



Connect

Distributed Tracing with Tempo and OpenTelemetry Auto-Instrumentation



Andreas Gerstmayr

Senior Software Engineer
Red Hat

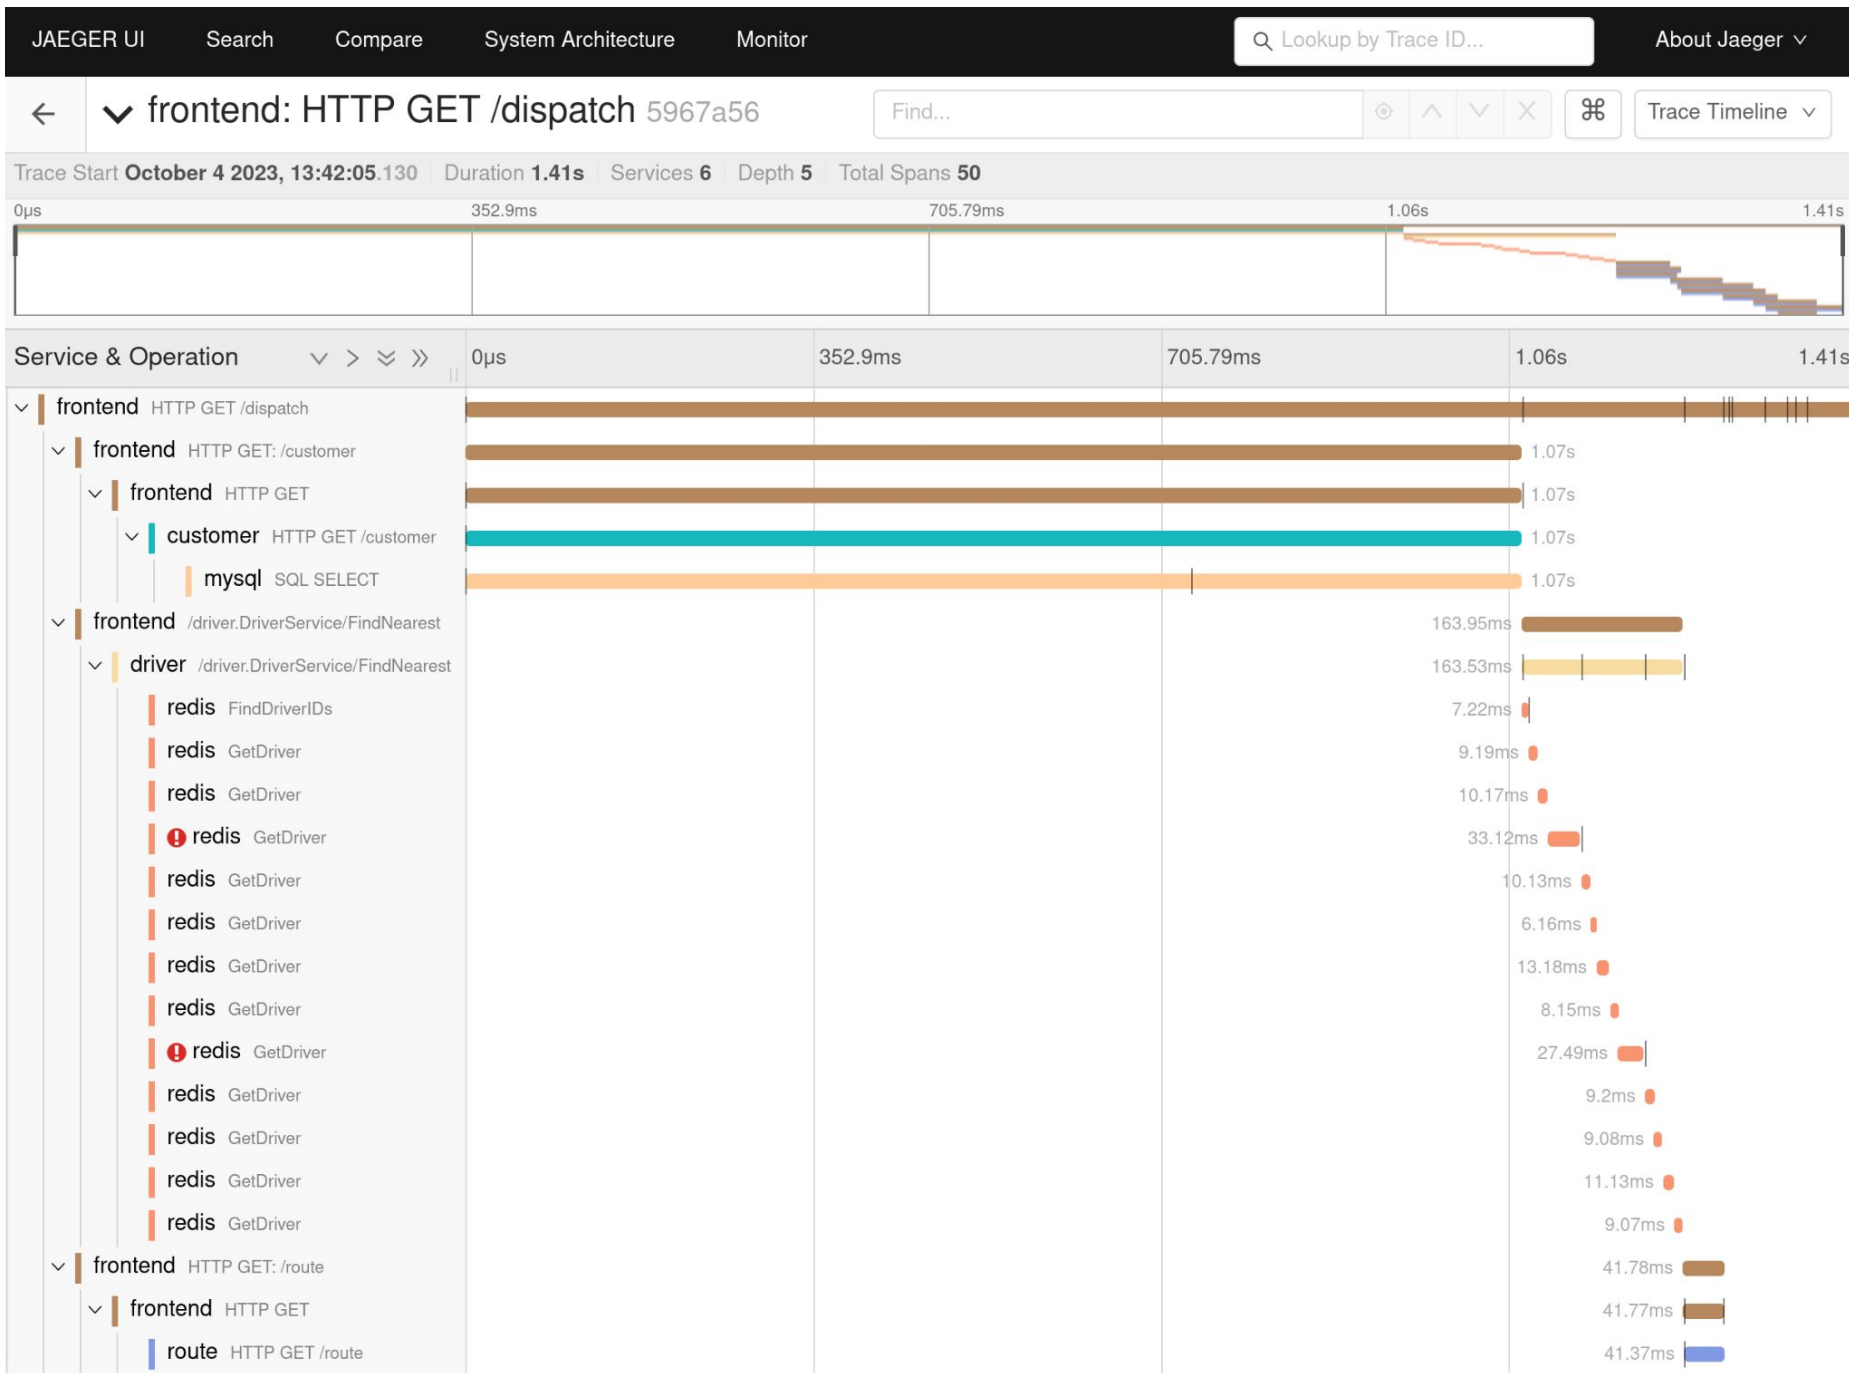
Agenda

- Introduction to Distributed Tracing
- Manual and Automatic Instrumentation
- New Distributed Tracing Stack on OpenShift
- Example

Distributed Tracing

What is distributed tracing?

