



Evolution of Observability Tools at Pinterest

Naoman Abbas
Engineering Manager, Observability

Naoman Abbas

Engineering Manager Observability

- Operational Metrics
- Alerting
- Log Search
- Distributed Tracing




Pinterest

Helping people
discover and do what
they love

+300M monthly active users

+200B pins saved

+2000 employees




Agenda

- 1 **Why Observability?**
- 2 **Observability Tools**
- 3 **Evolution**
- 4 **Lessons Learned**


Why Observability?



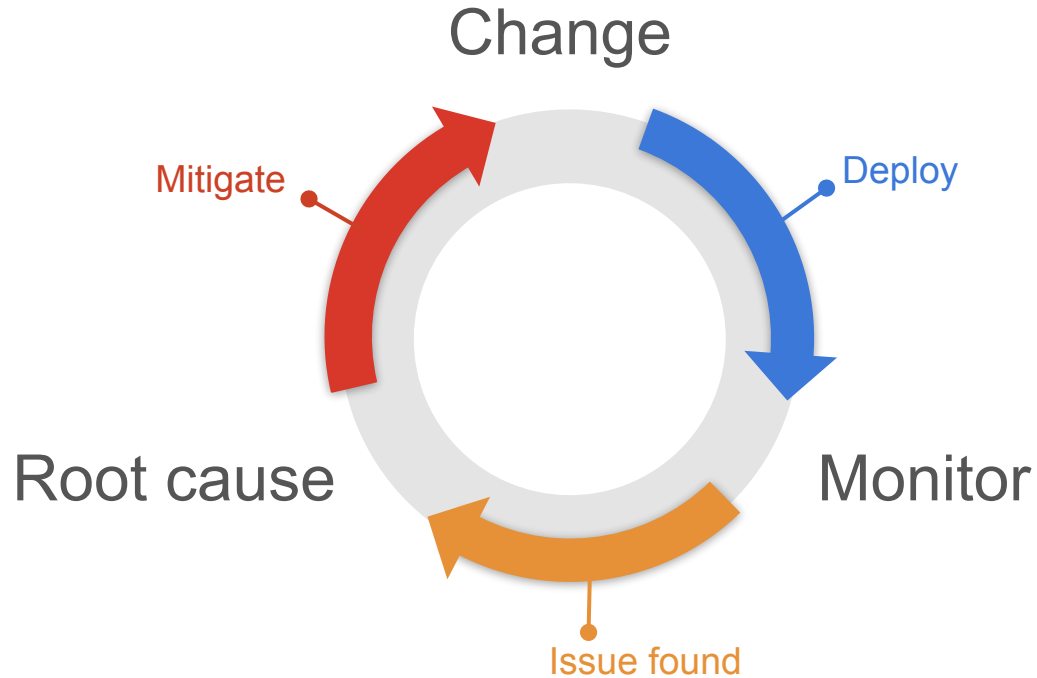


Microservice Architecture





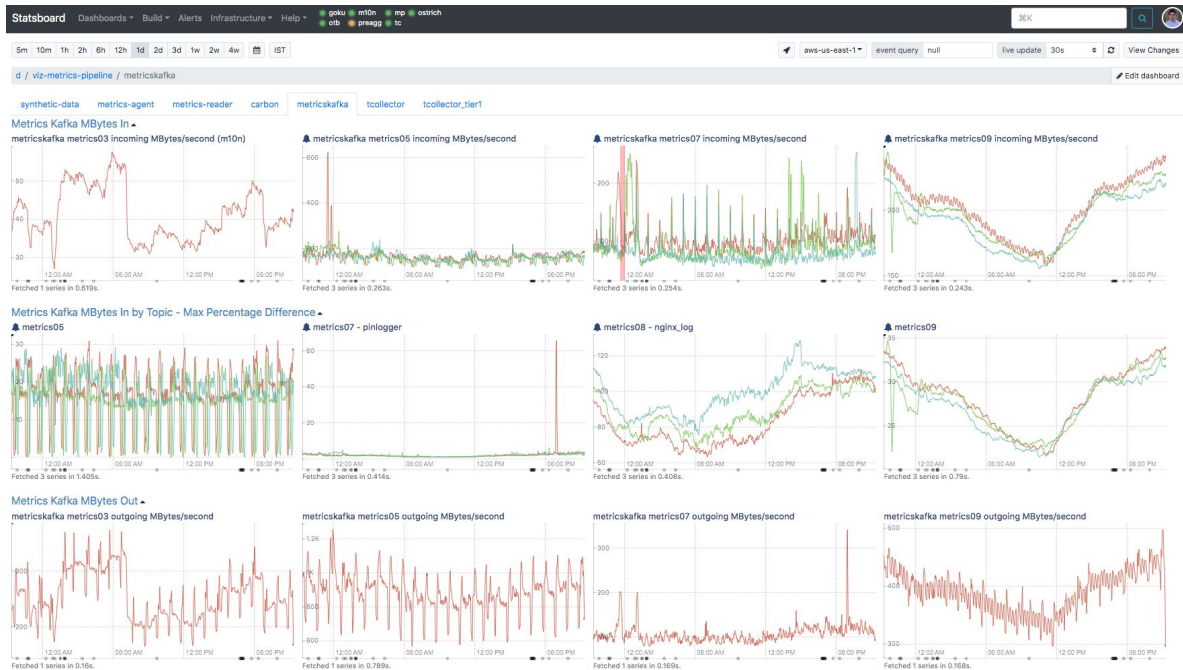
Change Cycle



Observability Tools



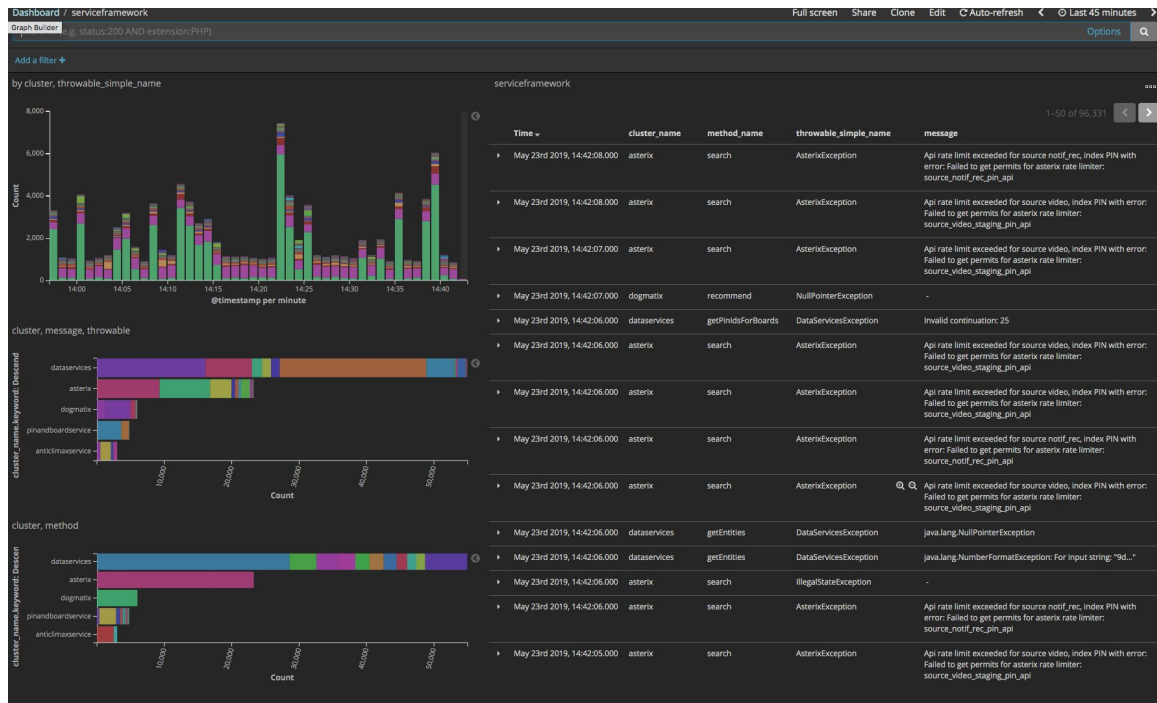
Statsboard



Operational metrics

- Service dashboards
- Alerts
- Debugging
- Performance tuning
- **2,000** dashboards
- **16,000** alerts

