Distributed Tracing in Splunk

Get end-to-end visibility into application performance with Splunk and OpenTracing

splunk> .conf19

Distributed Tracing in Splunk

Gary Burgett
Staff Sales Engineer | Splunk

David Cornette
Enterprise Monitoring Architect | T-Mobile

splunk> .conf19

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.

T-Mobile and OpenTracing

splunk> .conf19

T-Mobile: The Uncarrier's Growth

Q2 2019, the 25th consecutive quarter where T-Mobile has added more than 1 million customers

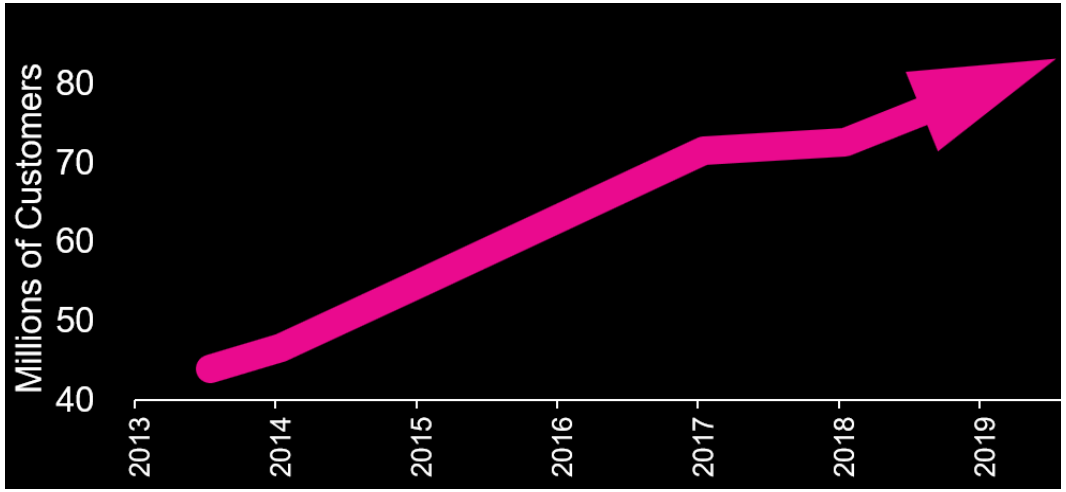

Leading the industry in growth, year after year…

Growth has raised the stakes on IT capacity, performance and availability

Direction from leadership to get bigger, better, faster, and for less money

Transition to Dev/Ops

- Shift to CI/CD and hybrid cloud
- Split core functionality into domain-owned microservices
- Embrace containerized platforms in the cloud and on-prem
- Shifting to an increasingly complex ecosystem

