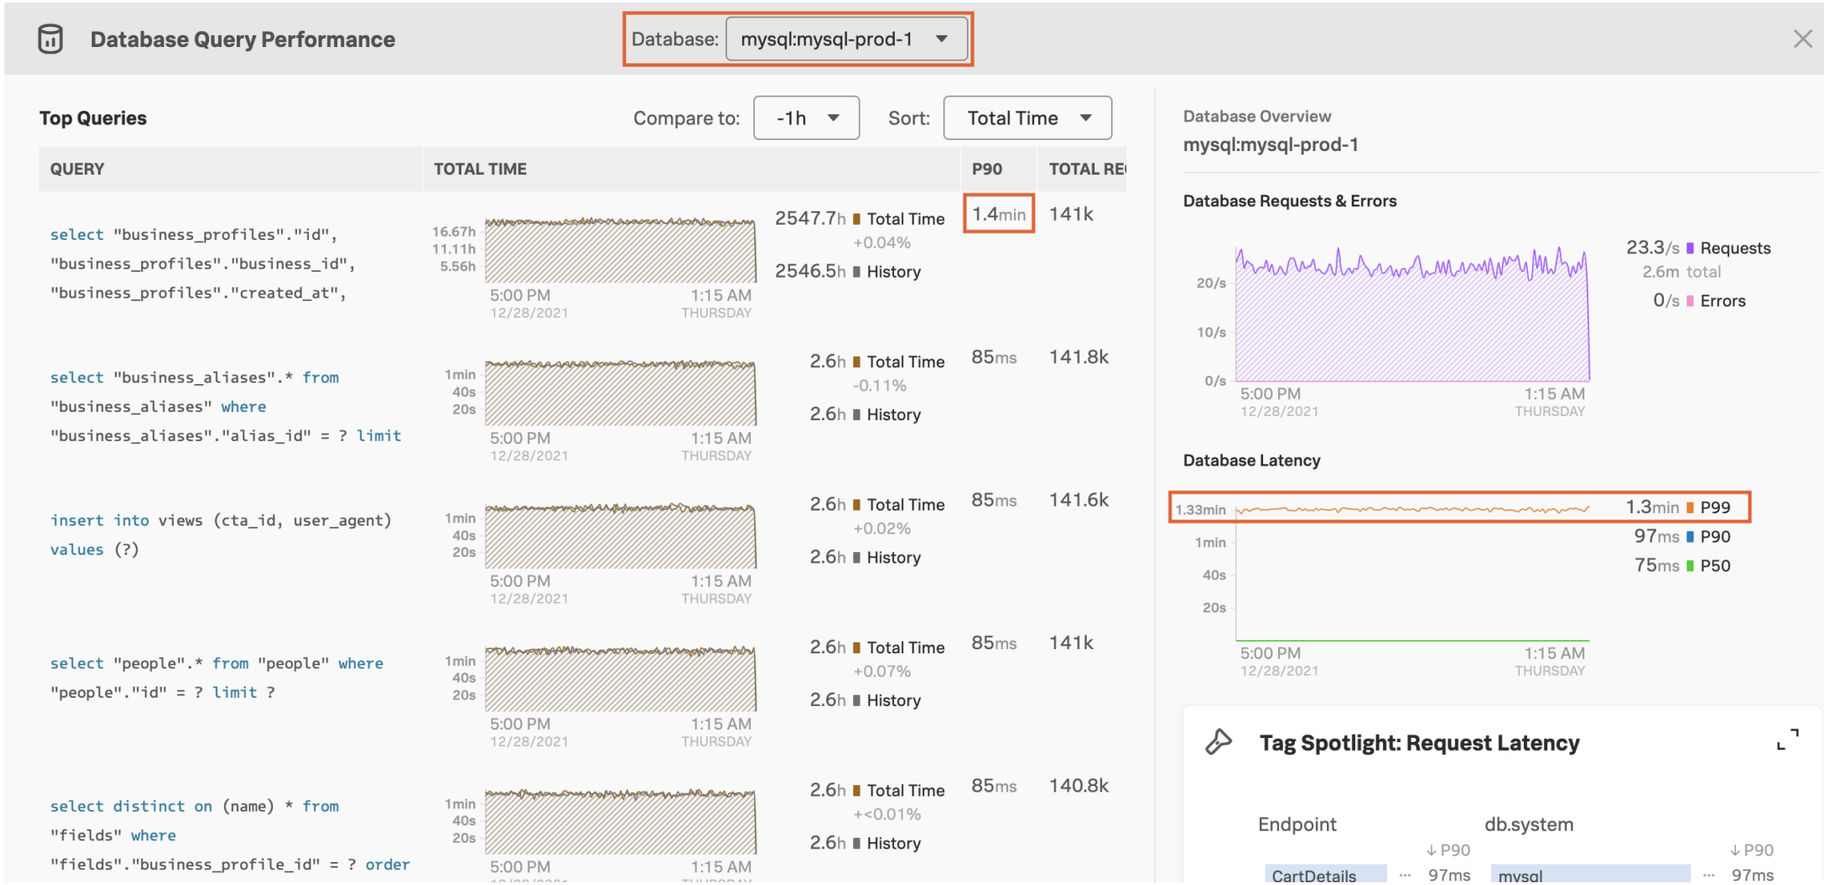

How Database Query Performance Works

- Splunk APM identifies databases as inferred services in your system using automatically generated span tags such as:
  - db.type
  - db.instance
  - db.statement
- Databases appear throughout APM, such as in their inferred locations in the service map, in the service filter list, and in trace view
- Database Query Performance is available by default, so all you need to do is enable the feature

DB Query and Related Content

From the DB Query view you can easily jump to relevant dashboards using related content

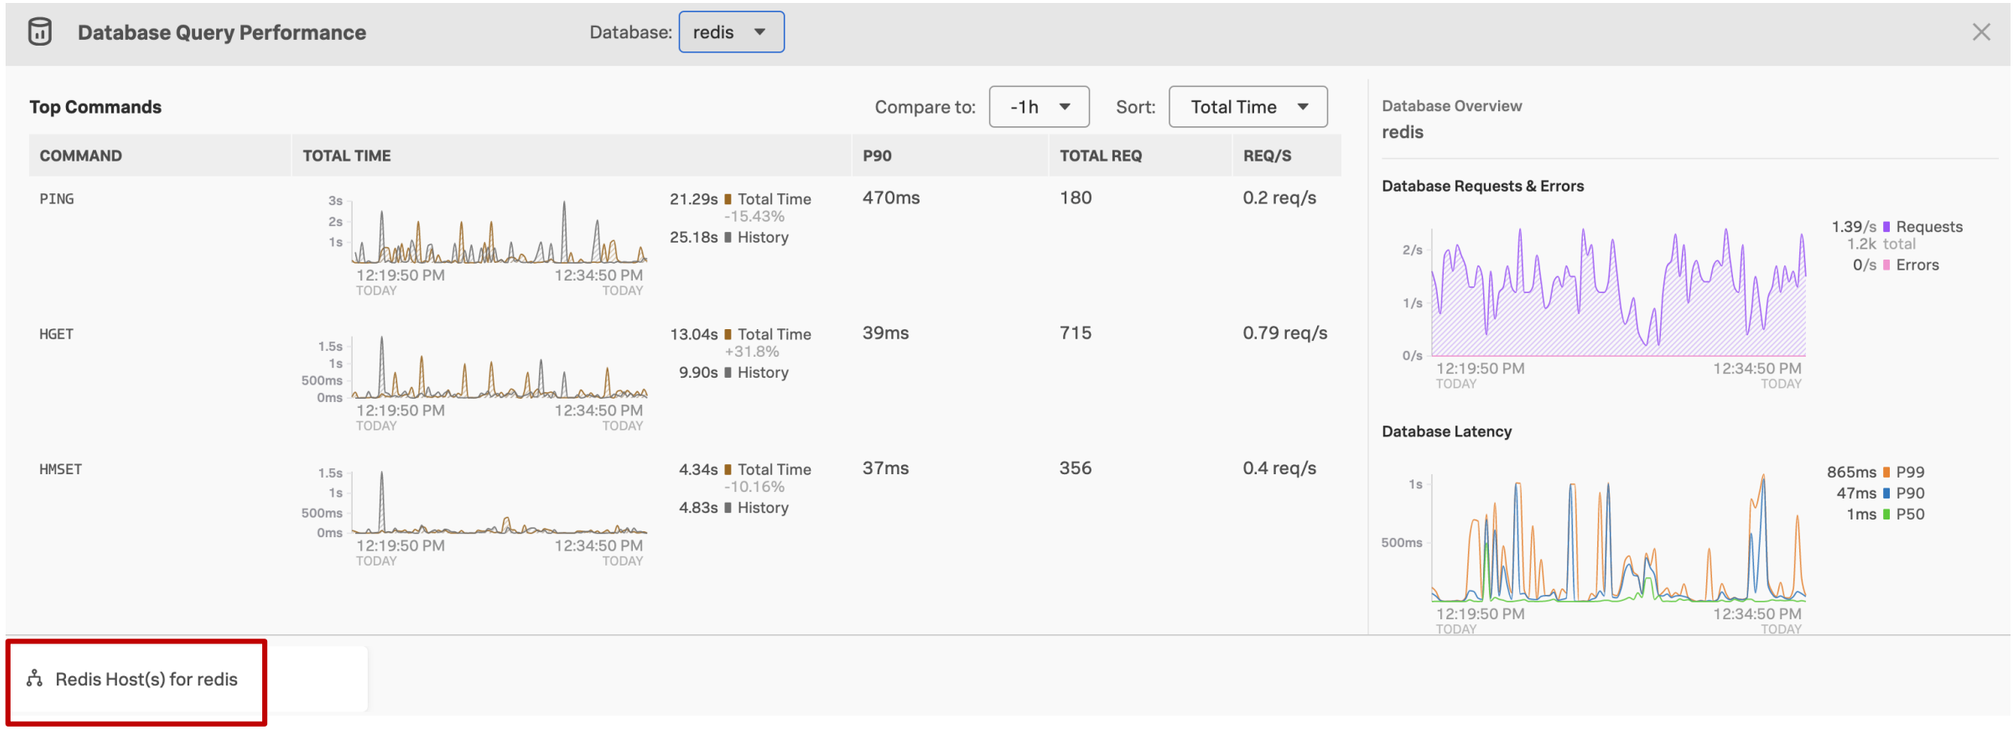

Take a break

Module 5: Troubleshooting Using Splunk APM

Module Objectives

- Create custom dashboards
- Create data links
- Create detectors to alert on Splunk APM metrics
- Troubleshoot an issue after alert triggers

Create Dashboards

- Use the built-in Splunk APM dashboards as templates
- Create a more comprehensive service level dashboard which includes metrics from your application, infrastructure, databases, webservers
- Clone and add charts to these dashboards
- Example: add chart for Week over Week (WoW) rate of change in latency

Working with Dashboards

- Creating dashboard group:
  - Create menu > Dashboard Group
  - Enter a name and click Create
  - A blank dashboard with the same name as dashboard group is added to dashboard
- Cloning dashboard:
  - Open dashboard to clone: Dashboard Actions menu > Save As
  - Specify name of clone and destination dashboard group

Working with Dashboards (cont.)