- Distributed Tracing records the execution of individual requests in distributed systems (through proxies, microservices, databases, etc.)
- a trace is a data/execution path through the system and contains one or more spans
- a span represents a single unit of work, with an operation name, start, duration, and optionally custom attributes and logs



Why distributed tracing?

- reduce mean time to detect (MTTD) and mean time to remediate (MTTR)
- optimize performance
- understand how data flows through a system

OpenTelemetry

- Collection of APIs and SDKs, Data Model and semantic conventions (“k8s.pod.name”)
- Protocol (OTLP)
- Collector (receive, process and export telemetry data)

Instrumentation

Manual Instrumentation with OTEL SDK

```
func rolldice(ctx context.Context) int {  
    ctx, span := tracer.Start(ctx, "rolldice")  
    defer span.End()  
  
    roll := 1 + rand.Intn(6)  
    span.SetAttributes(attribute.Int("roll.value", roll))  
    return roll  
}  
  
func rolldiceHandler(w http.ResponseWriter, r *http.Request) {  
    ctx, span := tracer.Start(r.Context(), "rolldiceHandler")  
    defer span.End()  
  
    roll1 := rolldice(ctx)  
    roll2 := rolldice(ctx)  
    roll3 := rolldice(ctx)  
  
    io.WriteString(w, fmt.Sprintf("%d,%d,%d\n", roll1, roll2, roll3))  
}
```

Manual Instrumentation with OTEL wrappers

```
import io.opentelemetry.instrumentation.jdbc.datasource.OpenTelemetryDataSource;

@Configuration
public class DataSourceConfig {

    @Bean
    public DataSource dataSource() {
        BasicDataSource dataSource = new BasicDataSource();
        dataSource.setDriverClassName("org.postgresql.Driver");
        dataSource.setUrl("jdbc:postgresql://127.0.0.1:5432/example");
        dataSource.setUsername("postgres");
        dataSource.setPassword("root");
        return new OpenTelemetryDataSource(dataSource);
    }
}
```

Manual Instrumentation with OTEL wrappers

```
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(properties);  
KafkaProducer<String, String> producer = new KafkaProducer<>(properties);  
  
KafkaTelemetry telemetry = KafkaTelemetry.create(GlobalOpenTelemetry.get());  
Producer<String, String> tracingProducer = telemetry.wrap(producer);  
Consumer<String, String> tracingConsumer = telemetry.wrap(consumer);
```

Automatic Instrumentation with OTEL

- Captures telemetry data from popular libraries and frameworks
- Red Hat OpenShift distributed tracing data collection supports injecting auto-instrumentation agents for the following languages/applications (Dev Preview):
 - Java
 - .NET
 - NodeJS
 - Go
 - Python
 - Apache HTTPD

Distributed Tracing on OpenShift

New Distributed Tracing Stack on OpenShift

- **Instrumentation:** OpenTelemetry, Jaeger, OpenCensus or Zipkin
- **Collection:** OpenTelemetry Collector (managed by OpenShift distributed tracing data collection operator)
- **Storage:** Grafana Tempo (managed by Tempo Operator), object storage
- **Visualization:** Jaeger UI (managed by Tempo Operator)

Tempo Operator

- Deployment of Grafana Tempo instances
- Authentication and Authorization, Multitenancy
- Jaeger UI
- Managed upgrades

Red Hat OpenShift distributed tracing data collection

- Deployment of OpenTelemetry Collector instances as sidecar, daemon set or regular deployment
- Auto-Instrumentation injection
- Managed upgrades

OpenShift Service Mesh

- OpenShift Service Mesh supports creating spans of the interactions between services in the service mesh
- Services must propagate trace context between inbound and outbound requests in order to correlate spans



Example

Installed Operators

Installed Operators are represented by ClusterServiceVersions within this Namespace. For more information, see the [Understanding Operators](#)

Name ▼

/

Name ↕	Managed Namespaces ↕	Status
<div><div><div>Red Hat OpenShift distributed tracing data collection</div><div>0.81.0-2 provided by Red Hat</div></div></div>	All Namespaces	<div><div>✓</div>Succeeded Up to date</div>
<div><div><div>Tempo Operator</div><div>0.3.0-2 provided by Red Hat</div></div></div>	All Namespaces	<div><div>✓</div>Succeeded Up to date</div>