splunk> .conf19

Platform as a Service at Scale in T-Mobile

3K+ applications on 50K+ containers

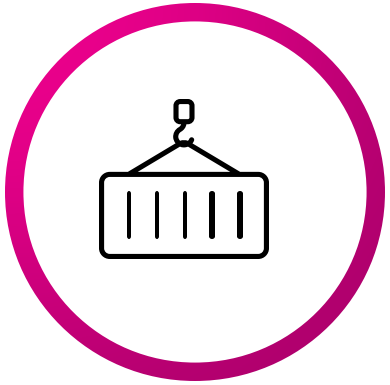

21 PCF foundations and 24 PKS clusters

700M+ transactions a day, doubled from a year ago

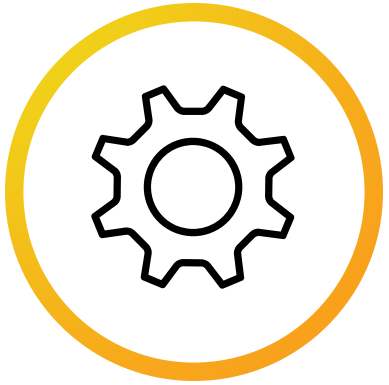

Operating in AWS and Azure, but with ~90% in on-prem PaaS offering

Enterprise APM at Scale in T-Mobile

9K Microservices

250 Monitored Applications

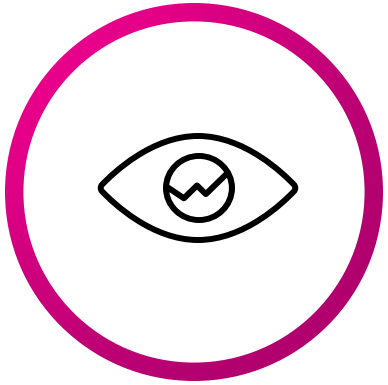

9K Microservices

35K Monitored Instances

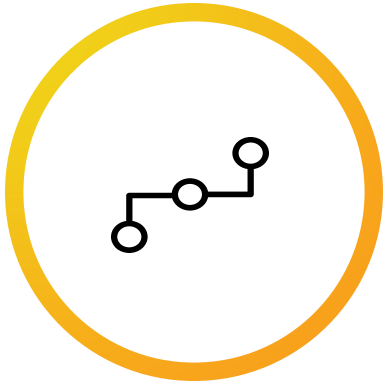

Distributed Tracing In a Complex Environment

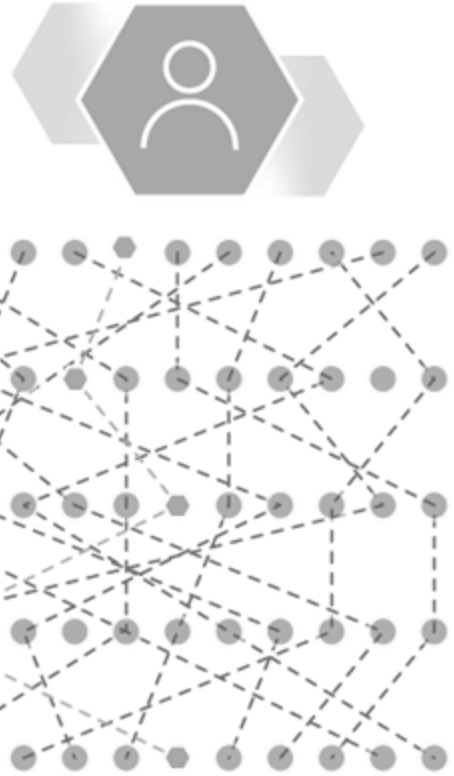

Traditional APM's agent based instrumentation comes with challenges in cloud services or with containerized platforms

Opensource specs are growing in popularity, though still fairly new

T-Mobile is exploring feasibility of OpenTracing across a complex stack, and exploring tooling options

Distributed Tracing with Splunk

splunk> .conf19

Terms and Concepts

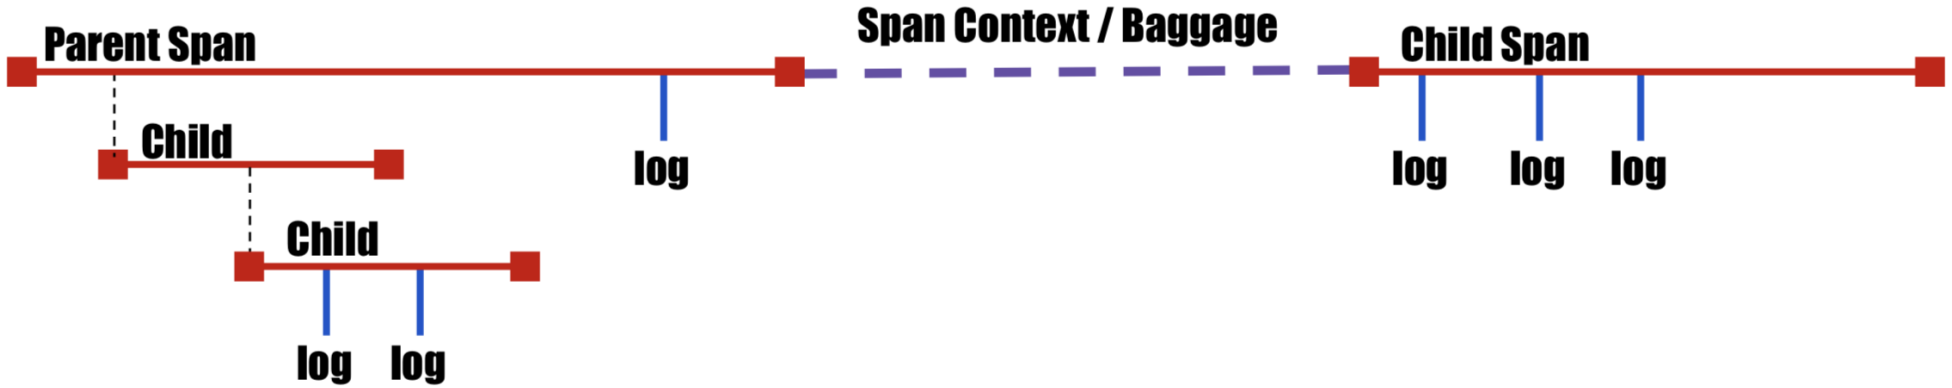

- **Trace:** All data pertaining to a distributed request; a collection of spans.
- **Span:** A logical unit of work within a request. Spans have start/end times, and may define relationships to other spans (e.g. parent/child).
- **Context:** Metadata that is propagated across spans, includes trace and span ids, as well as key-value pairs (AKA baggage).
- **Tags:** Additional key-value pairs that can be added to specific spans for more granular filtering or analysis (eg. error codes, container/host details). Not propagated to downstream systems.
- **Logs:** Good ole logs, but supercharged with the span context. Enables correlation of logs to specific transactions/requests.