- Applying filters:
  - Applied to all charts on dashboard
- Changing time range:
  - Applied to all charts, except text notes, single-value charts, list charts and heat maps

Adding Charts

- Creating a chart:
  - Create a chart from the Create menu
  - You can also click + New Chart on the dashboard
- Cloning a chart:
  - To copy a chart, select Copy from Chart actions menu on a dashboard OR
  - Select “Save As” from actions menu in the chart builder

Percent Change from Previous Time Period

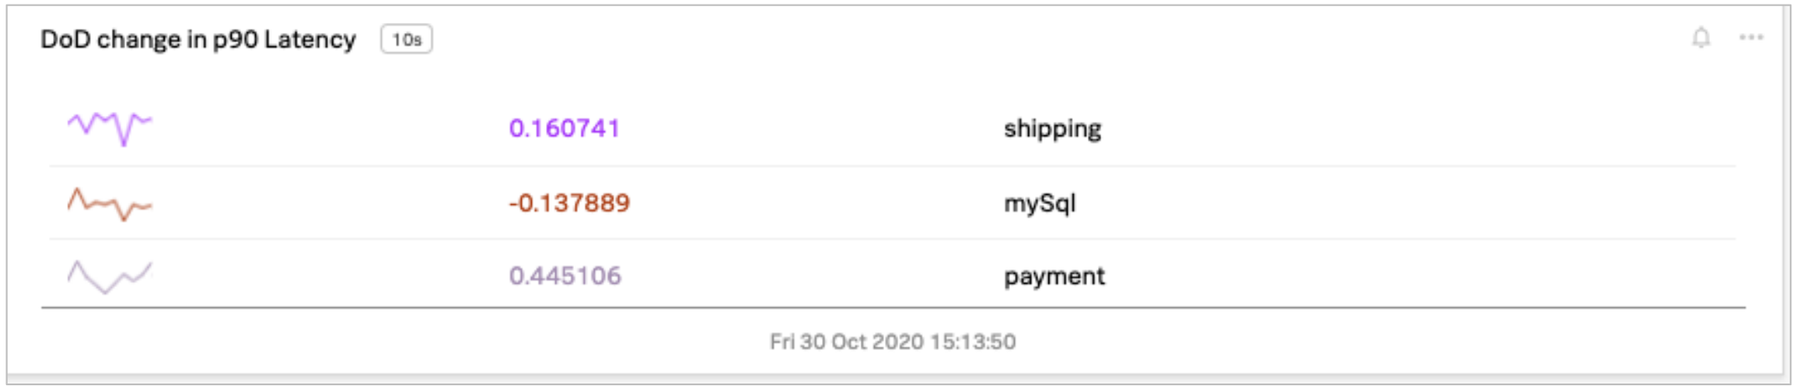

Examples

- Create a chart to view week over week change:
  - Is there a change in your traffic (request rate)
  - Is there a change in latency
  - Is there a change in error rate
- Create charts on metrics from your applications such as number of page views; login errors
- Create charts on metrics from your infrastructure, databases, etc.

Module 5 Lab Exercise 5a

Time: 20 minutes

Description: Create a custom dashboard

Tasks :

- Create dashboard group
- Add DoD percent change in latency chart

Linking Related Content

- You can create links for properties that appear:
  - In a chart's data table and in list charts
- Actions menu next to Services in the Splunk APM pages
- Create links to:
  - A dashboard
  - Splunk instance
  - External system (custom-defined URL)
- Keeps the context of metadata that you selected

Configure Data Links from Splunk APM

- You can configure Data Links from Splunk APM pages
- Only a Splunk Observability Cloud Admin can perform these actions