Storage Configuration for Tempo

```
apiVersion: v1
kind: Secret
metadata:
  name: tempo-storage
type: Opaque
stringData:
  endpoint: http://minio:9000
  bucket: tempo
  access_key_id: tempo
  access_key_secret: supersecret
```

Tempo Deployment

```
apiVersion: tempo.grafana.com/v1alpha1
kind: TempoStack
metadata:
  name: prod
spec:
  storage:
    secret:
      name: tempo-storage
      type: s3
  storageSize: 1Gi
  template:
    queryFrontend:
      jaegerQuery:
        enabled: true
      ingress:
        type: route
```

```
apiVersion: opentelemetry.io/v1alpha1
kind: OpenTelemetryCollector
metadata:
  name: otel
spec:
  config: |
    receivers:
      otlp:
        protocols:
          grpc:

    exporters:
      otlp:
        endpoint: tempo-prod-distributor:4317
        tls:
          insecure: true

    processors:
      batch:

  service:
    pipelines:
      traces:
        receivers: [otlp]
        processors: [batch]
        exporters: [otlp]
```

OpenTelemetry Collector Deployment

Auto-Instrumentation

```
apiVersion: opentelemetry.io/v1alpha1
kind: Instrumentation
metadata:
  name: petclinic
spec:
  exporter:
    endpoint: http://otel-collector:4317
```


