

March 8th, 2018 – QCon London



How to build observable
distributed systems



@PierreVincent



pvincent.io





Pierre Vincent

SRE Manager at Poppulo

 @PierreVincent

 pvincent.io

Poppulo

 @PierreVincent

**Reaching production is
only the beginning**


 [@PierreVincent](https://twitter.com/PierreVincent)




**No system is immune to failure
Be ready to recover**



**When distributing a system,
we're also distributing the places
where things might go wrong**



**Monitoring only applies to
known failure modes.
What about everything else?**




“ Monitoring tells you whether the system works. Observability lets you ask why it's not working. ”

– Baron Schwarz



Healthchecks



Metrics



Logging



Tracing

Healthchecks





Healthchecks

Is it running?


**Can it perform
its task?**

**Can it accept
more work?**




Healthchecks


Broadcast



Register




Expose





Healthchecks




```
GET http://1.2.3.4:8080/health
```

```
200 OK
```


```
{  
  "service": "registration-service",  
  "healthy": true,  
  "workload": { "healthy": true },  
  "dependencies": [  
    { "name": "cassandra", "healthy": true },  
    { "name": "billing-svc", "healthy": true  
  },  
  ]  
}
```


Overzealous Healthchecks can be counter-productive



Source: HTTP Healthchecks for a Resilient Platform - **Chris O'Dell**
skeltonthatcher.com/blog/http-healthchecks-for-a-resilient-platform