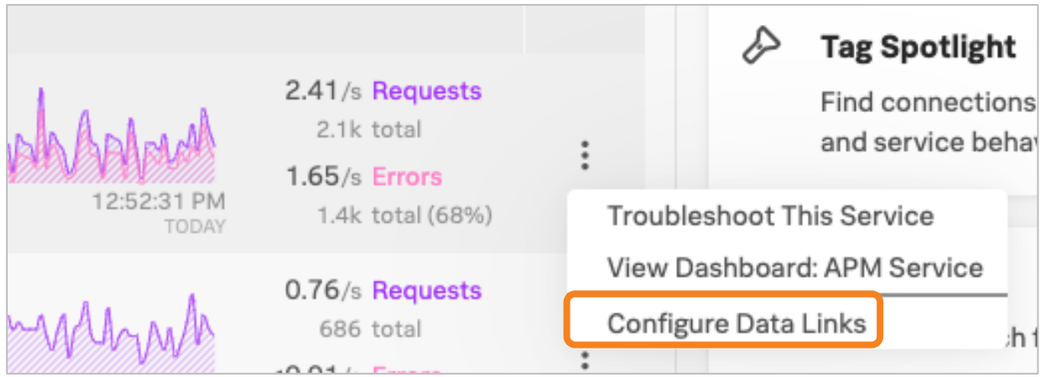

APM > Overview
Services tab > Service actions

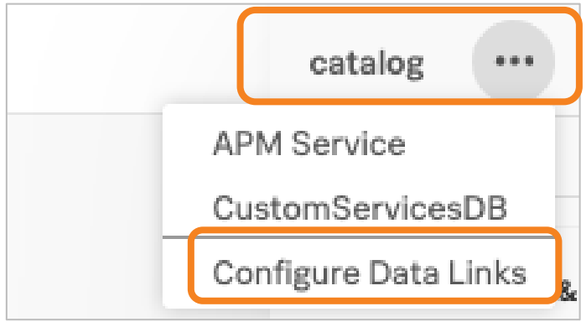

APM > Overview > Explore > Select Service > Service name

Defining The Data Link

Menu > Organization Settings > Global Data Links > New Link

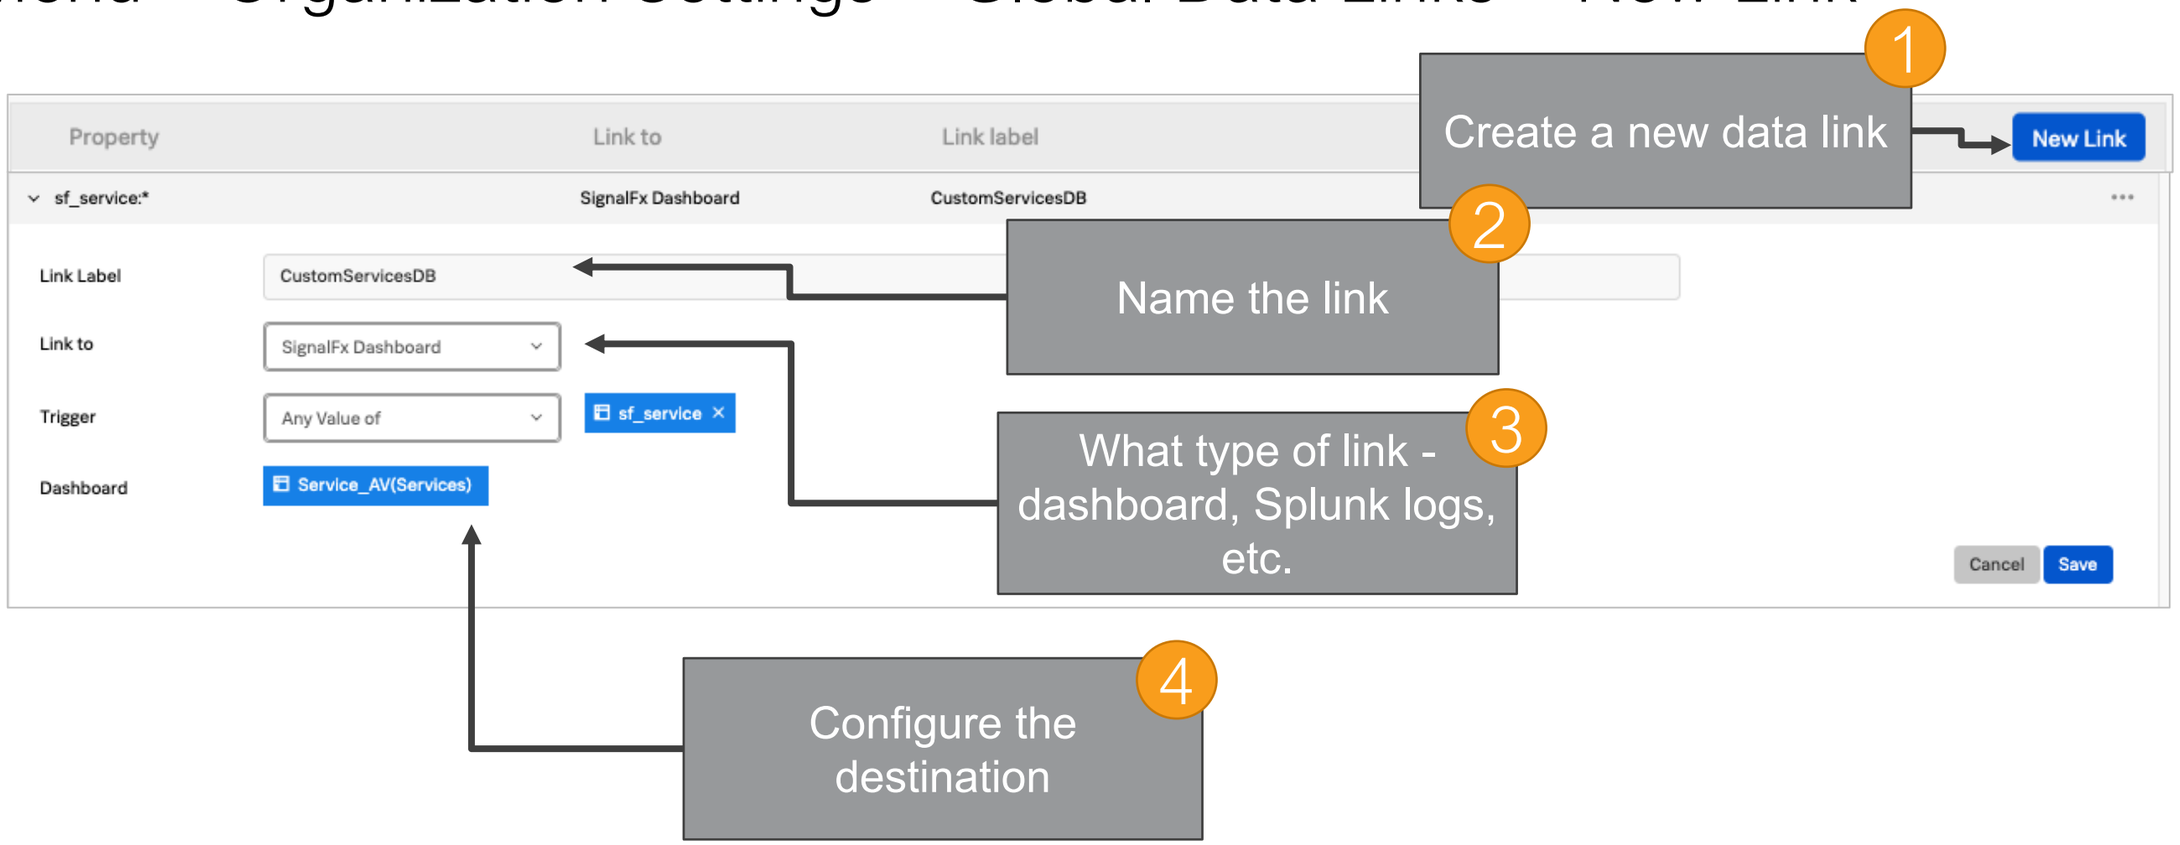

Accessing Data Links

If links have been created, you can access links from Splunk APM > pages

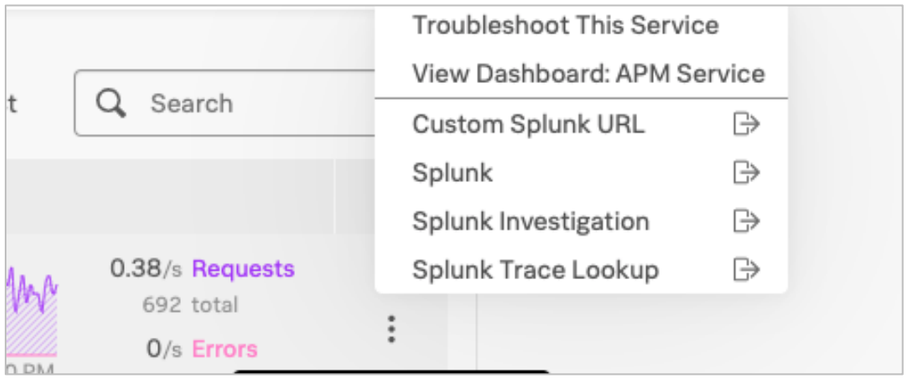

APM > Overview
Services tab > Service actions

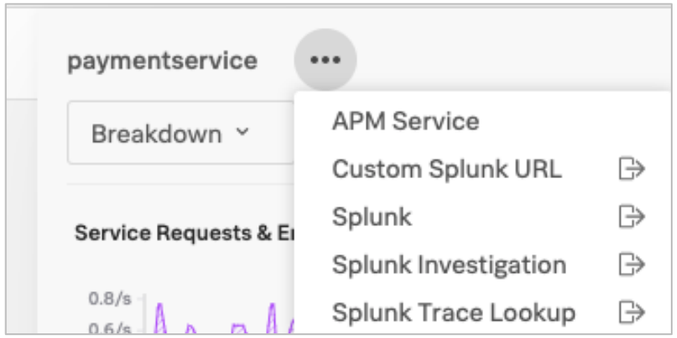

APM > Overview > Explore
Select Service > Service name

If There Are Multiple Links To Dashboards …

- More specific links override less specific links
  - Example: links with property key:value pair overrides links with property:*
- Local links override equally specific global links
- Default links override non-default links

Do I Have a Problem?

- You have deployed Splunk APM
- You have instrumented applications and services to send in traces
- You have defined tags to be indexed
- How do you know if there is a problem?
  - You need to create detectors

APM Detectors

Use built-in algorithms to generate alerts about sudden spikes, historical anomalies, or a static threshold in your APM metrics or Business Workflows

Create an APM Detector Entry Points

- The create menu (+ icon)
  - select APM detector
- The bell icon on a chart in an APM dashboard
  - select new detector from chart
- Tag spotlight in APM three dot menu
  - Select create detector
- APM landing page three dot menu
  - Select create detector
- Service map in Splunk APM select servie then three dot menu
  - Select create detector

Create an APM Detector

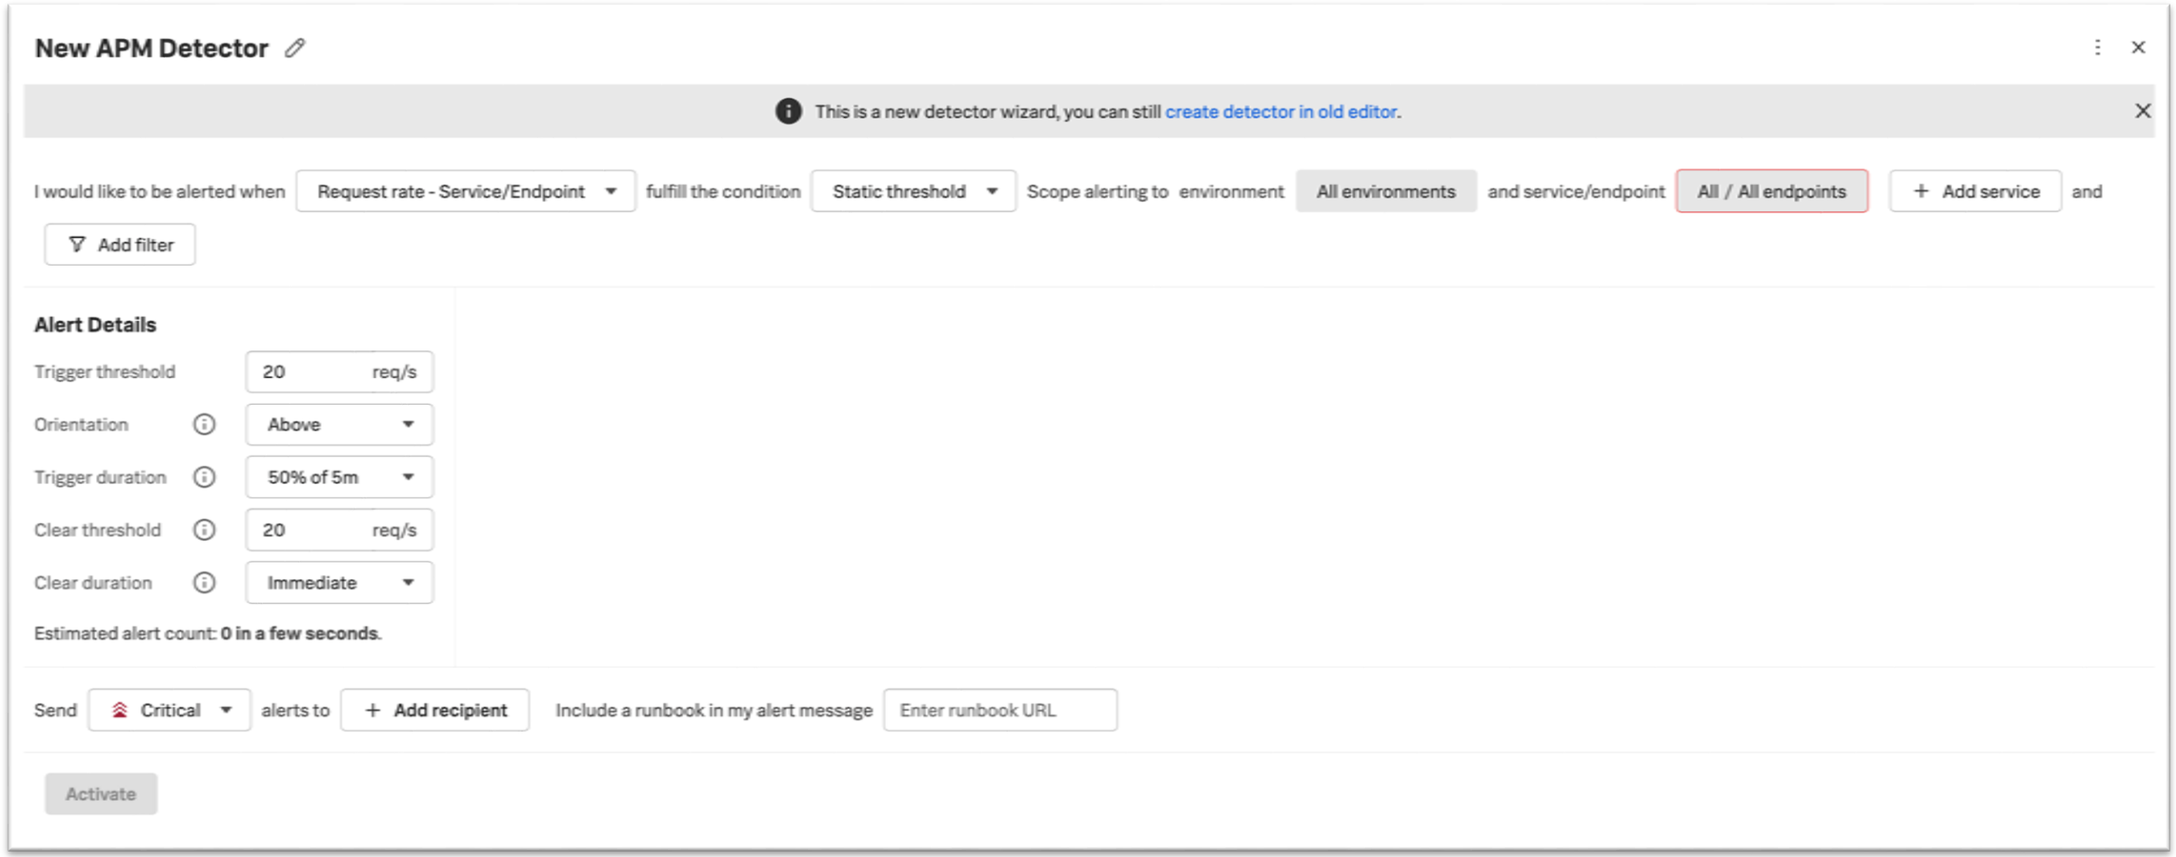

Create an APM Detector Steps

1. Name your detector
2. Select the metric you want to alert on
3. Set the condition for your alert:
    1. Static threshold
    2. Sudden change
    3. Historic anomaly
4. Select the scope of your alert
    1. You can select specific environments, workflows, services, and endpoints.