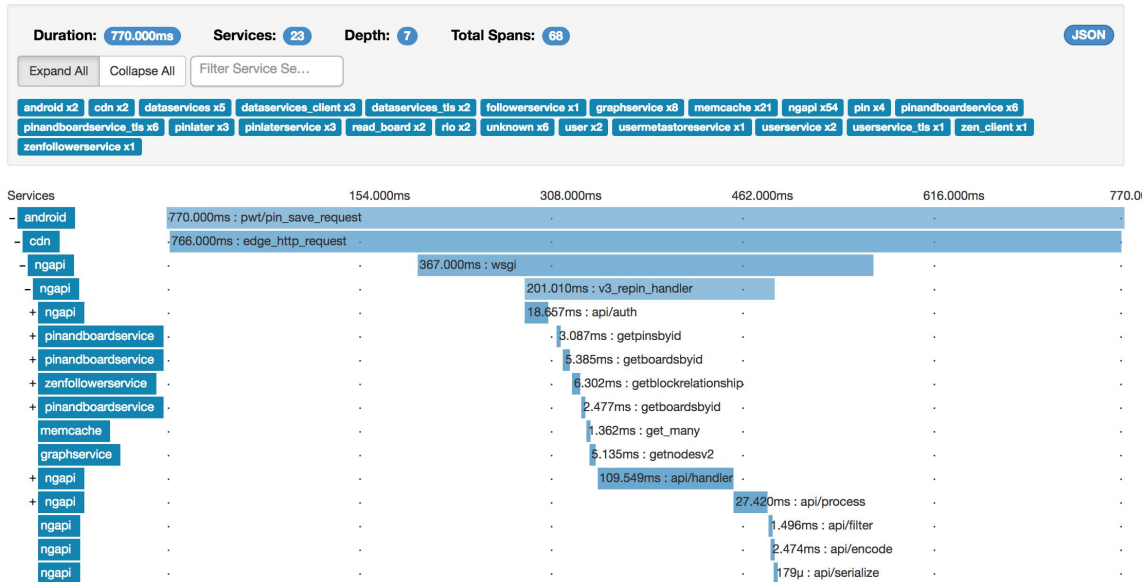
Logsearch



Debug logs

- Root cause analysis
- Alerting

Pintrace



Distributed tracing

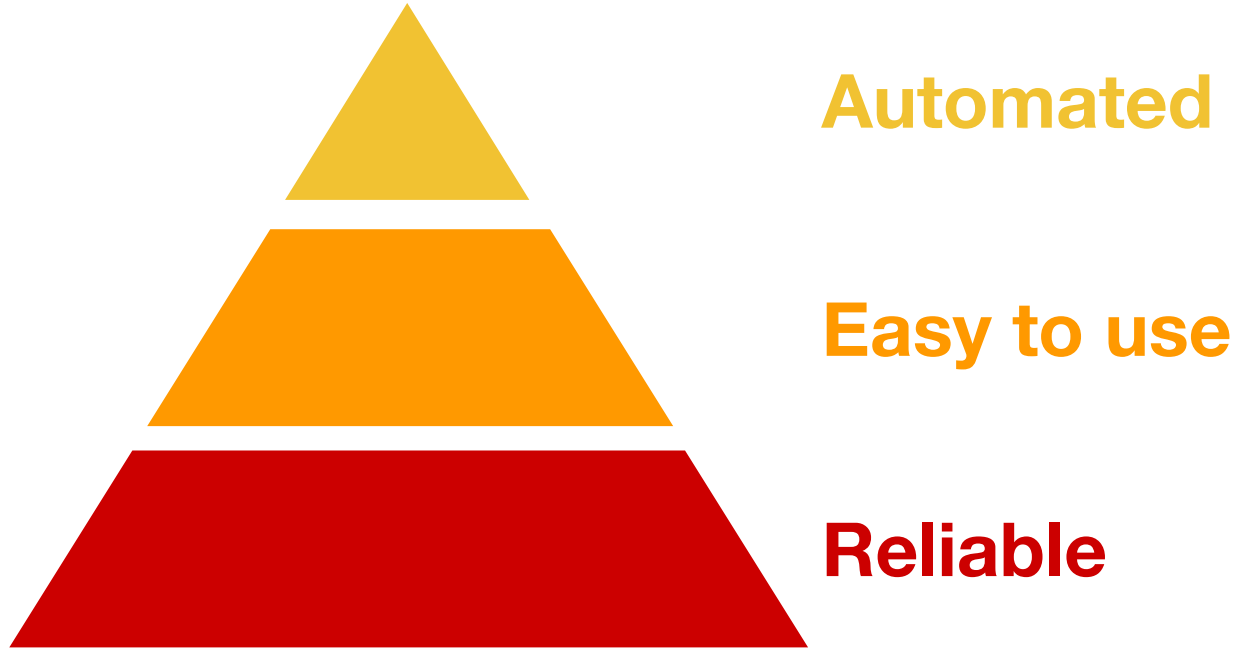
- Performance tuning
- Root cause analysis