How do I do that?

- Code Instrumentation
- Context Management
- Context Propagation
- Transport

OpenTracing

- **Standardized span management:** programmatic APIs to start, finish, and decorate timed operations (“spans”)
- **Standardized inter-process propagation:** programmatic APIs to aid in the transfer of tracing context across process boundaries
- **Extensible Transport Architecture:** extensible API for transporting data to various storage/analytics backends (like Splunk!)

Data Collection

OpenTracing Compatible Tracing Implementations!

- Python, Golang, Javascript, Java, CSharp, PHP, Ruby
- Mobile Coming Soon? (Android, Obj-C)

They all Speak HEC!

- Send data directly to your Splunk deployment on-prem or in the cloud
- Send to a local forwarder to consolidate traffic and extra enrichment

Beautifully Formatted JSON Lands in Splunk

Data Collection

But what if I'm already tracing with another implementation?

Zipkin, Jaeger, Lightstep, etc. all have different transport mechanisms and naming conventions

No problem, with Splunk you always have options

- **Data Collection:** Depending on your implementation, different output types are supported (File, HTTP). Just needs to get to Splunk in human readable format
- **Data Formatting:** there's a CIM for Tracing
  - Use Splunk field extractions, aliases, calculated fields, etc. to match the fields in the CIM, and all of the Tracing Viz and content that's being built will work OOTB.

Analytics and Visualization

Splunk App for Distributed Tracing

- **Visualization:** New visualizations for plotting service topology, trace waterfalls and latency plots with percentile bands
- **Workflow:** Quickly drill from a high level overview of your service health and performance to quickly identify errors, outliers and anomalies.
- **Machine Learning:** Leverage the machine learning capabilities of Splunk to detect anomalous changes in request volume, error rate, and latency for smarter alerting.
- **Extensible Analytics:** Dashboards and workflow are all built using core Splunk tools, so you are free to clone, copy, modify and extend the analytics to meet your organizations needs

Service Topology

Service Mapping

- Leverage tracing data to discover app flow and dependencies
- Overlay health, latency, component metrics
- Drilldown to quickly troubleshoot slow or unhealthy services
- Augmented with high level KPIs and detected anomalies