Create an APM Detector Steps (cont.)

5. Configure your alert details
6. Select the alert severity
7. Select Activate

AutoDetect

- alerts and detectors that Splunk Observability Cloud automatically creates when you have supported integrations configured
- AutoDetect detectors are available for Splunk APM and Splunk Infrastructure Monitoring
- To use AutoDetect alerts and detectors, you must first send data for integrations and instrumented services
- Here is a list of available detectors

How to Use AutoDetect

- Navigate to the **Alerts** page
- Select the Active Alerts or Detectors tab on the Alerts page
- AutoDetect components are indicated by the Auto badge

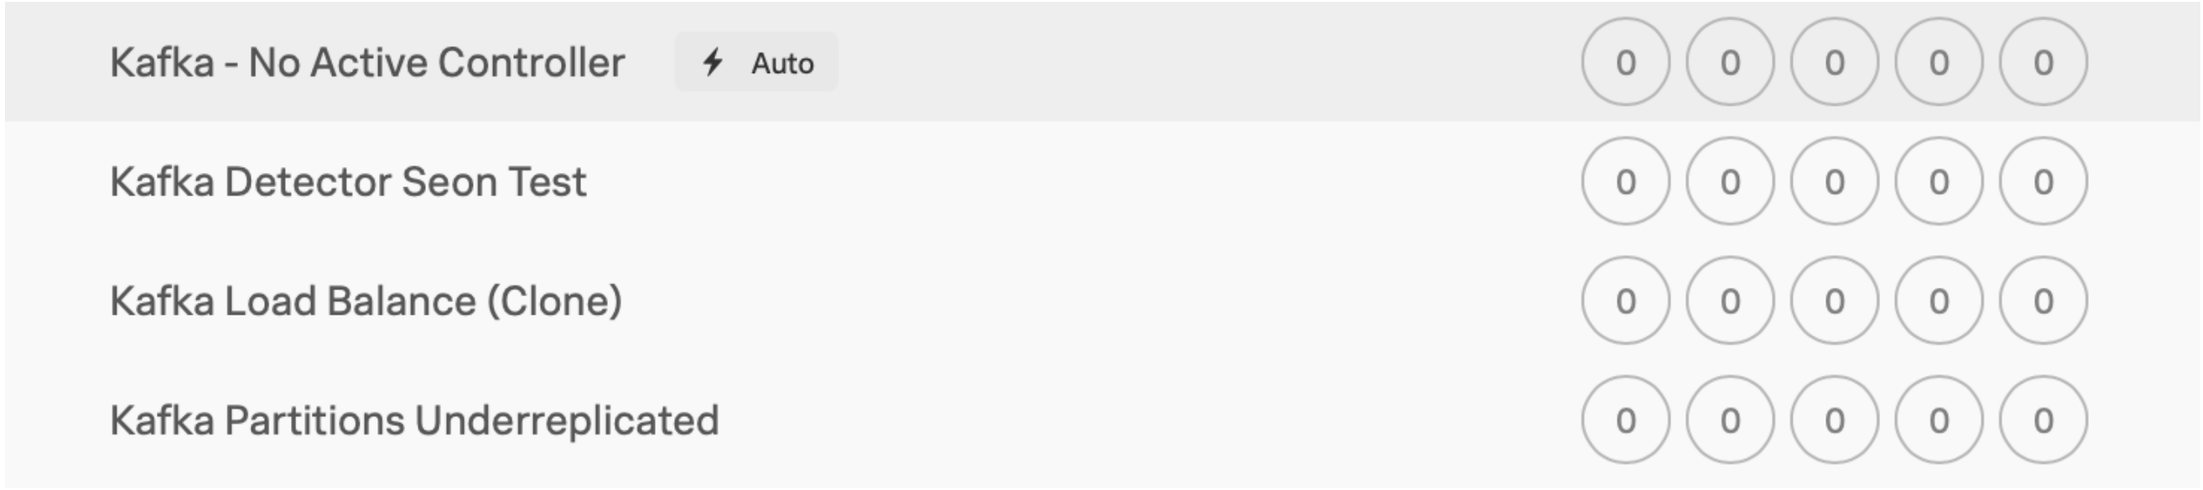

AutoDetect Detectors

- If available, AutoDetect detectors are connected to a chart by default
- On the Dashboards page you can select the bell icon on a chart to see AutoDetect detectors linked to that chart
- You can also subscribe to these alerts and detectors
- You are able to create a custom copy of detectors or disable them
  - Edits to the custom copies will not affect the original detector

Customizing AutoDetect

- Select the AutoDetect detector you want to customize
- Select Create a Customized Version
- Make edits to the detector
- Save and Activate
- Customized detectors created from AutoDetect detectors are indicated by the Custom badge

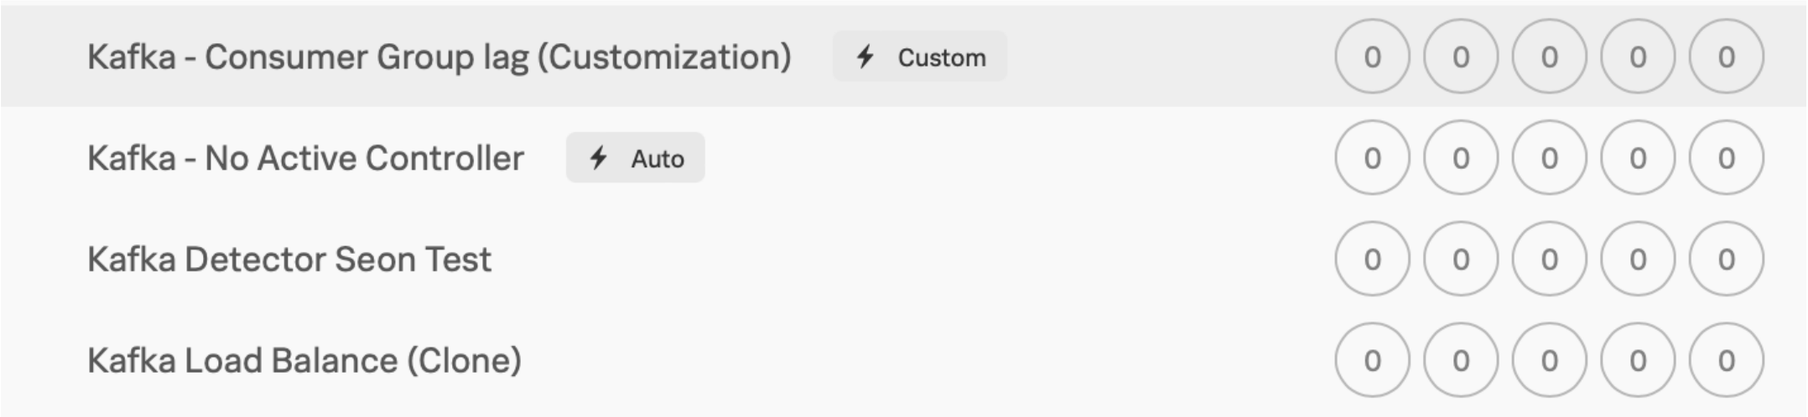

Module 5 Lab Exercise 5b

Time: 25 minutes

Description: Create a detector to monitor error rate

Tasks:

- Open Service dashboard for api service
- Create detector for error rate on api service
- Create a custom version of an AutoDetect detector