PetClinic Deployment

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: petclinic
spec:
  selector:
    matchLabels:
      app: petclinic
  template:
    metadata:
      labels:
        app: petclinic
      annotations:
        instrumentation.opentelemetry.io/inject-java: petclinic
    spec:
      containers:
        - image: springcommunity/spring-framework-petclinic:6.0.3
          name: petclinic
```

New Pet

Owner

George Franklin

Name

Birth Date

Type

bird

cat

dog

hamster

lizard

Add Pet

Search

Upload

Service (1)

petclinic ▾

Operation (49)

all ^

all

CrashController.triggerException

GET /

GET /oups

GET /owners

GET /owners/{ownerId}

Max Duration

Min Duration

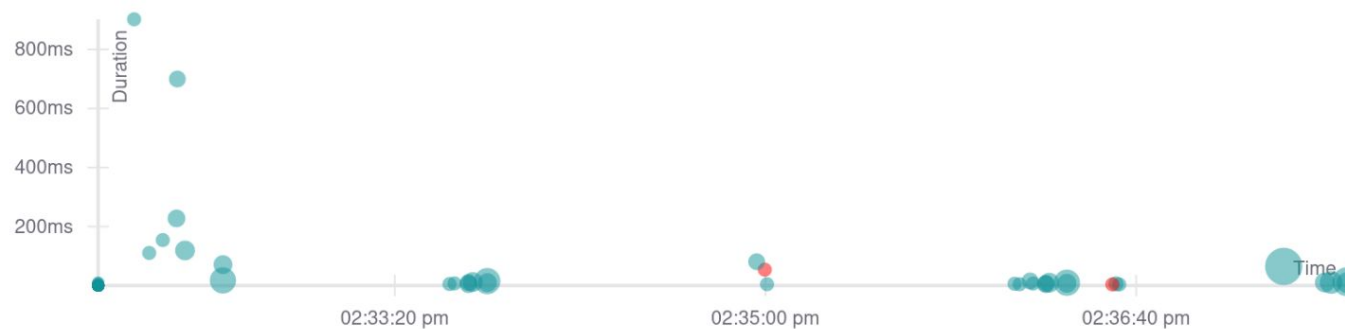
e.g. 1.2s, ...

e.g. 1.2s, ...

Limit Results

200

Find Traces



107 Traces

Sort: Most Recent ▾

Download Results

Deep Dependency Graph

Compare traces by selecting result items