Metrics





Metrics



**System
metrics**


**Application
metrics**

**Business
metrics**

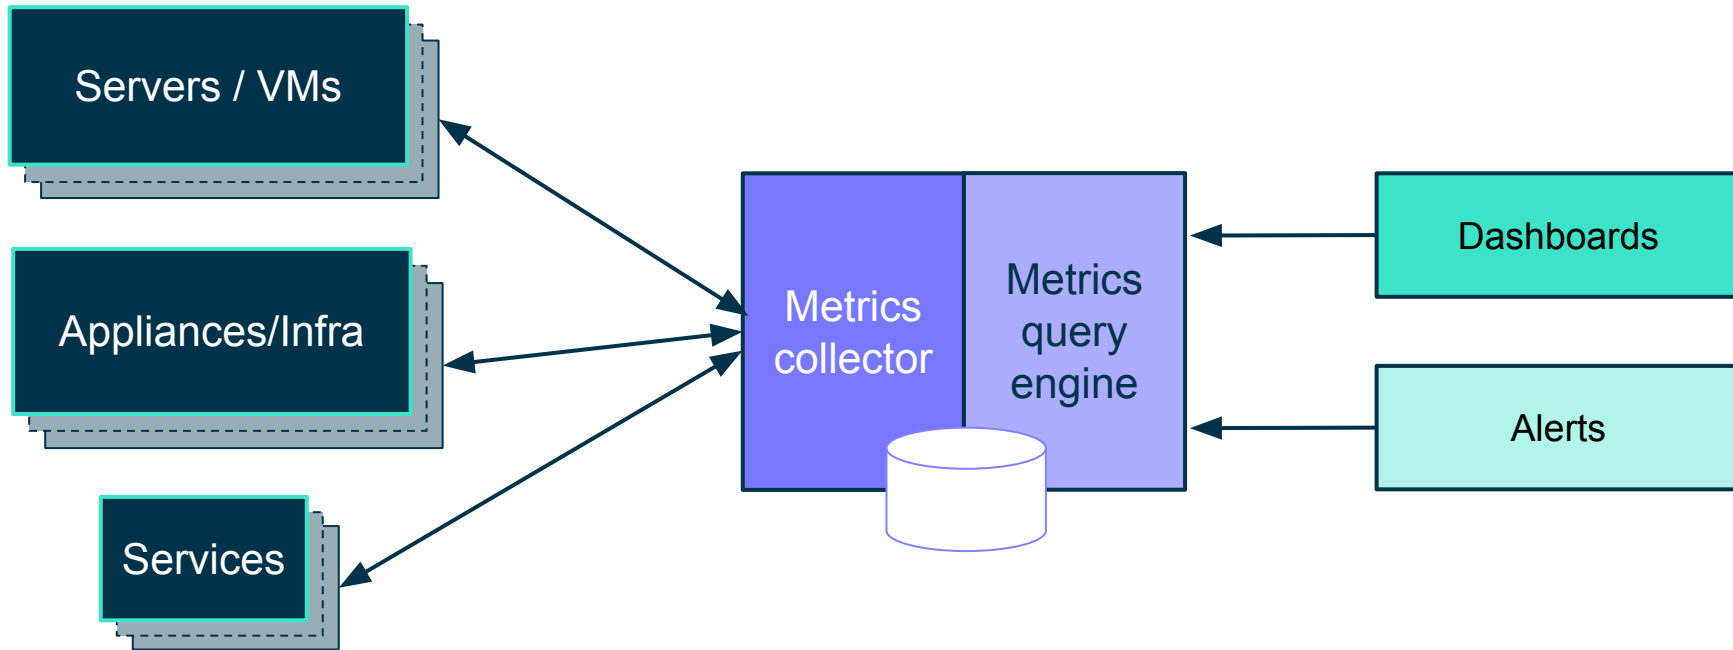
CPU usage


Error rates

Customer conversions

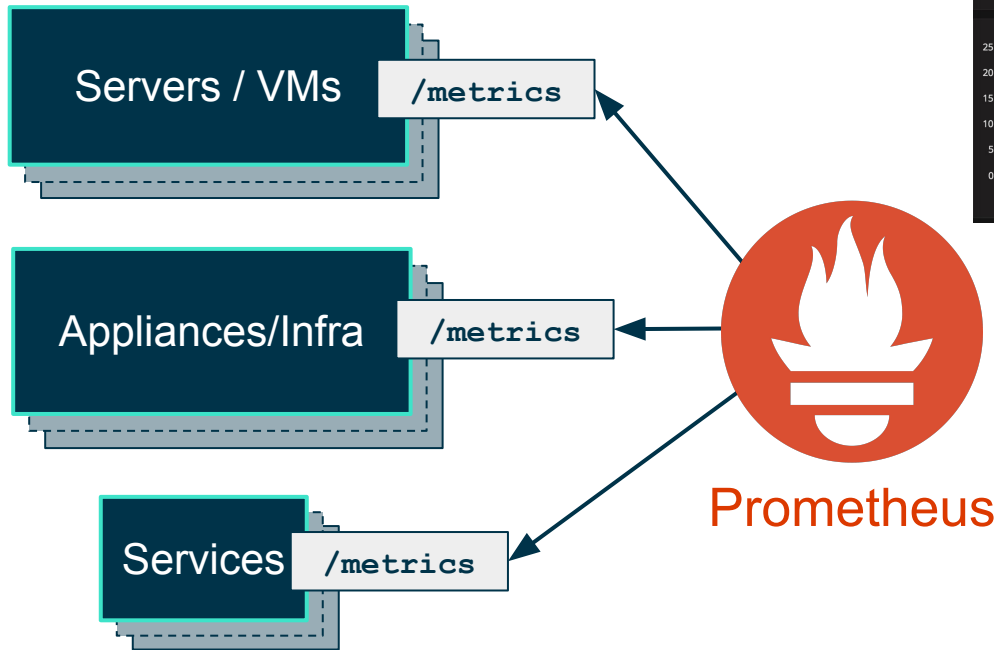


Metrics






Metrics



Prometheus





Watch out for over-reliance on metrics

Not every metric
deserves an alert

**Limit alerting to
user-impacting symptoms**

Limitations at high-cardinality

Poor fine-grained debugging
e.g. CustomerId

Real-time querying means
some trade-offs on retention

**Not suitable for long-term
trend analysis**

Logging