Take a break

When The Alert Triggers

View the high error rate alert from the APM > Overview page

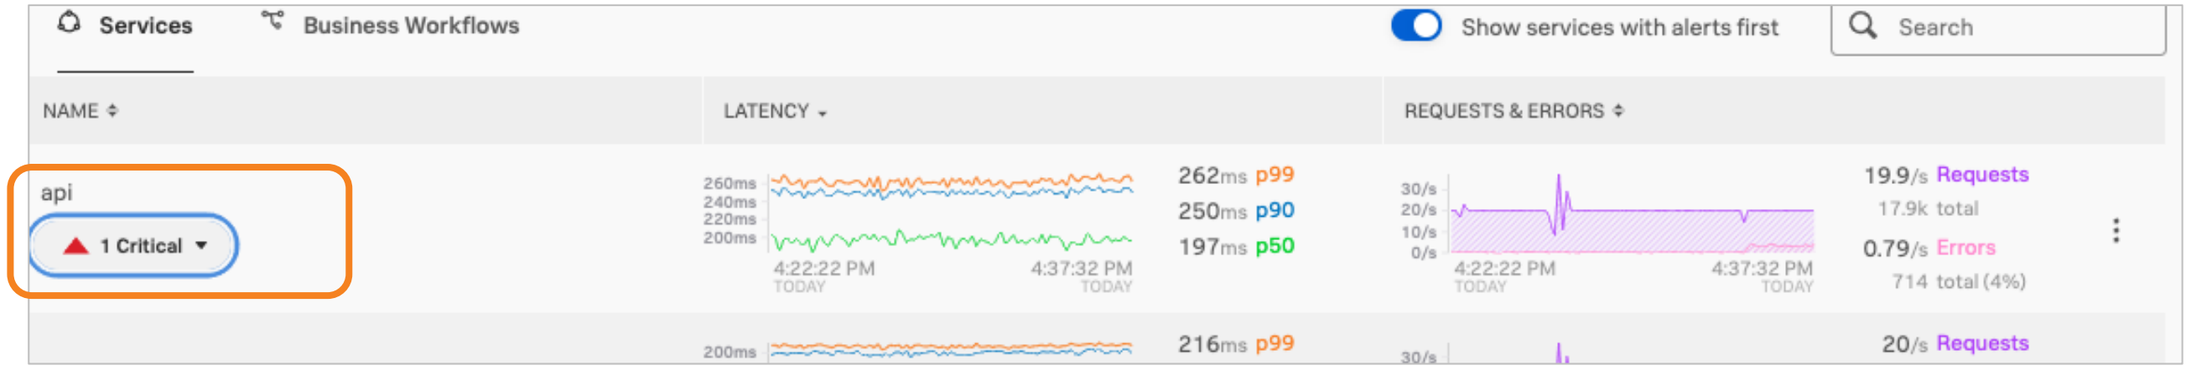

Alert Details

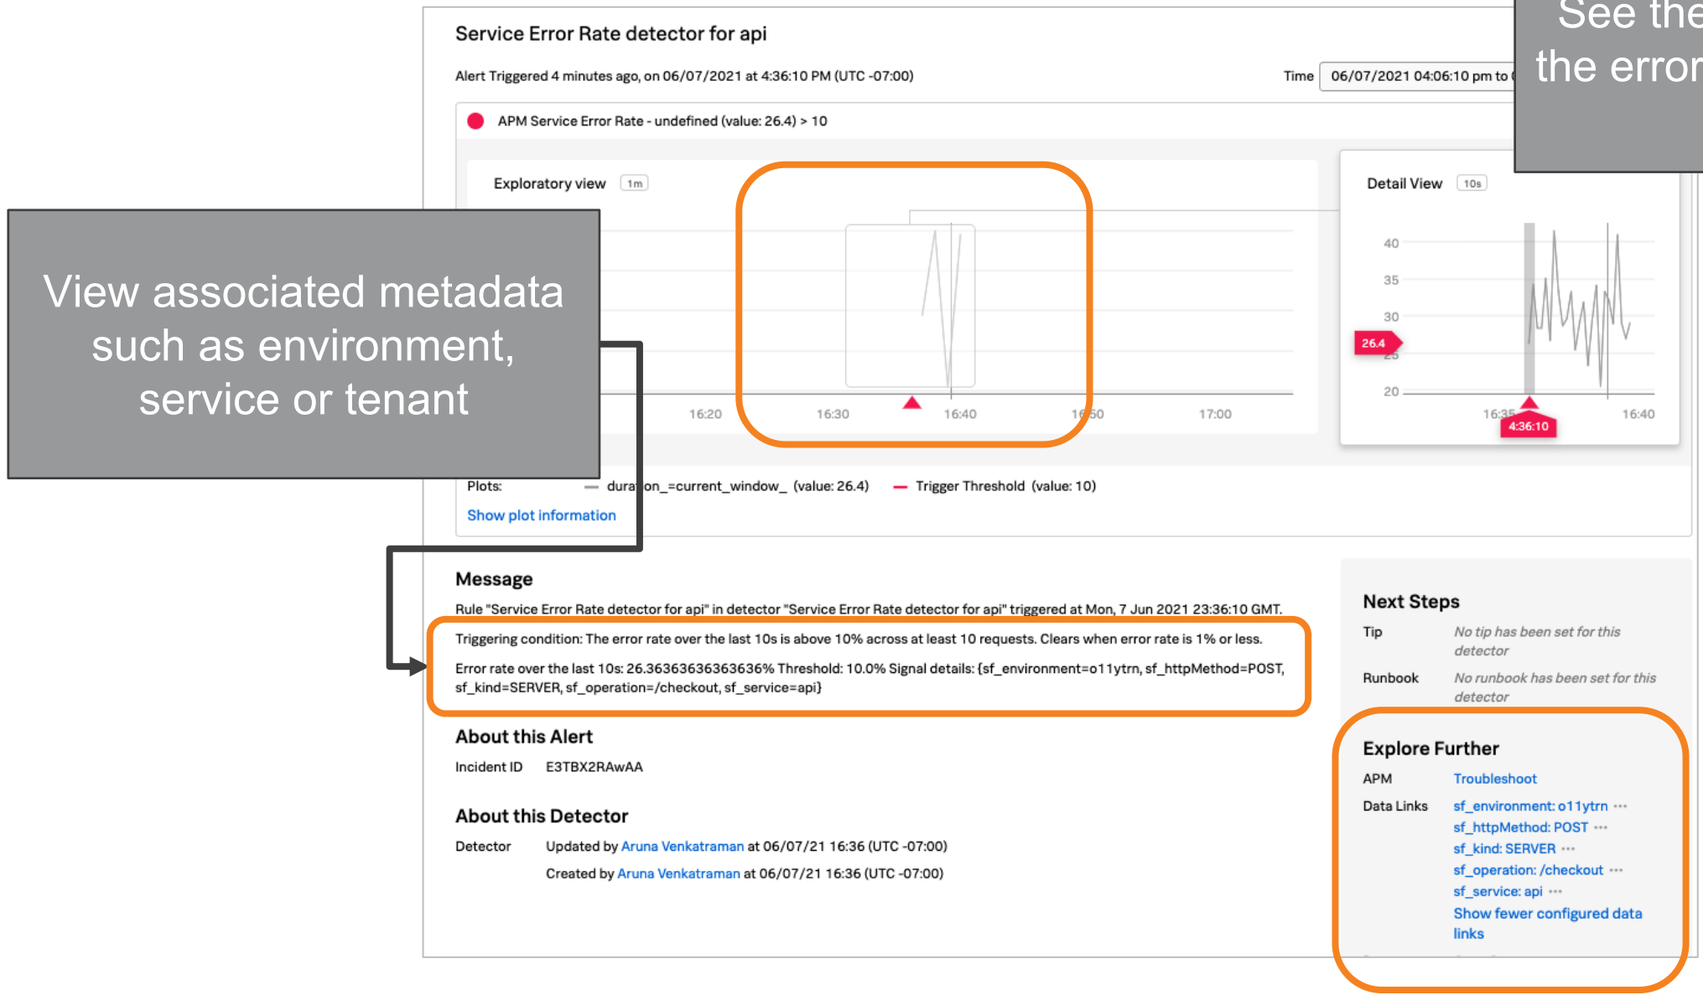

See the historical trend of the error rate for the service

View associated metadata such as environment, service or tenant

View APM Service Dashboard

- To investigate further, view Service dashboard
- Do you see any sudden spikes in the infrastructure and other resource metrics in the dashboard?

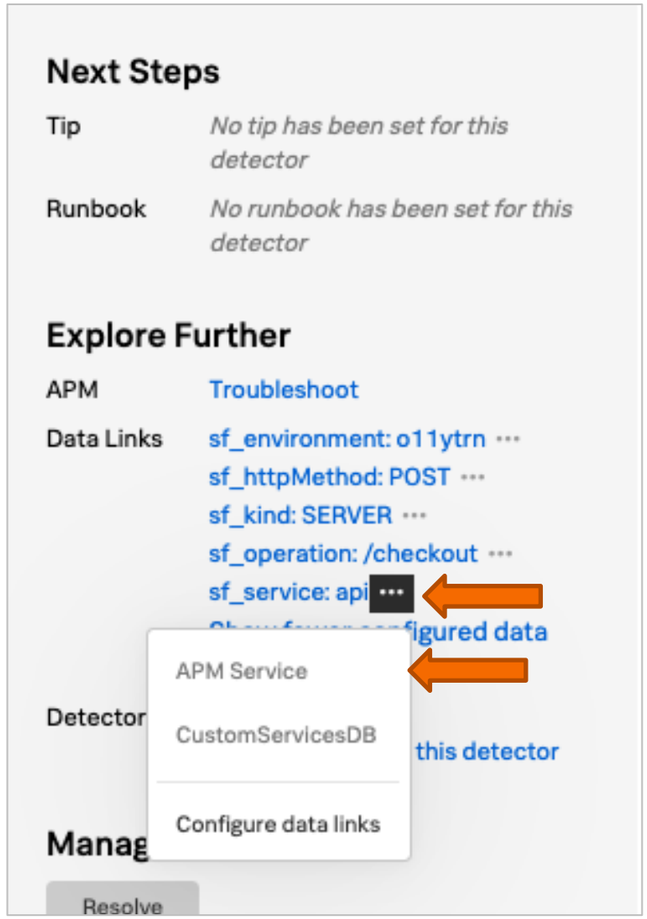

View APM Service Endpoint Dashboard

- Is there an endpoint that has a higher error rate?
- Go to the Endpoints dashboard for that endpoint

Troubleshoot Endpoint

Go to the Explore page for that endpoint

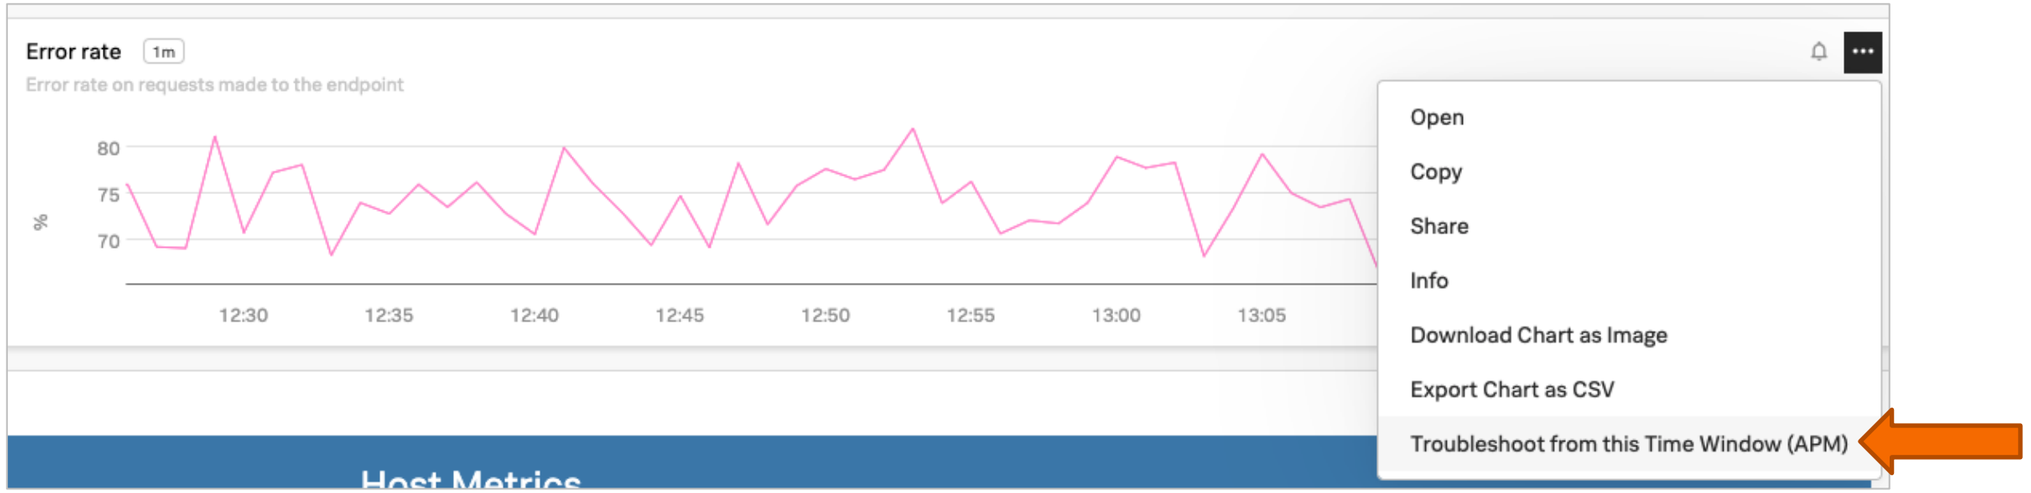

Explore Page

- Review the Tag Spotlight area on this page
- Note the top tags associated with error spans and high latency spans
- Click any chart to see example traces (Requests & Errors or Latency)
- Click a trace to go to the Waterfall view