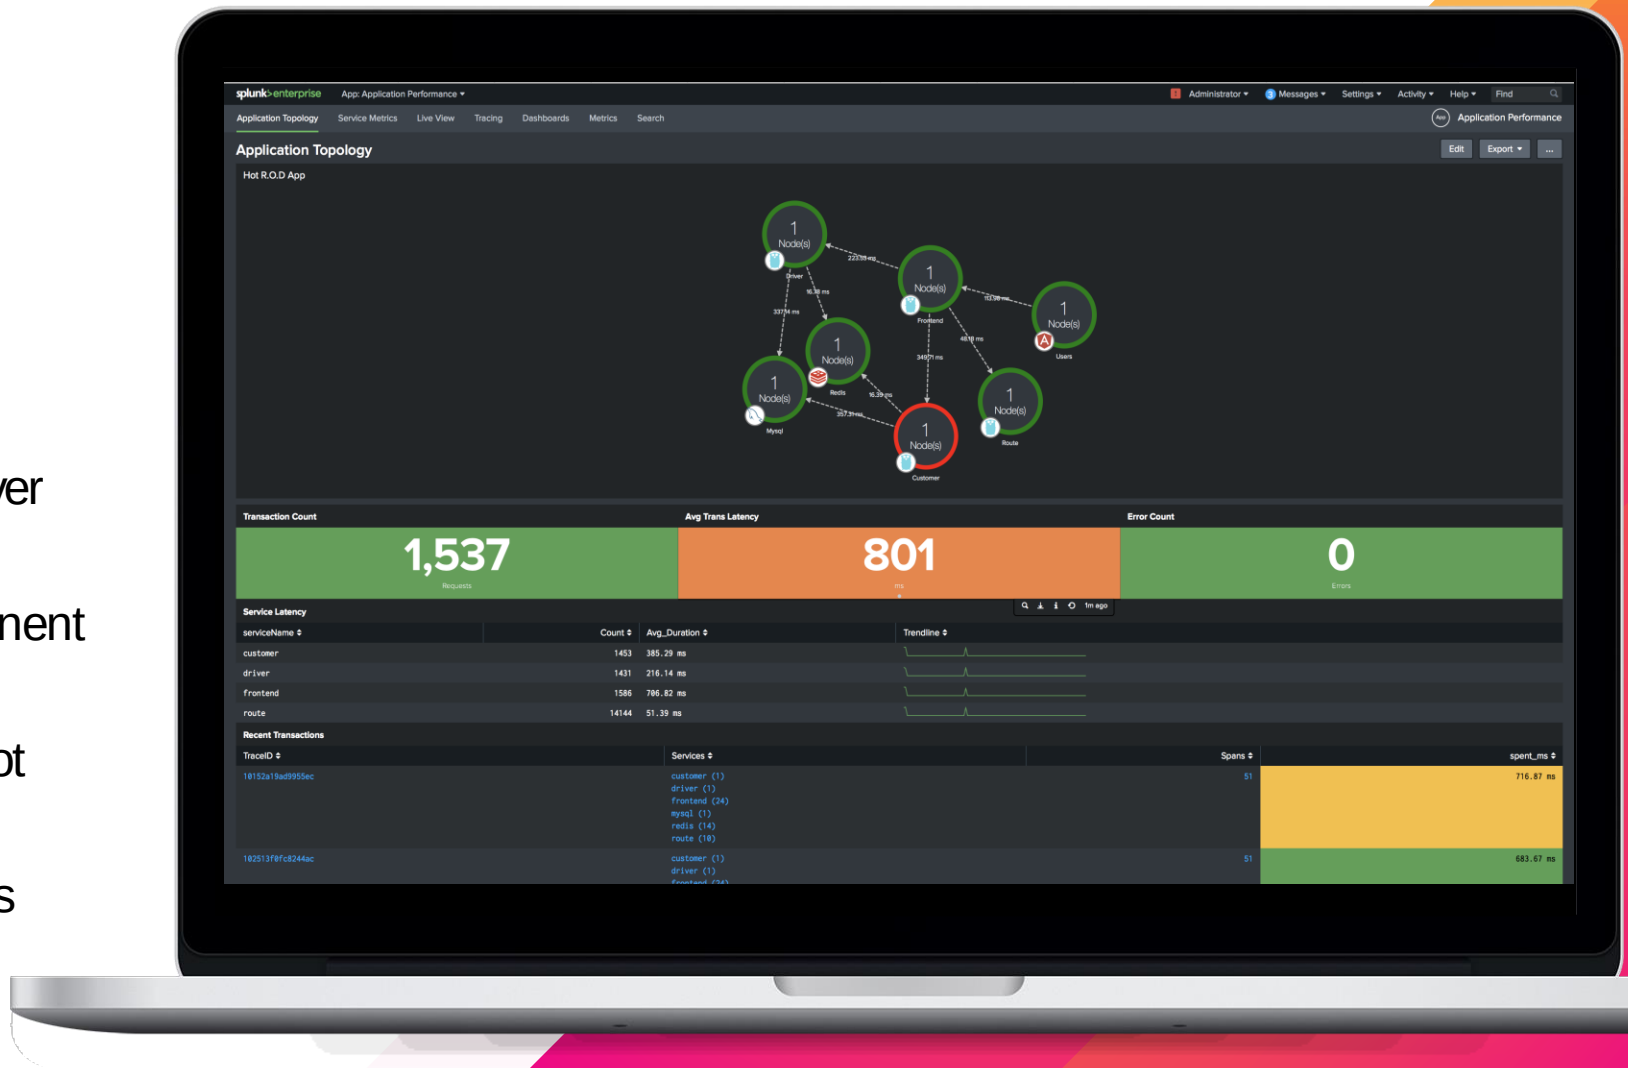

Latency Distribution

Frequency Analysis

- At a glance, understand the average performance of a service
- Quickly identify and filter on outlier bands
- Find common patterns in slow trace bands
- Drill down to look at details for specific traces