Evolution



Observability Needs




Tools

Reliability

Usage Growth

Metrics per Minute



Tool Architecture

- **Metrics Storage**

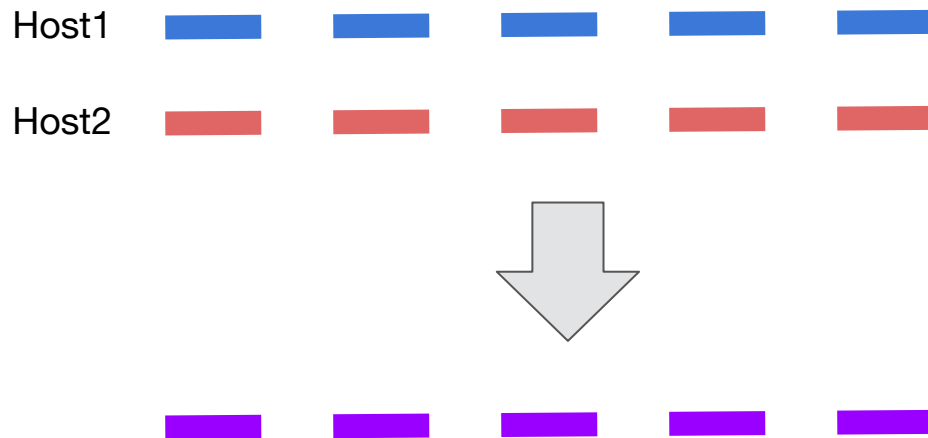
- Graphite -> OpenTSDB -> Sharding -> Goku (in-memory storage)

- **Metrics Processing**

- Storm -> Spark Streaming -> Job Stream (custom streaming)


Data Reduction

Metrics Aggregation by Service



Data Reduction

Chargeback



Tools Usability

TScript

Scripting Language for Time-series

Before

```
divideSeries (abs (diffSeries (timeShift  
(tc.kafka.stats.kafka.server.brokert  
opicmetrics.mbytesinpersec.perTopic.  
OneMinuteRate.metrics07{topic=pinlog  
ger}, '7d'), tc.kafka.stats.kafka.serv  
er.brokertopicmetrics.mbytesinpersec  
.perTopic.OneMinuteRate.metrics07{to  
pic=pinlogger}))), 100)
```

After


```
today = metric  
one_week = today.timeShift(Week)  
return today.pctDiff(one_week).abs()
```

Integrated Alerts

Health Score (combination of Correctness, Query SR, and Write Delay) ▲



Fetches 2 series in 0.223s.



Fetches 3 series in 0.209s.



Fetches 2 series in 0.209s.

Quick Dash


5m 10m 1h 2h 6h 12h 1d 2d 3d 1w 2w 4w  PDT

 event query live update 30s  

View Dashboard  Save Changes


Kafka 

Kafka Mbps




Fetches 13 series in 0.13s.

Latency




Fetches 8 series in 0.121s.




JVM 

CPU



Fetches 4 series in 0.126s.



Add Section

RunDash

Visibility spark runbook

All metrics in this dashboard are expressed as a per-minute rate

[show / hide runbook](#)

viz-eagleboard-aggregation ▾

Kafka writes by upstream



Fetches 1 series in 0.7s.

Kafka reads by spark



Fetches 3 series in 0.7s.

Kafka writes by spark




Fetches 2 series in 0.128s.

Kafka messages in kafka



Fetches 1 series in 0.14s.

Kafka reads by ingestors



Fetches 4 series in 0.322s.

[More on viz-eagleboard-aggregation](#)

viz-edge-requests-aggregation ▾

Kafka writes by upstream



Fetches 1 series in 0.128s.

Kafka reads by spark



Fetches 5 series in 0.29s.

Kafka writes by spark



Fetches 2 series in 0.129s.

Kafka messages in kafka



Fetches 1 series in 0.137s.