Logging

Making sense of **(a lot of)** logs

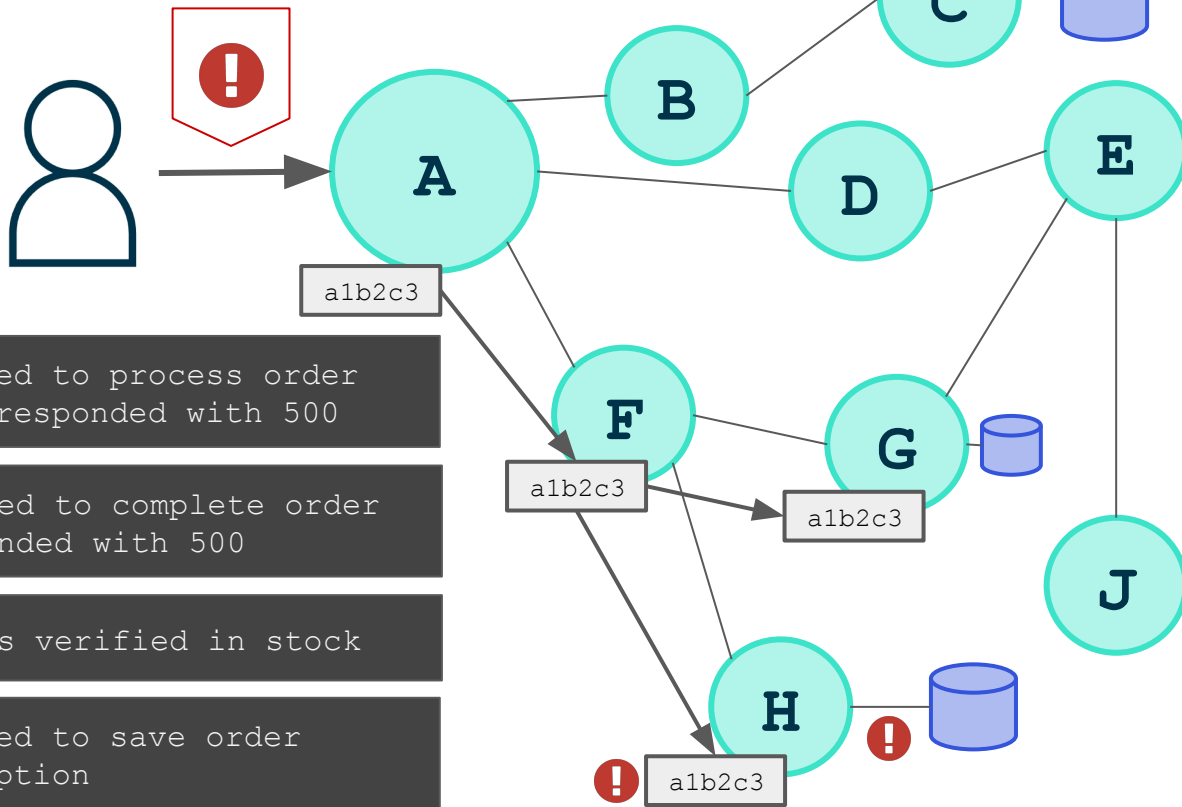
Centralised

Searchable

Correlated



Log Correlation



```
ERROR [svc=A][trace=a1b2c3] Failed to process order  
Cause: Order process manager responded with 500
```

```
ERROR [svc=F][trace=a1b2c3] Failed to complete order  
Cause: Shipping service responded with 500
```

```
INFO [svc=G][trace=a1b2c3] Items verified in stock
```

```
ERROR [svc=H][trace=a1b2c3] Failed to save order  
Cause: Cassandra timeout exception
```

2018-02-20T16:38:23+00:00

ERROR

Read timed out

Hmmm thanks... ?