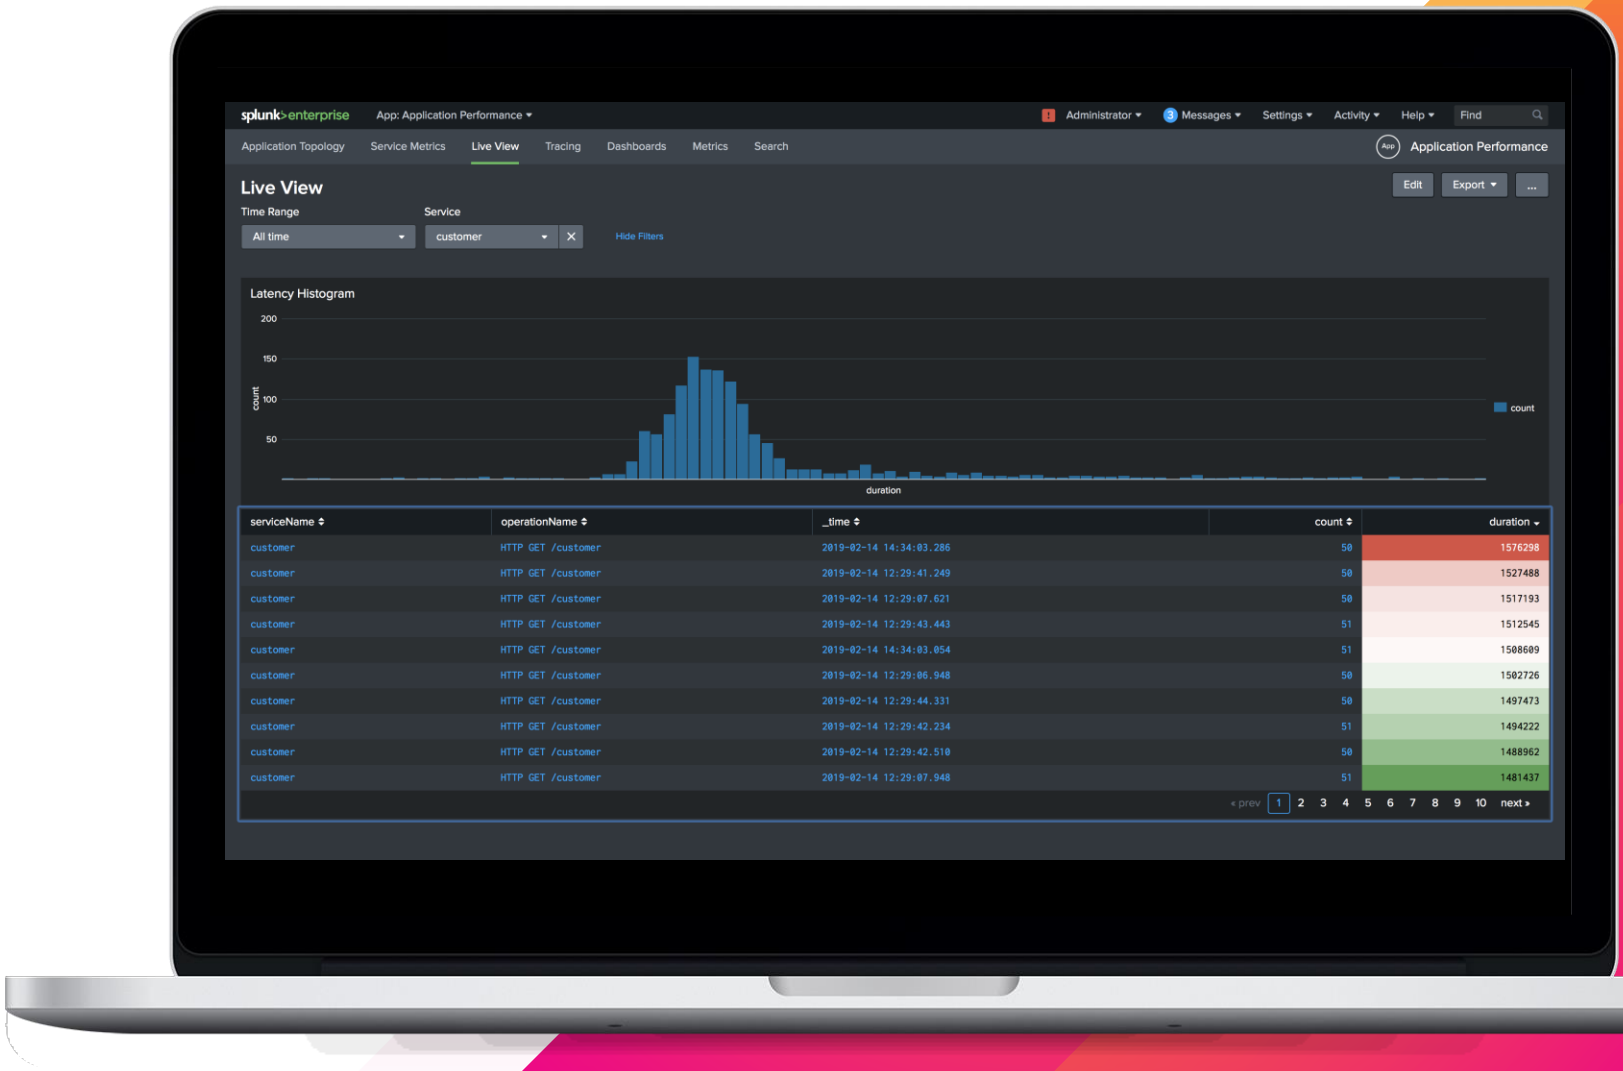

Service Level Metrics

Time Series Metrics

- View latency, request thruput, and error rate over time
- Latency includes points for each individual trace, overlaid with p50, p90, p95, p99 bands
- Quickly identify anomalous traces
- Drilldown to investigate individual traces

Trace View

Waterfall Visualization

- Quickly identify spans or service calls that are contributing to latency
- Overlaid with critical path bands to show working time vs waiting time
- Drilldown to correlate with application associated with trace
- Extensible to include any tags, meta information included with spans