Kafka reads by ingestors



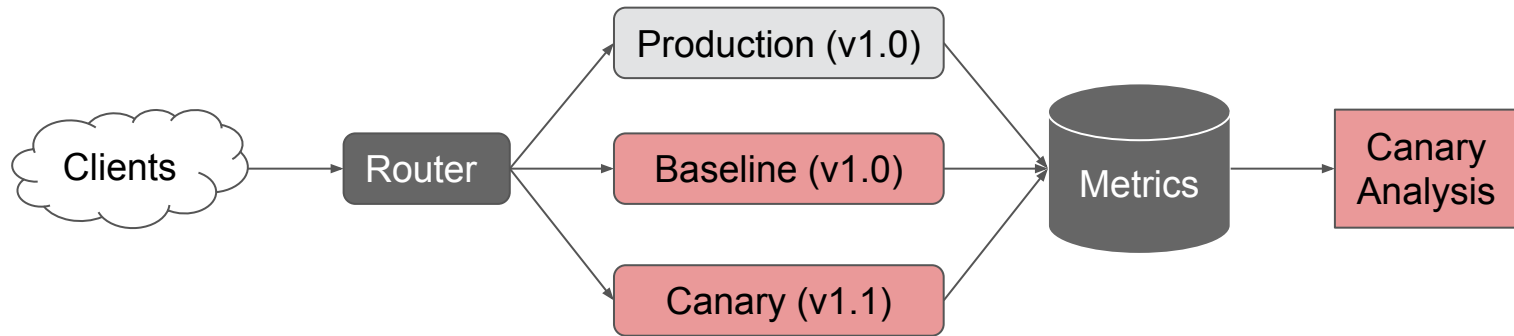
Fetches 4 series in 0.175s.

Tools Automation

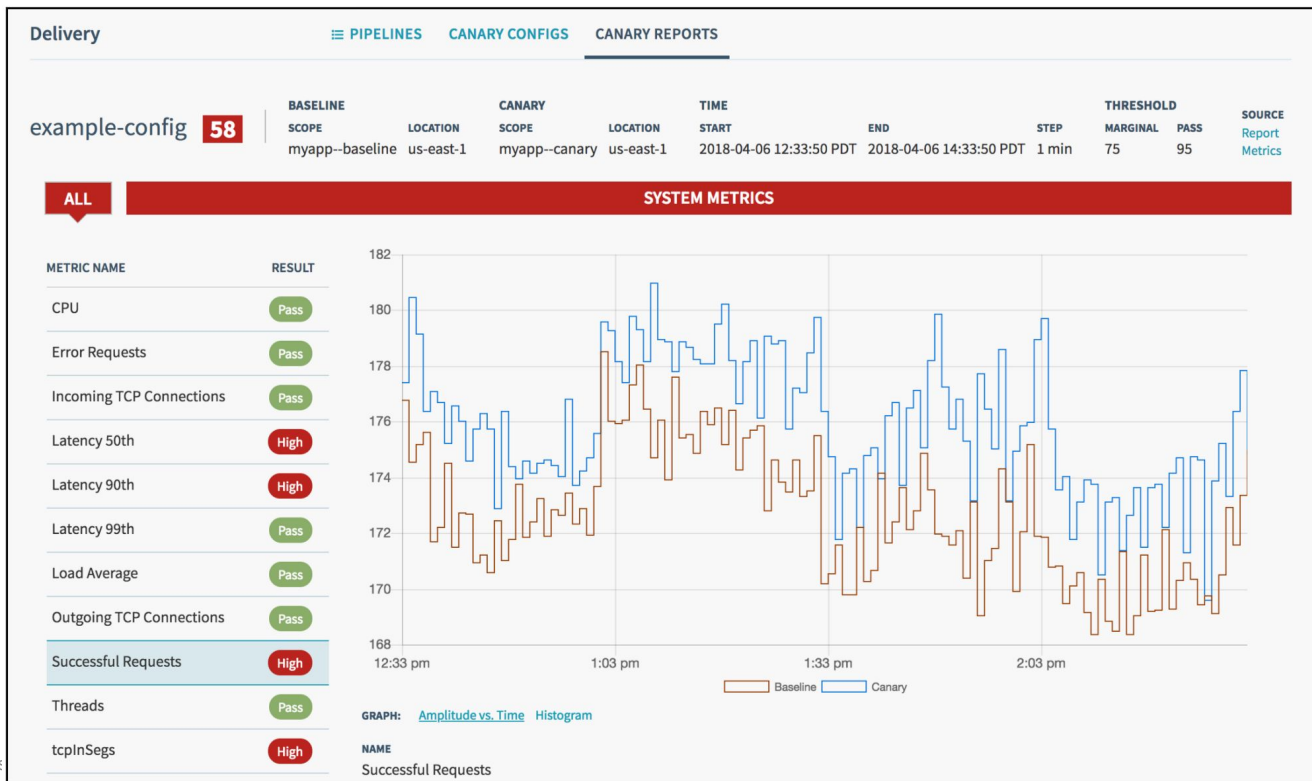
**Roughly 70% of
outages are due to
changes in a live
system**