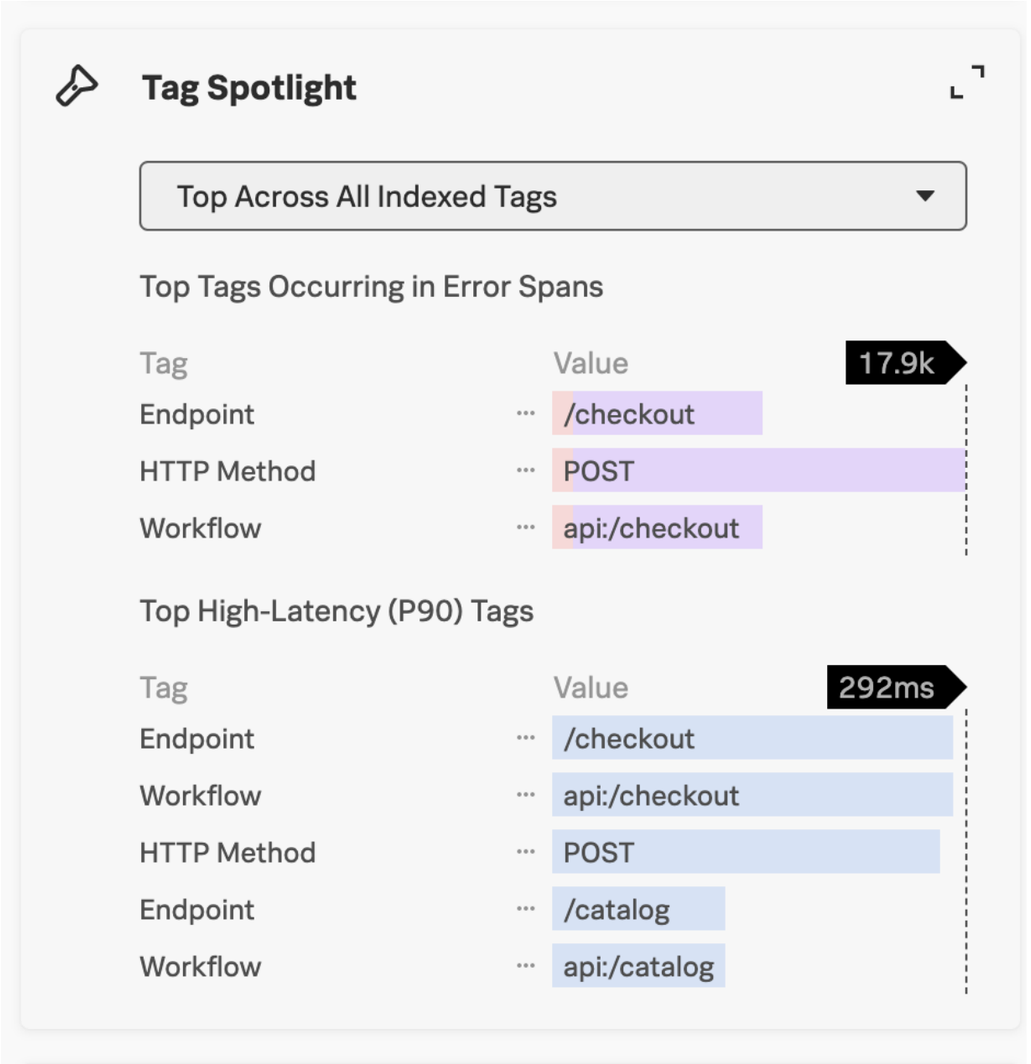

Trace and Span Details

- Waterfall view of trace - similar to a call graph
- If there are errors, you will see the red icon next to the operations that have errors
  - For example: payment: /payment/execute
- Look at Span details for this span
- What other information do you get from the details page?
- What is the error?

Tag Spotlight

At a glance, you can identify which components are affected

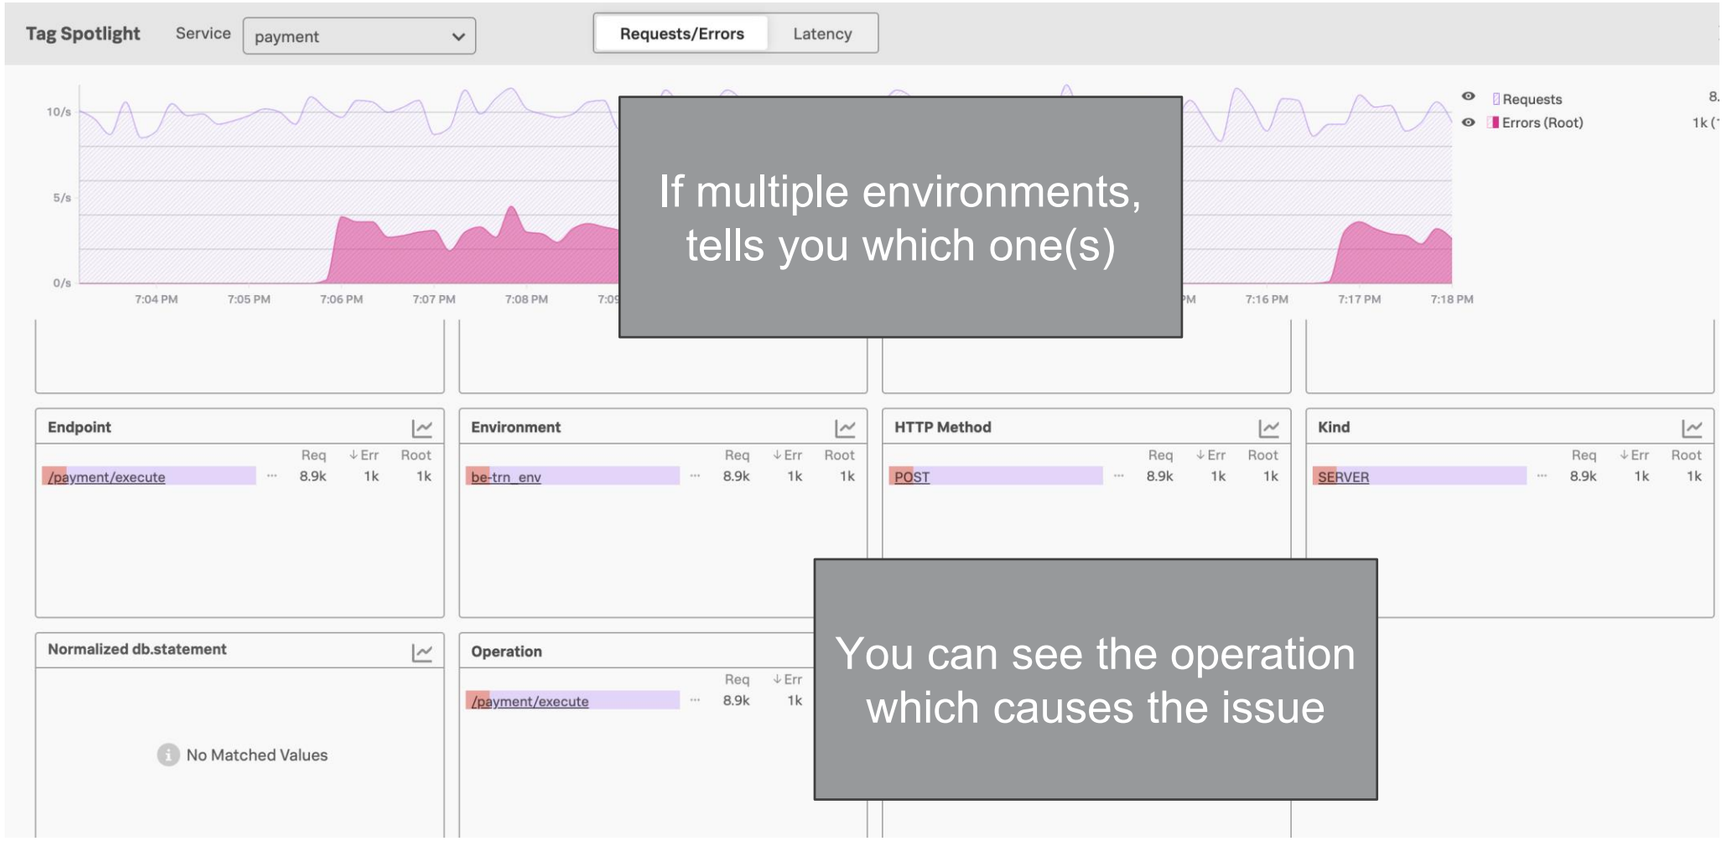

Related Content

- Allows you to jump across views of related data
  - For example, from APM to Infrastructure or Log Observer
- Keeps the same filters and context applied
- Related content relies on specific metadata keys to maintain context across different views

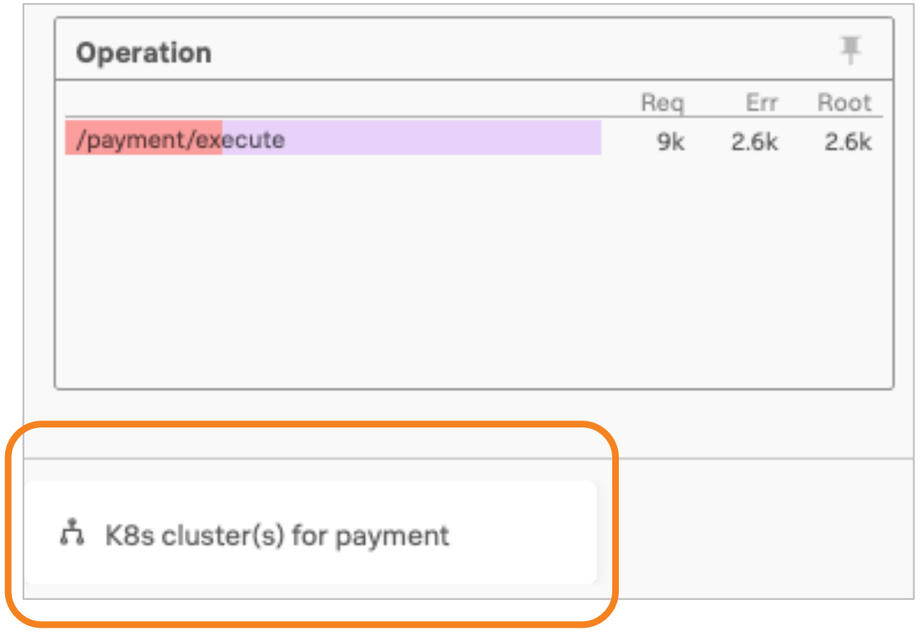

Exploring Infrastructure

- Click the Related Content link
- It takes you to the Cluster map for Kubernetes clusters
- Node Details page in the Infrastructure Navigator
- Do you notice any issues with the Payment container
- Look at the other containers running on this node
- Is there anything unusual with any of the other containers
- Get the details for a container by clicking the container name
- The problem is caused not by the payment container, but one of the other containers on the same node (noisy neighbor issue)