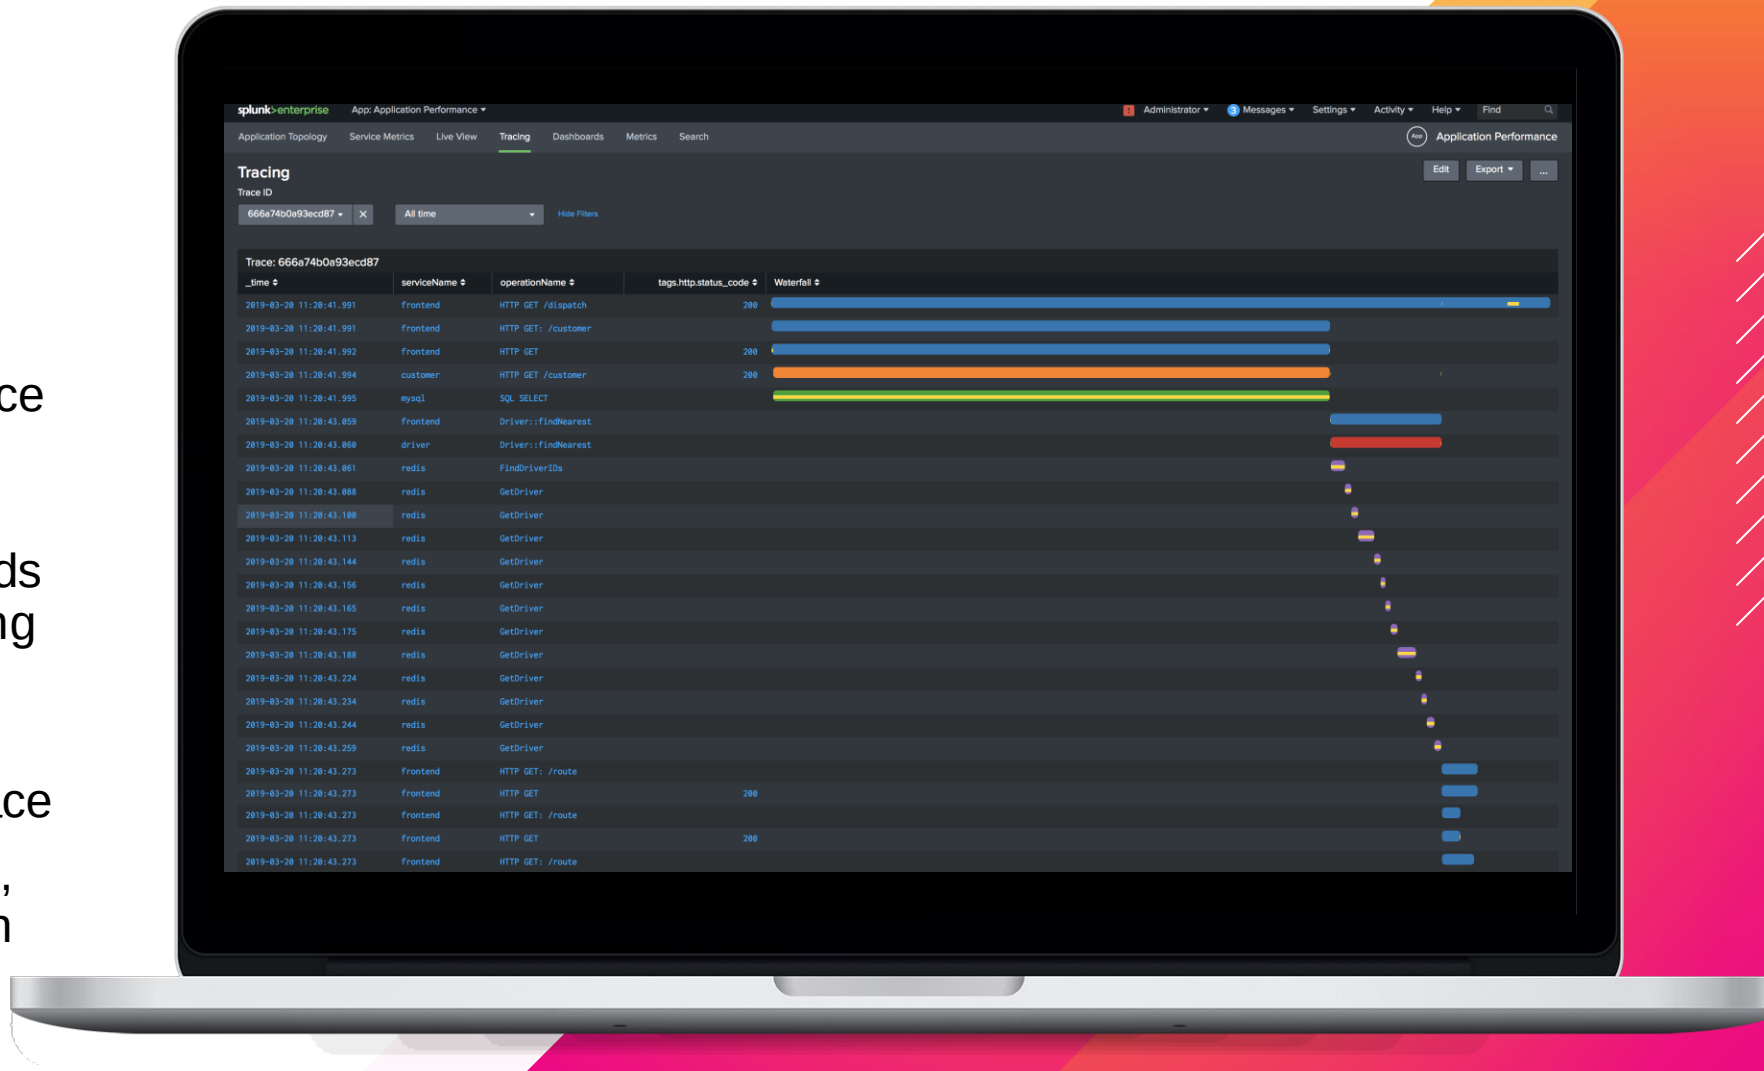

What's Next?

OpenTracing + OpenCensus = OpenTelemetry

- Instrumentation for app traces, metrics and logs
- Same tracing APIs w/additional features for metrics
- Tools for standardizing instrumentation, simplifying collection of data
- Collector supports advanced features for enrichment, sampling, aggregation
- Simplified collection architecture

Improvements to App for Distributed Tracing

- Better correlation between traces, metrics and logs
- ITSI Integration for service mapping, KPIs
- Leverage infrastructure metrics (K8s, AWS, etc.) for more complete visibility