When did it happen?

timestamp 2018-02-20T16:38:23+00:00

What is the message?

level ERROR log Read timed out

What is it?

service registration-service team events

What is it running?

commit 542a8b8e build 542a8b8e.7 runtime java-1.8.0_161

Where is it?

region europe-west2 node node_e79f3e52

Who caused it?

customerID 55123 userID 458

Can I trace it?

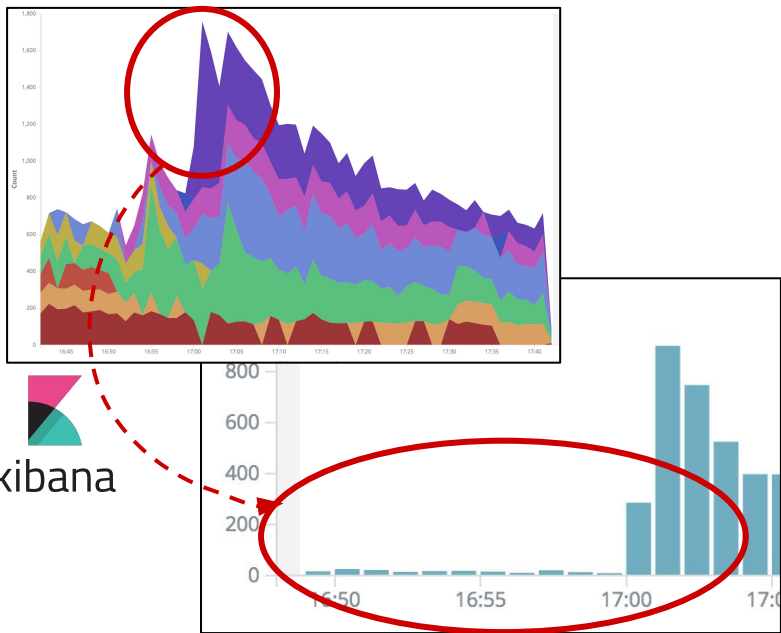
requestID ec667cb45

Any other info?

... ...