Module 5 Lab Exercise 5c

Time: 35 minutes

Description: Troubleshoot Errors

Tasks:

- View alert details
- View Service Dashboard
- Navigate Explore page
- View trace and span details for error trace
- Investigate using Tag Spotlight
- View related content

Quiz

Appendix

Custom Metric Sets

splunk>

Custom Monitoring Metric Sets

- You can create Monitoring MetricSets with custom dimensions
- You can use this custom dimension to create charts, dashboards, and alerts by leveraging the Infrastructure Monitoring platform
- You can filter and aggregate the generated metrics by a specific indexed span tag such as customer.id, version, or cloud.provider
- In order to create a new Monitoring MetricSet based on a span tag, you need to first index the span tag

How to Generate MetricSets

- Navigate to the APM MetricSets configuration page
  - From the left navigation panel, select: **APM > APM Configurations > APM MetricSets**
  - From the APM landing page, select APM Configuration and select APM MetricSets

How to Generate MetricSets (cont.)

- Make sure the span tag is indexed
- Once indexed, it appears in the list of MetricSets and is already generating Troubleshooting MetricSets
- Select the edit icon for that span tag to open the **Edit MetricSet** modal

Add a Monitoring Metric Set

- Add Monitoring MetricSet to your configuration using the following steps:
  - In the Service field, enter the service or services for which you want to create a Monitoring MetricSet
  - In the Add MetricSet or Edit MetricSet modal, select the checkbox for Also Create Monitoring MetricSet

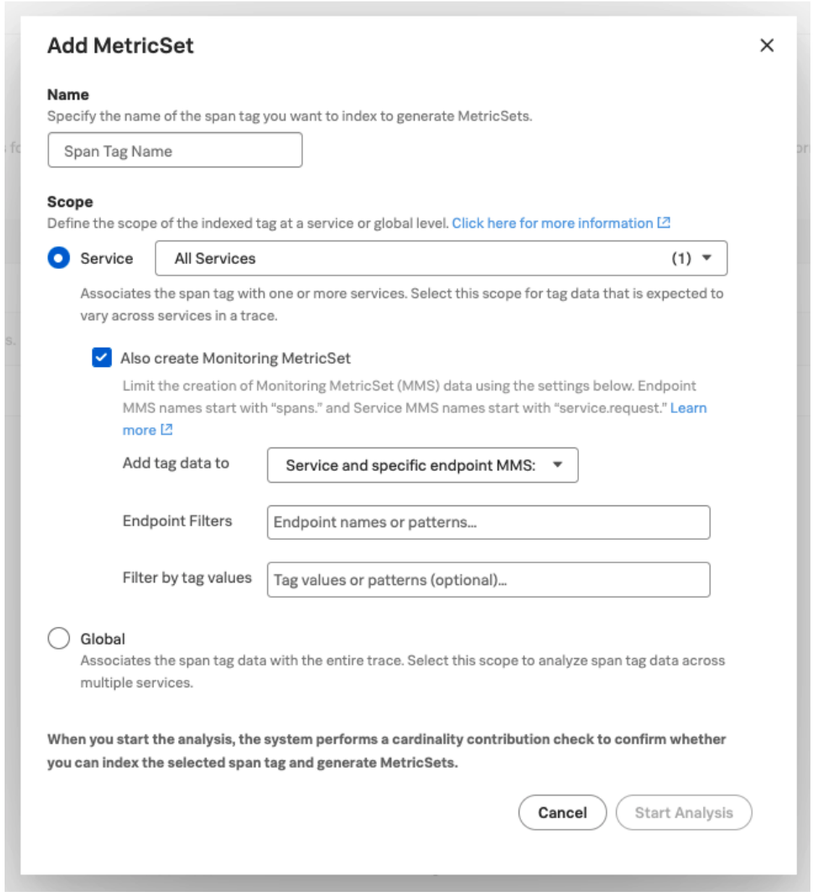

Add a Monitoring MetricSet (cont.)

- Choose how you want to add tag data to your Monitoring MetricSet from the drop-down list
  - Service and all endpoint MMS
    - Create an MMS for each of the selected service(s), as well as an MMS for each endpoint in each selected service
  - Service and specific endpoint MMS
    - Create an MMS for each of the selected service(s) and an MMS for specific endpoints you select
  - Service MMS only
    - Create an MMS for each of the selected service(s) and no endpoint-level MMS

Add a Monitoring MetricSet (cont.)

- (Optional) In the field labeled **Filter by tag values**, add tag filters to narrow the scope of your MMS to include only data associated with the tag values you enter
- Select **Start Analysis** and wait for the cardinality check to run

Use the Custom MetricSet

- You can create charts, dashboards, and alerts based on your custom Monitoring MetricSet
- To use the custom dimensionalized Monitoring MetricSets you have created, apply the filter sf_dimensionalized:true
- This filters out the metrics generated by the default Monitoring Metricset. To filter your metrics even more, use the new dimension you have created which is the tag name

Troubleshooting MetricSets

- Used for filtering service-graph and breaking down SLIs, historical comparison for spans/workflows
- Created out-of-the-box for combinations of:
  - Service
  - Endpoint
  - Workflow
  - Edges
  - Operation

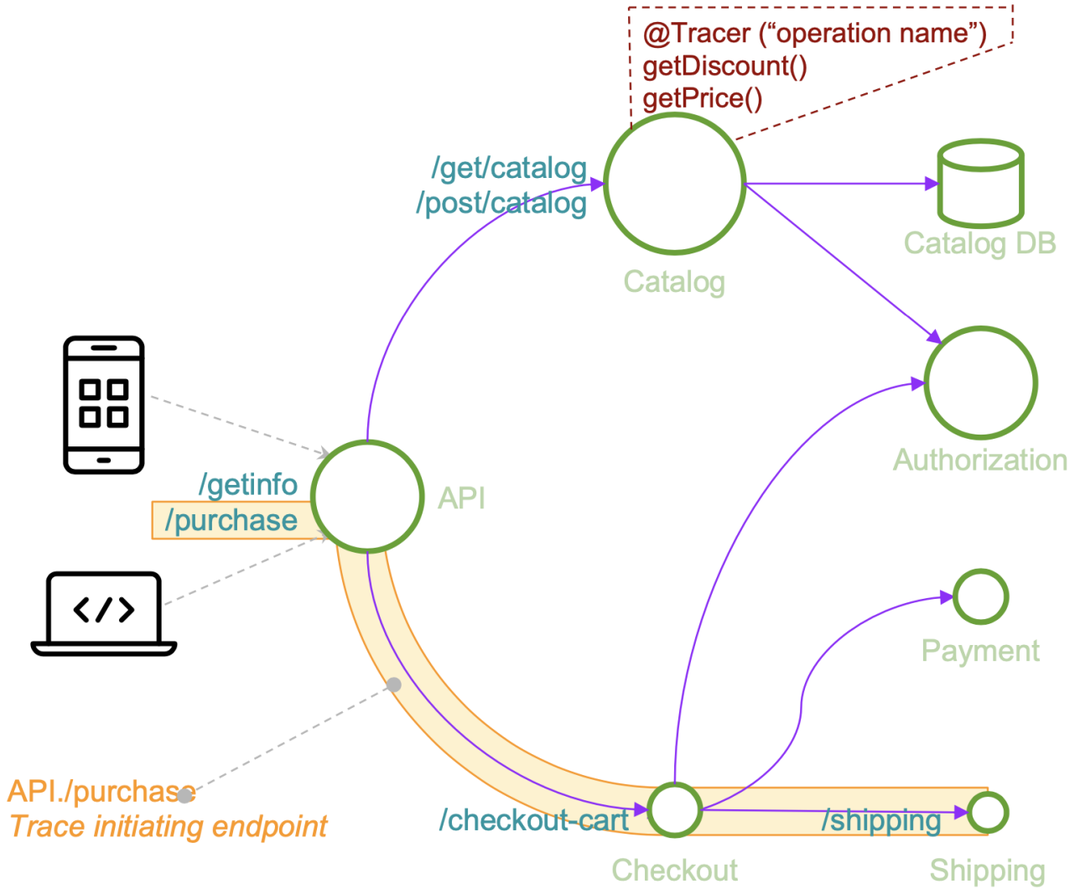

Troubleshooting MetricSets (cont.)