DEMO: Tracing MAST as a Proof of Concept

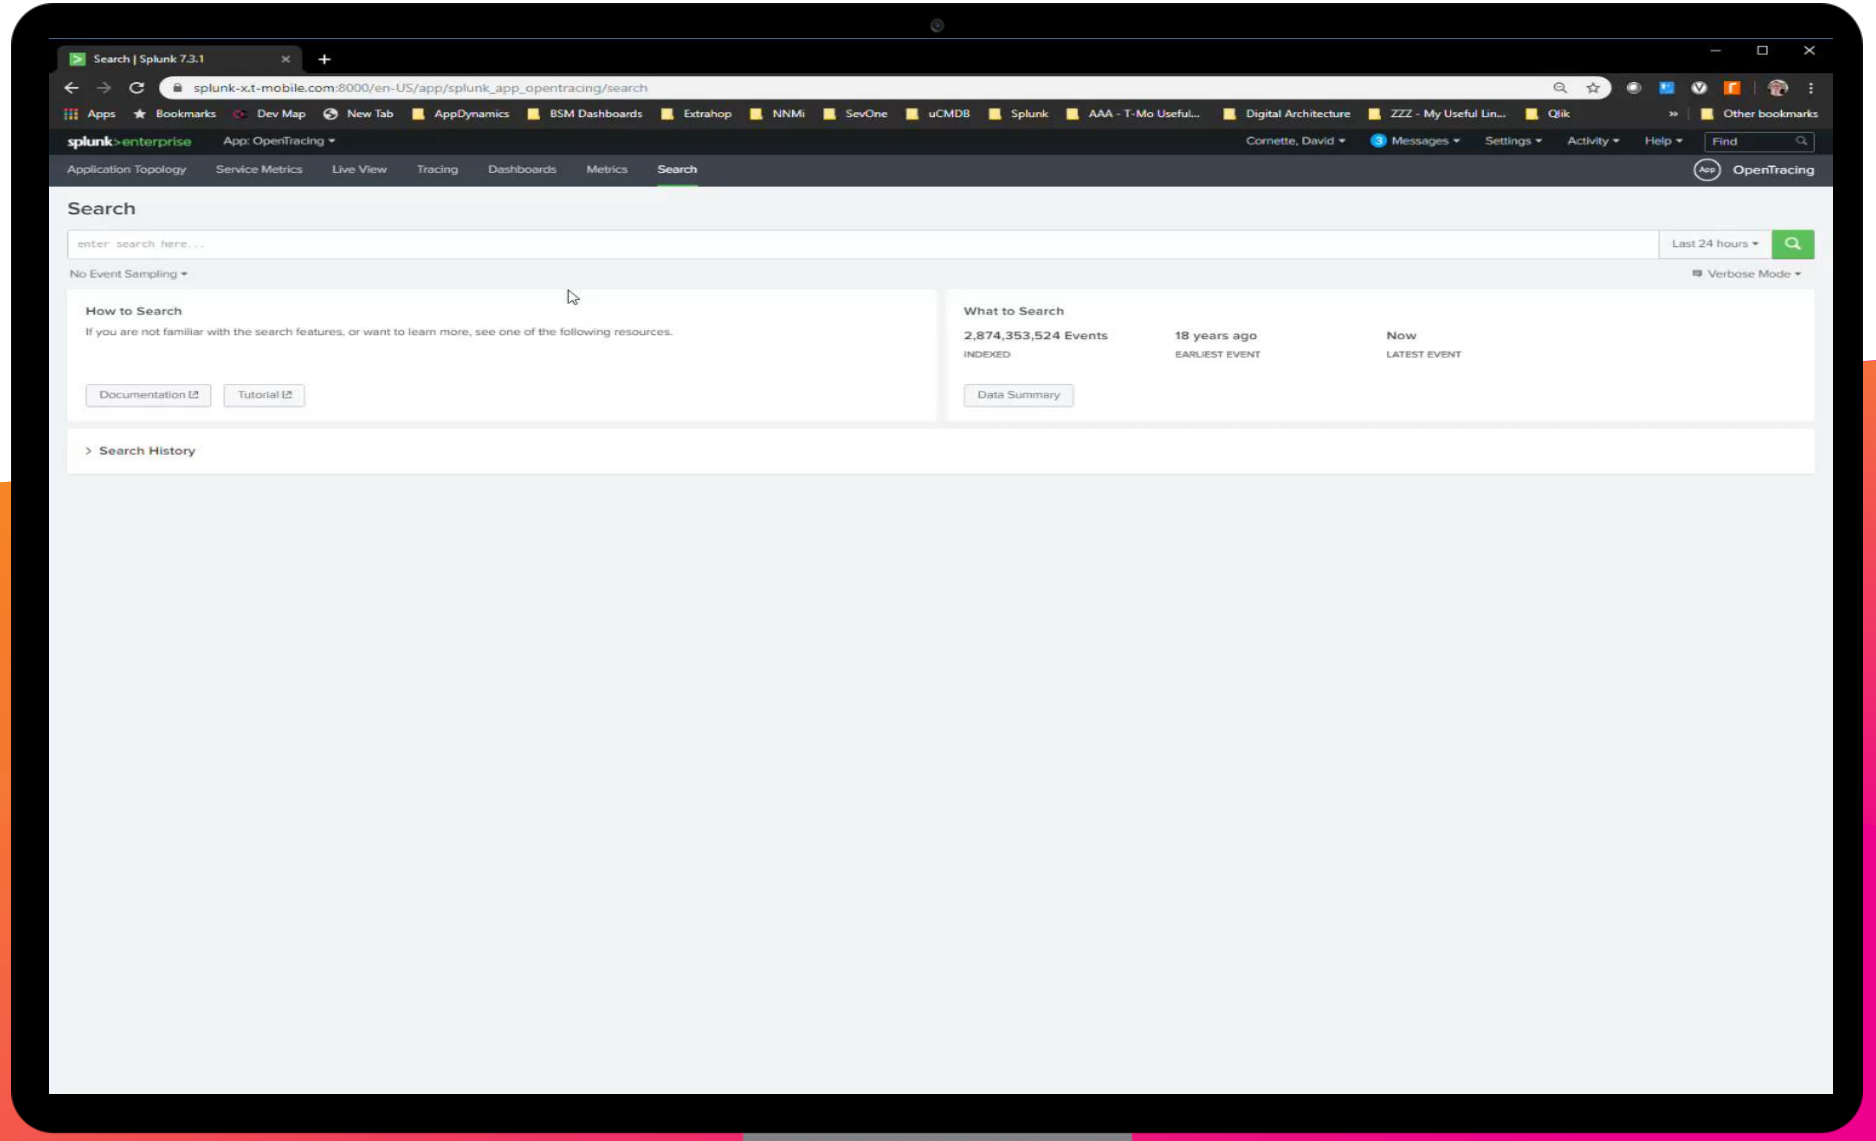

Splunk and OpenTracing at T-Mobile

Key Findings

Ability to keep and analyze every trace

- Splunk has a proven track record for handling the scale to incorporate all of trace data

Flexibility in visualizing trace data

- Adjusts to any tracing library
- Customizable to whatever tags, dashboards, drills are most meaningful for you

Potential for realizing the 3 Pillars of Observability in one place

- Ability to contextualize traces with logs and metrics
- Events to metrics for real time RED metrics from trace data
- Output of trace/span ids to logs allow ability to create a consolidated view

Q&A

Gary Burgett | Splunk Staff Sales Engineer

David Cornette | T-Mobile Enterprise Monitoring Architect

splunk> .conf19

.conf19

splunk>

Thank You!

Go to the .conf19 mobile app to

RATE THIS SESSION