☐ petclinic: GET /owners/{ownerId} a52d96c

9.85ms

10 Spans

petclinic (10)

Today |

2:37:36 pm

4 minutes ago

☐ petclinic: POST /owners/{ownerId}/pets/new 4c06eac

13.17ms

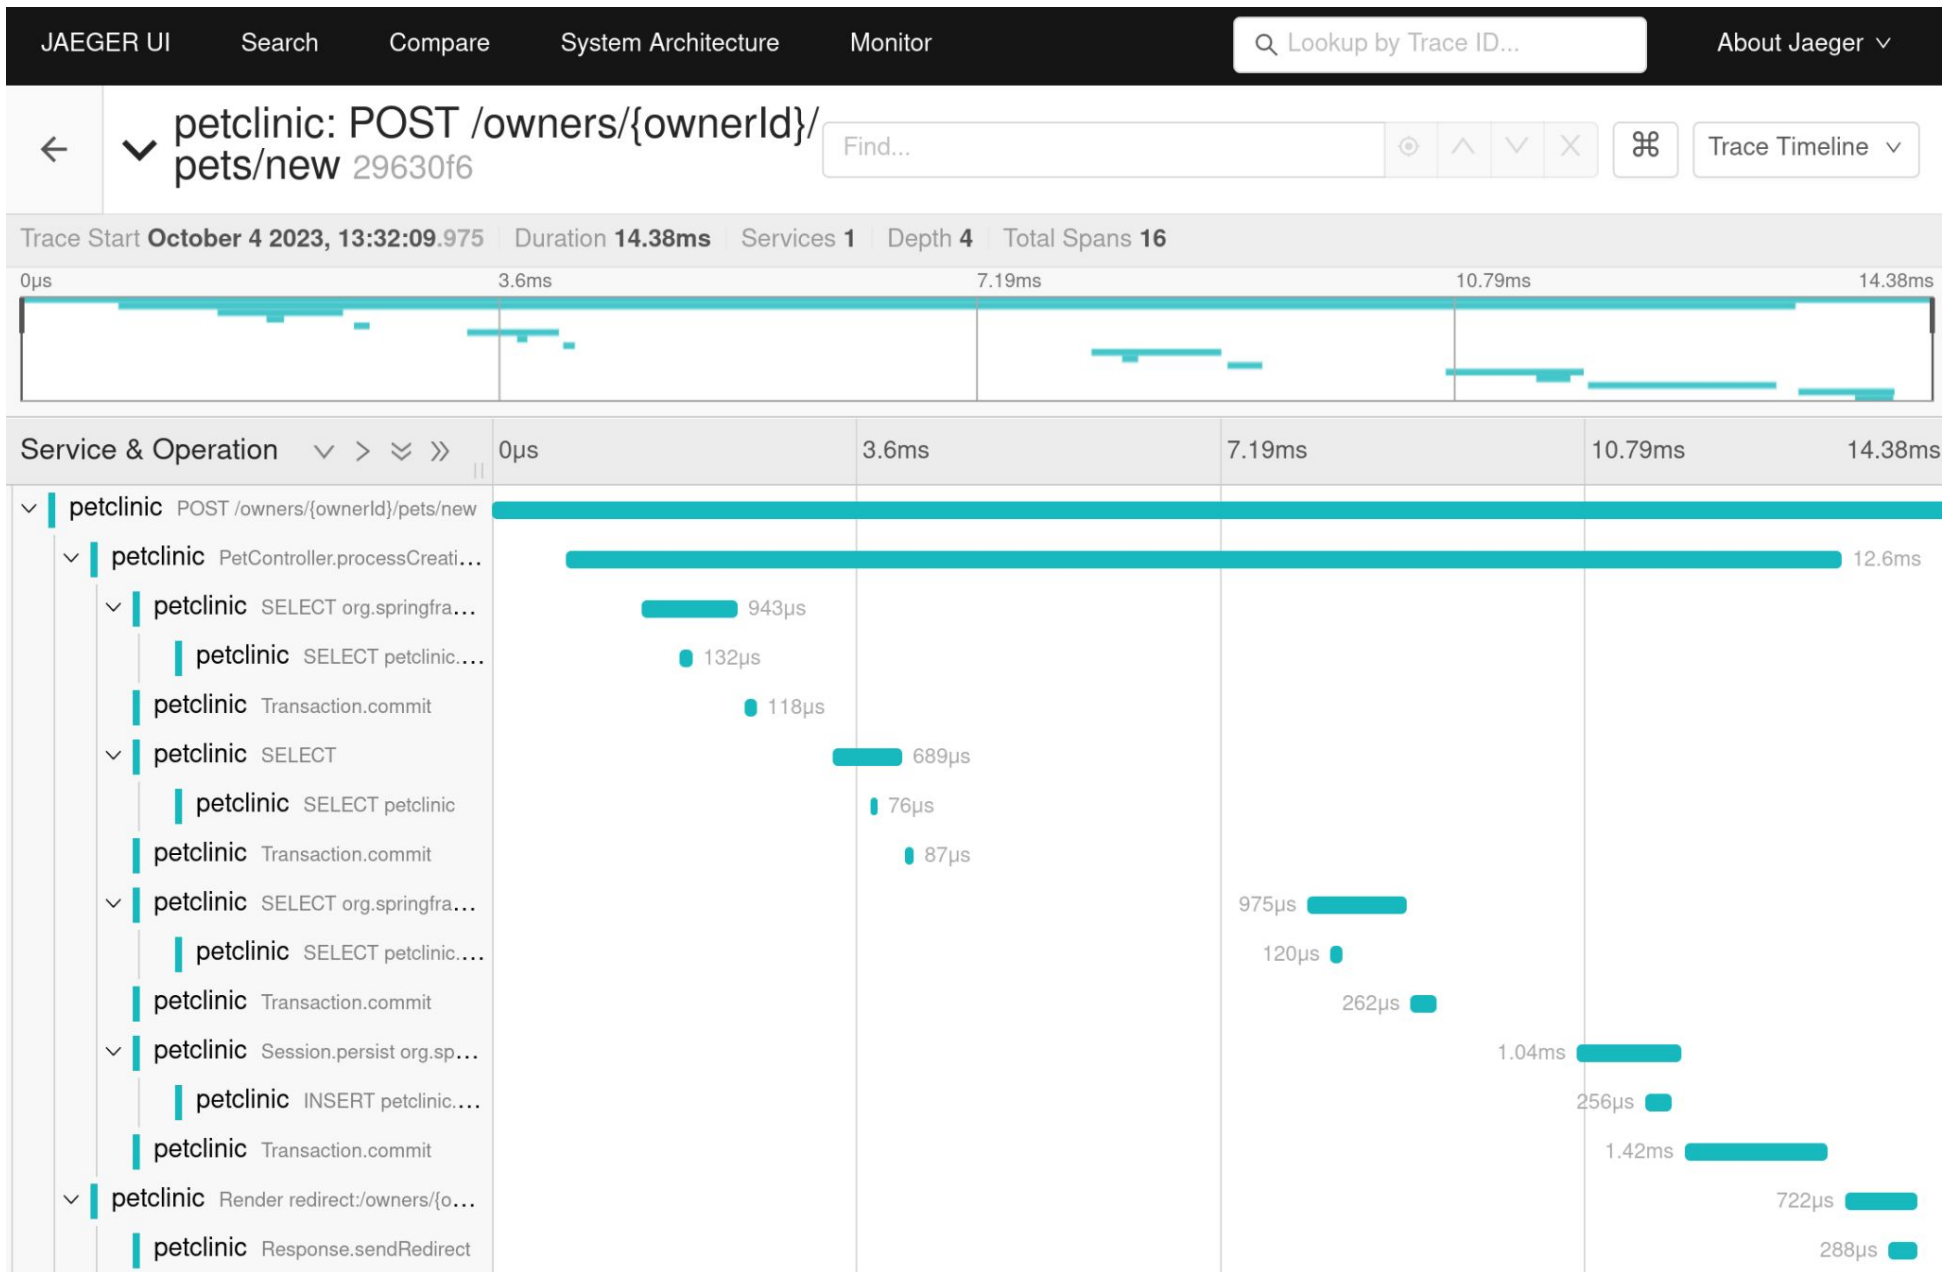
19 Spans

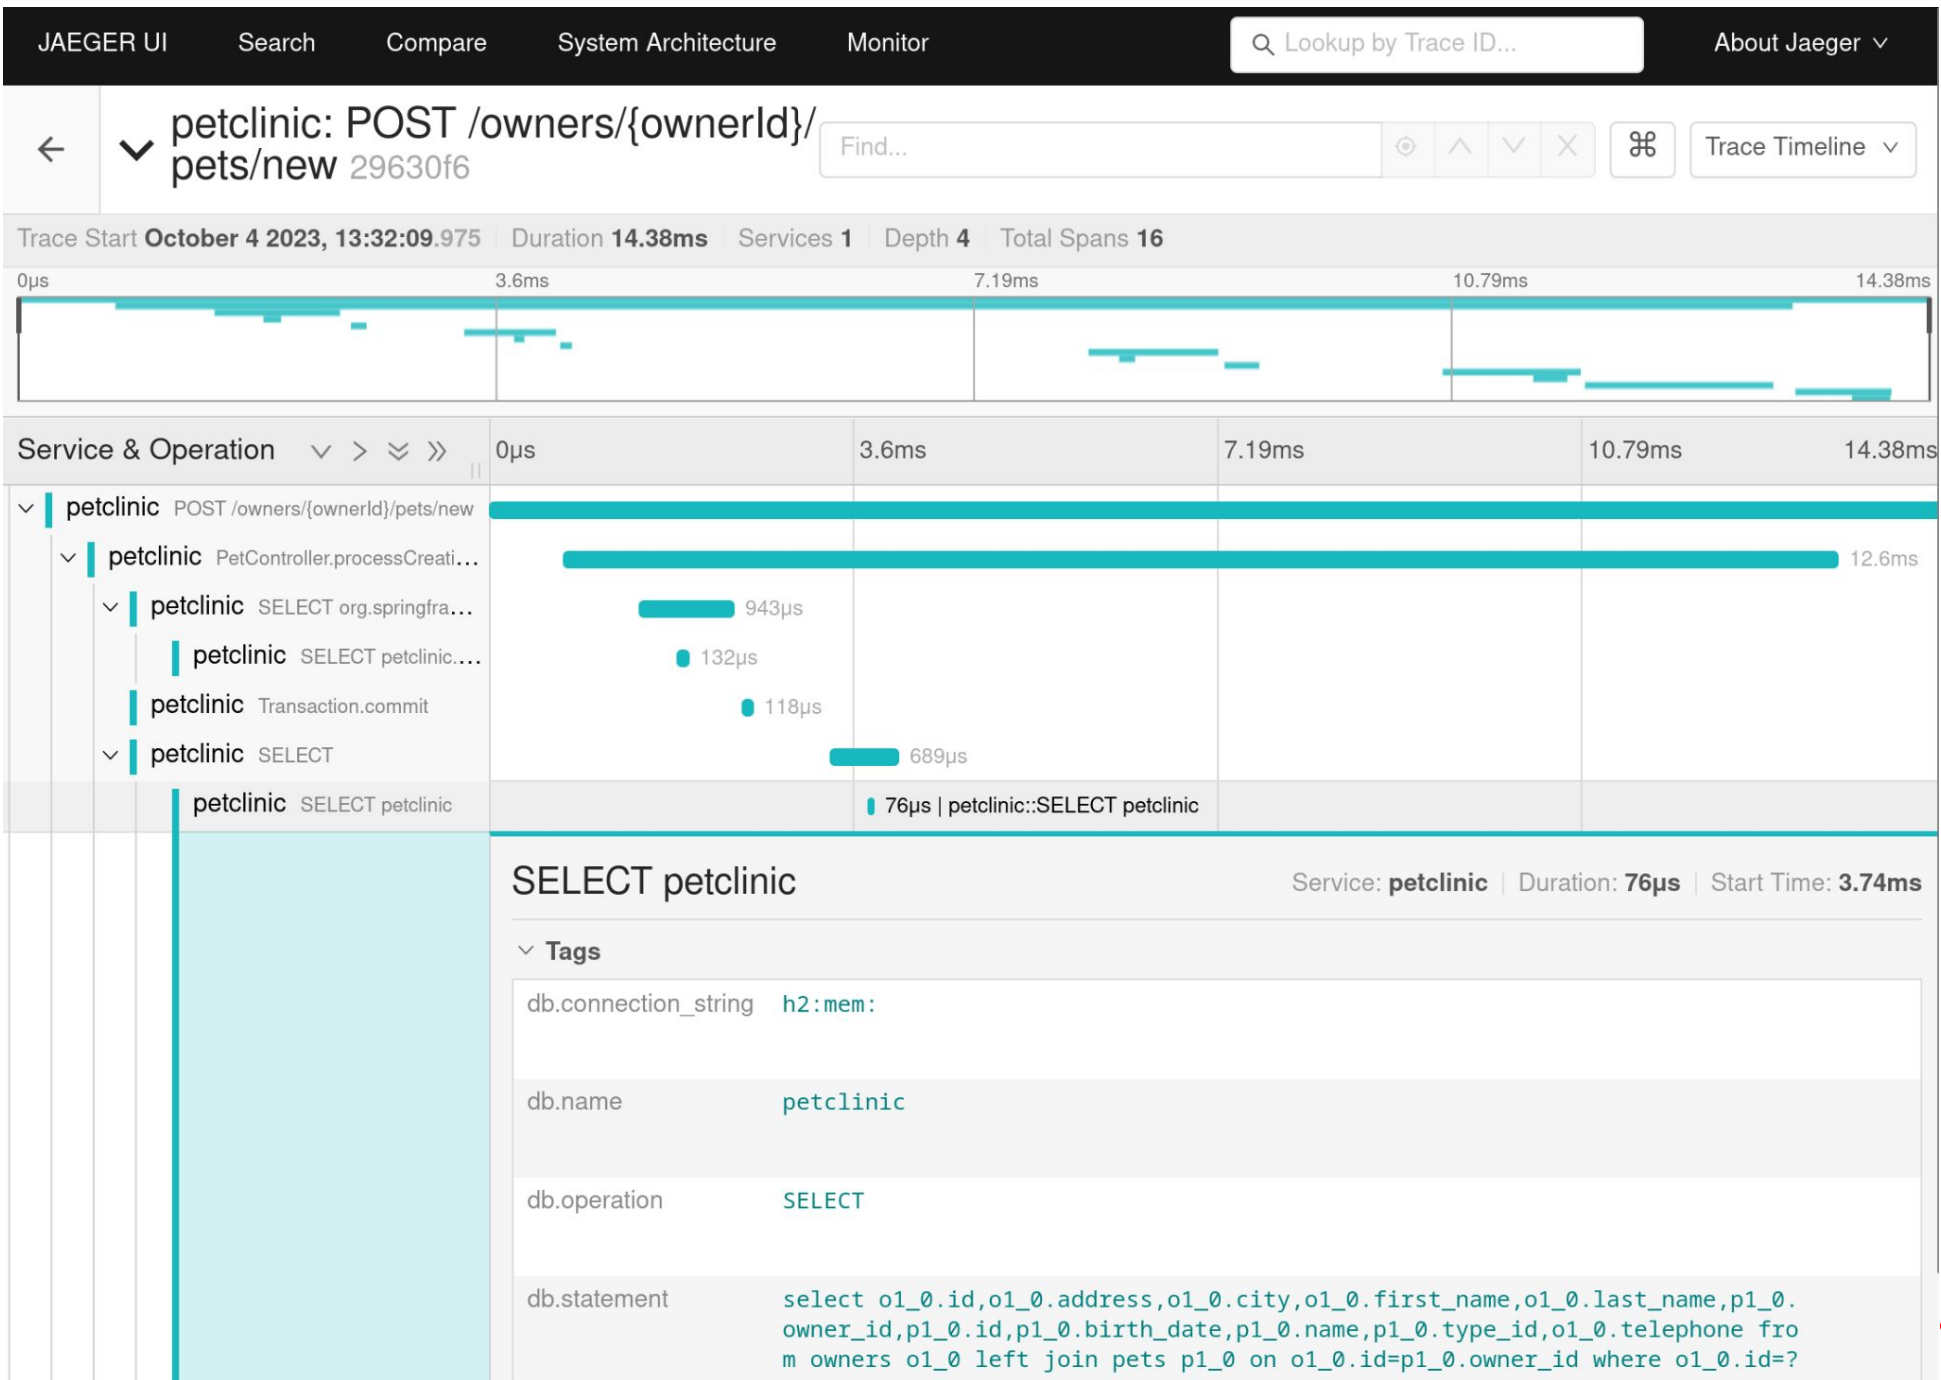
petclinic (19)

Today |

2:37:36 pm

4 minutes ago





Java Auto-Instrumentation: Behind the Scenes

The OpenShift distributed tracing data collection operator performs the following modifications to a pod:

- It attaches a new emptyDir volume
- It adds a new init container, which copies javaagent.jar to this volume
- This volume is mounted in the container of the application
- The JAVA_TOOL_OPTIONS environment variable is modified to load javaagent.jar

Q & A

Red Hat
Summit

Connect

Thank you



linkedin.com/company/red-hat



facebook.com/redhatinc



youtube.com/user/RedHatVideos



twitter.com/RedHat