- Each Troubleshooting MetricSet contains the following metrics:
  - Request-Rate
  - Error-Rate
  - Root-Cause Error-Rate
  - Latency: Min, Max, P50, P90, P99
- Stored for 8 days by default along with full-fidelity traces

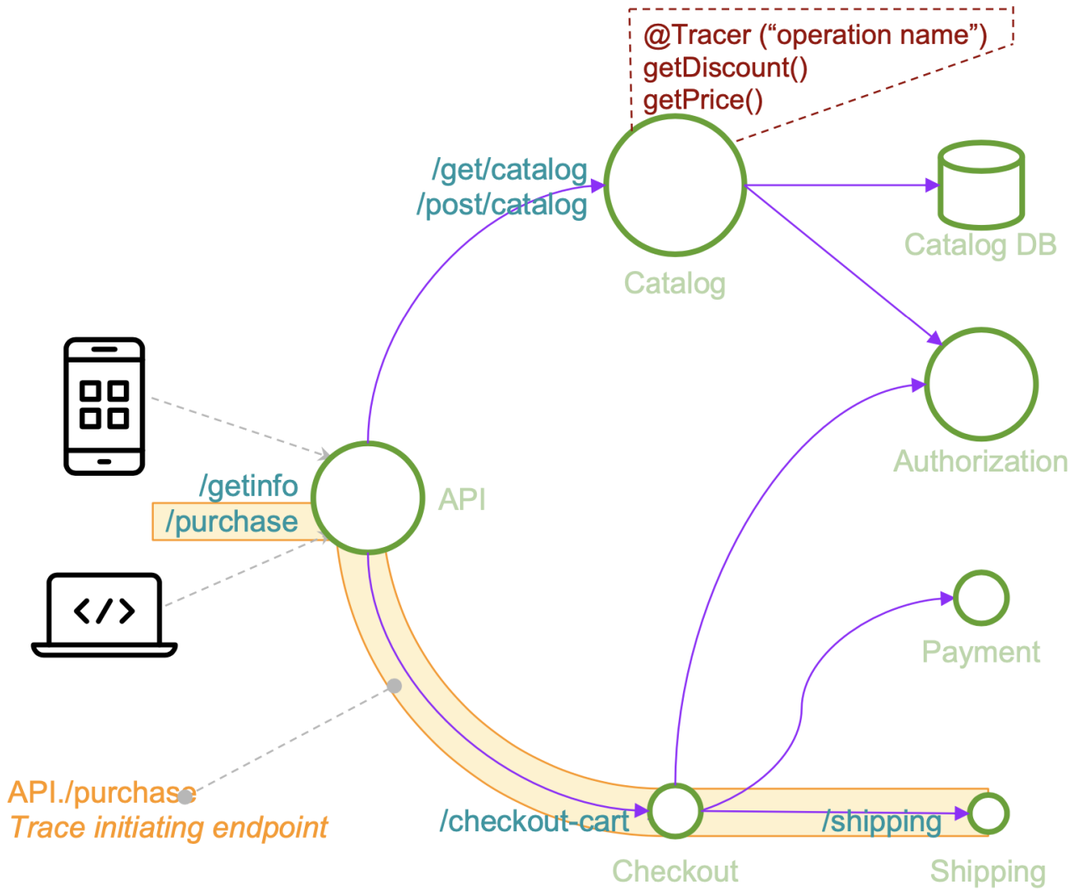

Profiling Lab

splunk>

Profiling and DB Query Scenario

- Customers are complaining about slowness on item pages with ads
- Use AlwaysOn Profiling to pinpoint stack traces that contain more information about where the problem is occurring
- Use DB query to check if there are any heavy queries that are causing the slowdown

Profiling and DB Query Steps

- Open the service map for the service with high latency (ad service)
- Navigate to AlwaysOn Profiling to view stack traces
  - Investigate CPU profiling
  - Investigate Memory Profiling
- Navigate to the service map for all service
- Select a database on the service map
- Navigate to DB Query Performance to view more information about the queries

Module 4 Lab Exercise 4

Time: 20 minutes

Description: Investigate further using profiling and DB Query Performance

Tasks:

- View CPU and memory profiling details
- View DB Query Performance details

Community

- Splunk Community Portal – community.splunk.com
  - Answers
  - Discussions
  - Splunk Trust
  - User Groups
  - Ideas
- Splunk Blogs – splunk.com/blog/
- Splunk Apps – splunkbase.com
- Splunk Dev Google Group – groups.google.com/forum/#!forum/splunkdev
- Splunk Docs on Twitter – twitter.com/splunkdocs
- Splunk Dev on Twitter – twitter.com/splunkdev
- Splunk Live! – splunklive.splunk.com
- .conf – conf.splunk.com

Splunk How-To Channel

- Check out the Splunk Education How-To channel on YouTube:

splk.it/How-To

- Free, short videos on a variety of Splunk topics

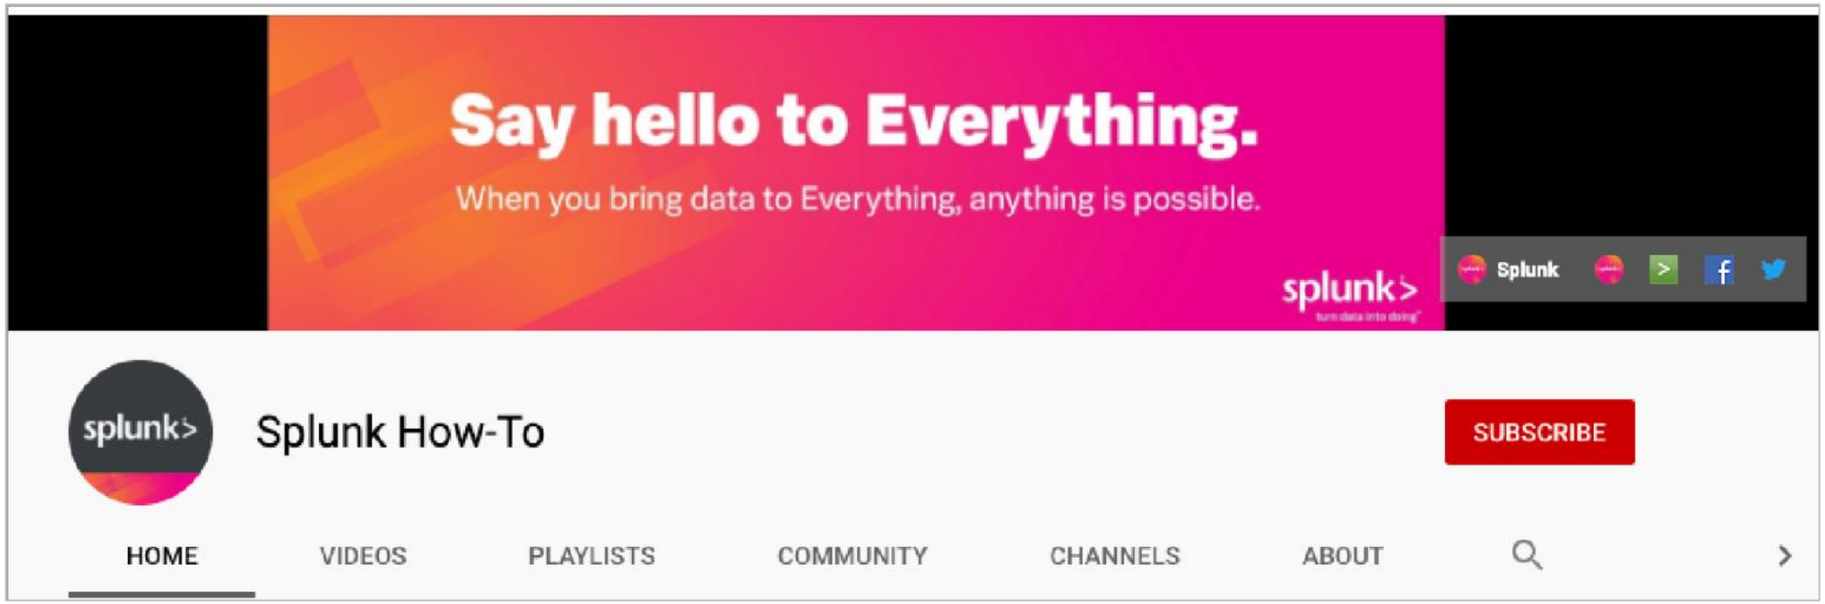

Support Programs

Web

Documentation: dev.splunk.com and docs.splunk.com

Wiki: wiki.splunk.com

Splunk Lantern

Guidance from Splunk experts

lantern.splunk.com

Global Support

Support for critical issues, a dedicated resource to manage your account – 24 x 7 x 365

Web: splunk.com/index.php/submit_issue

Enterprise, Cloud, ITSI, Security Support

Web: splunk.com/en_us/about-splunk/contact-us.html#tabs/customersupport

Phone: (855) SPLUNK-S or (855) 775-8657

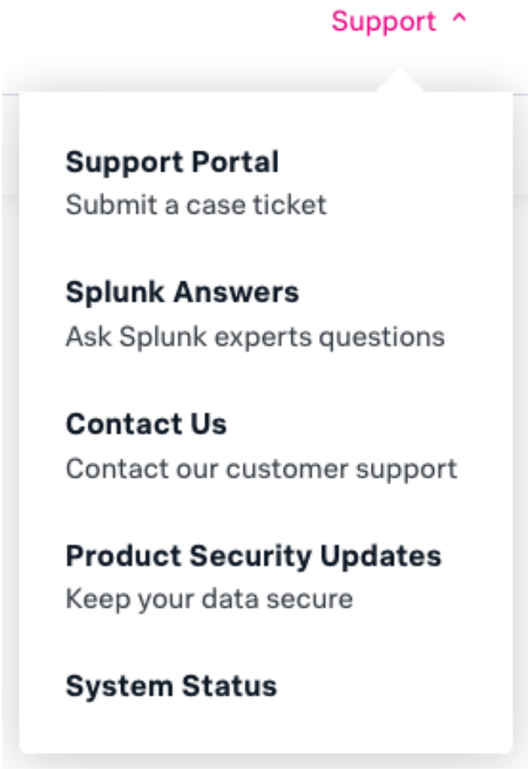

OnDemand Services for Expert Assistance

Direct Access

- Credit-based service accessible through the Support Portal.

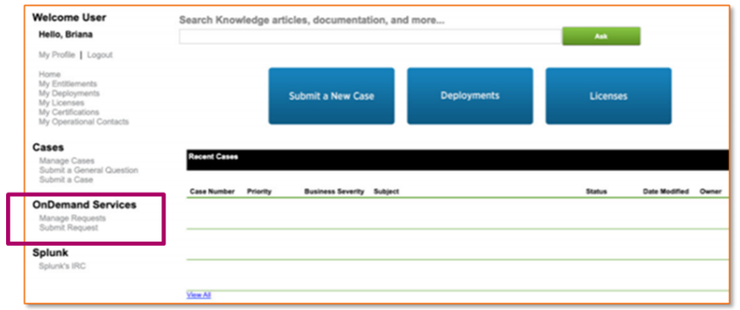

Get Started

- Choose your product and desired task and get access to Splunk Experts!

Continued Help

- 20+ tasks available for continued growth and help:
- General consultations
- Adoption and onboarding
- …and more!

OnDemand Requests

- How to Open a Case
  - Most customers have OnDemand Services included as a part of their license purchase. Follow the instructions in the OnDemand Services Portal End User Guide, pick the product you need help with, open a request under **Pick Your Product** > **Splunk Core - Enterprise/Splunk Cloud** and task **"Build a Simple Dashboard"**
- Issue Opening a Case?
  - Contact the ODS team at OnDemand@splunk.com OR contact your Customer Success Manager/Advocate or Account Team

Splunk Mobile

- Free app available to all Splunk Cloud and Splunk Enterprise customers
- Analyze data and receive actionable alerts on-the-go with mobile-friendly dashboards
- iOS and Android
- See the Product Brief
- Download for iOS **splk.it/ios** and Android **splk.it/android**

Thank You