JSON

Structured logs unleash high-cardinality exploration



Activity spike isolated for single customer

Error rate spike isolated by build version



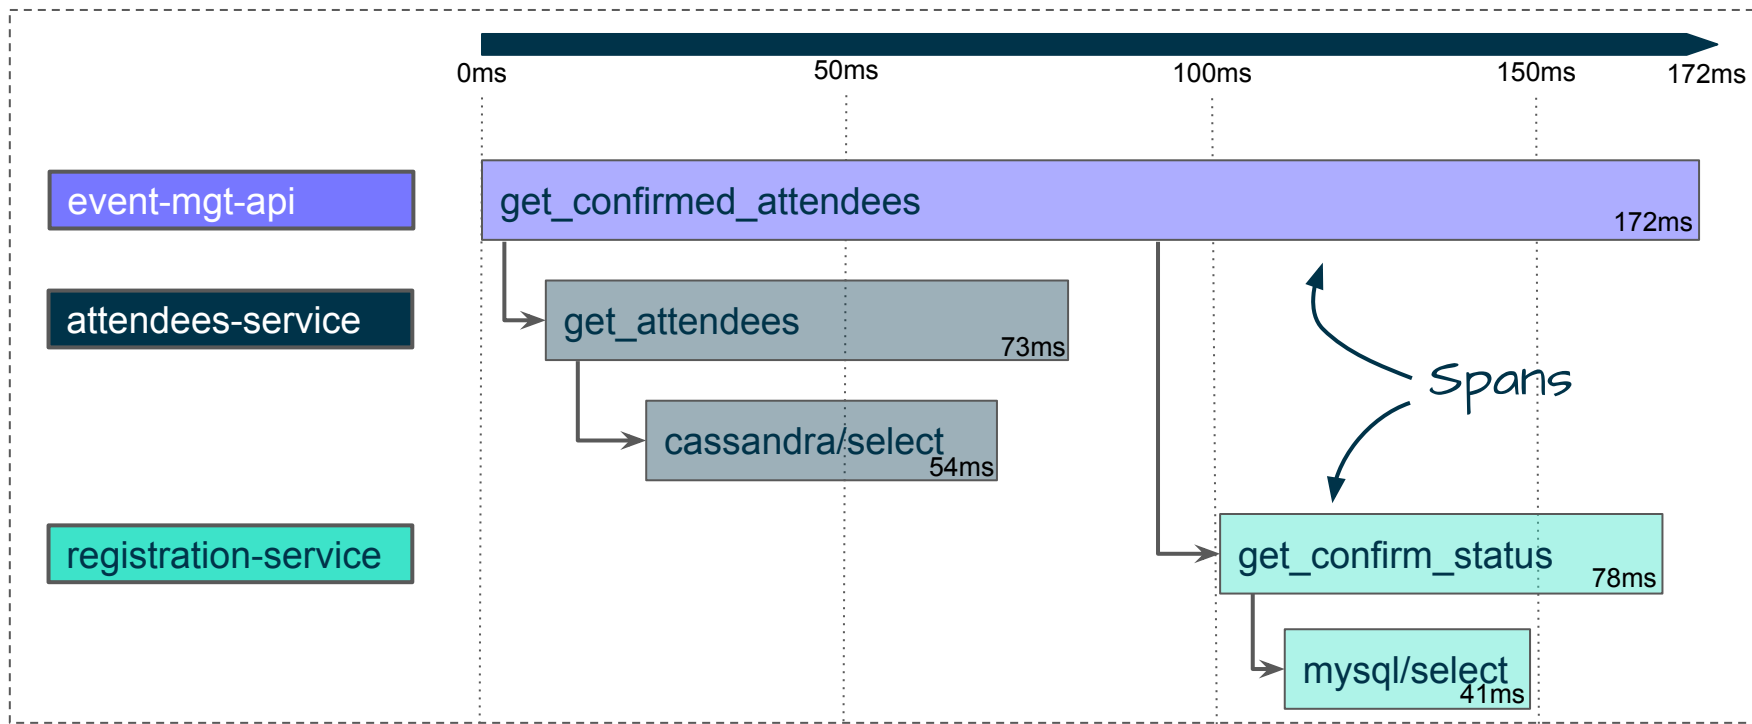
Tracing





Tracing

Trace



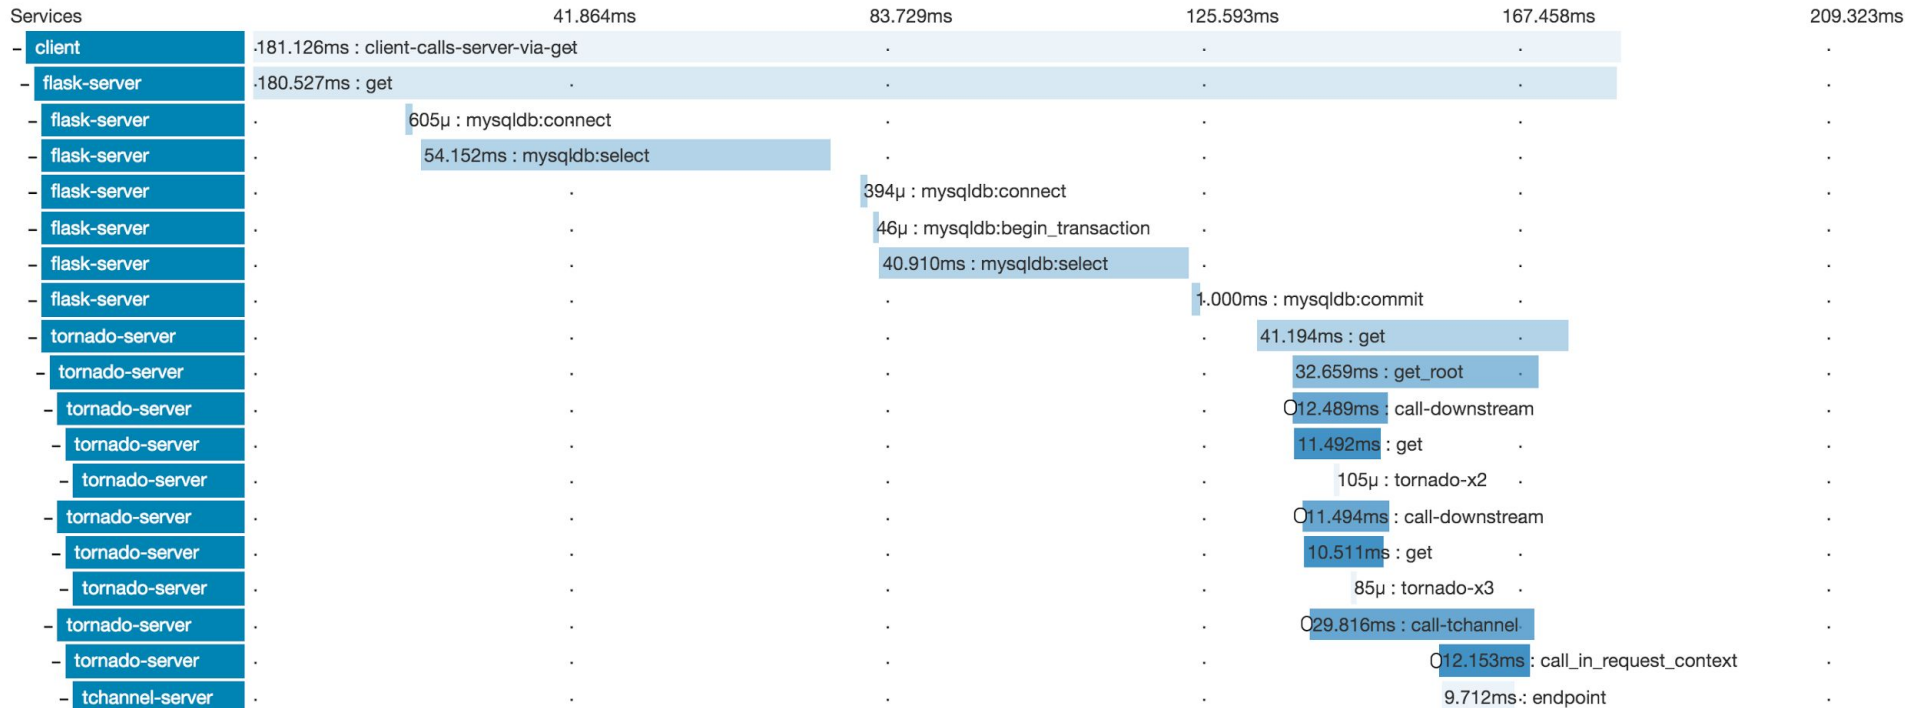
Duration: 209.323ms Services: 5 Depth: 7 Total Spans: 24

Expand All Collapse All Filter Service Se...

client x4 flask-server x10 missing-service-name x2 tchannel-server x2 tornado-server x11



JSON



Known Unknowns

Unknown Unknowns



Healthchecks


Metrics (Alerting)

Metrics (Queries)

Tracing

Logs

Events



Monitoring & Resiliency

Debugging & Exploration

Usability of tooling is key to adoption



Cheap to
instrument



Easy to
explore



Reliable &
Trustworthy

Visibility builds trust but requires safety





Visibility helps justify decisions



Visibility enables operability

Don't spend all your time testing

...keep some for instrumenting

Thank you!



 @PierreVincent

 pvincent.io

Poppulo

Thank you!



Pierre Vincent
SRE Manager at Poppulo

 @PierreVincent

 pvincent.io

Poppulo

 @PierreVincent