Google SRE Book


Automated Canary Analysis



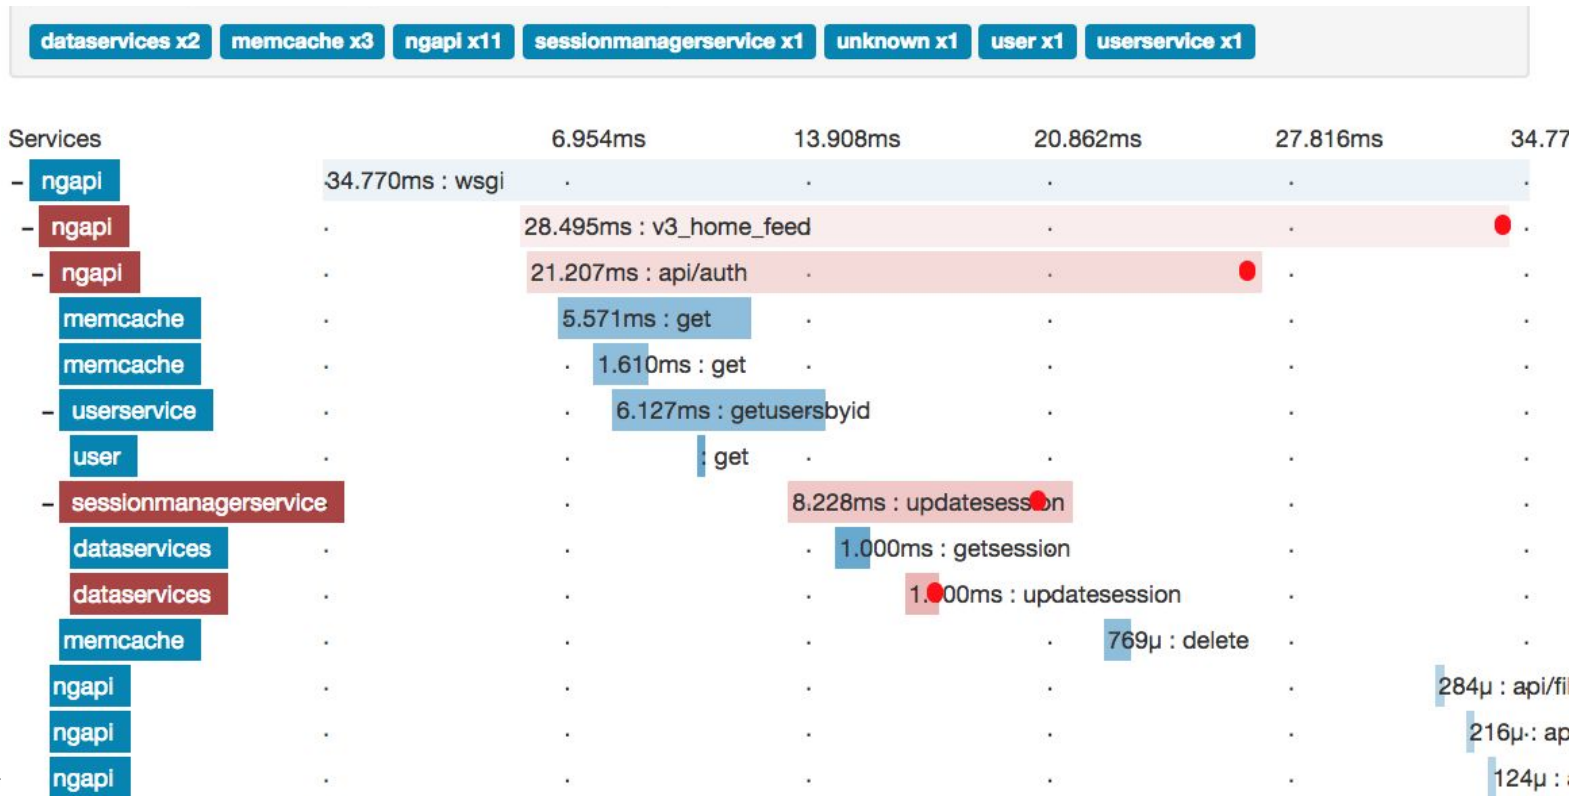
Automated Canary Analysis



Automated Root Cause Analysis

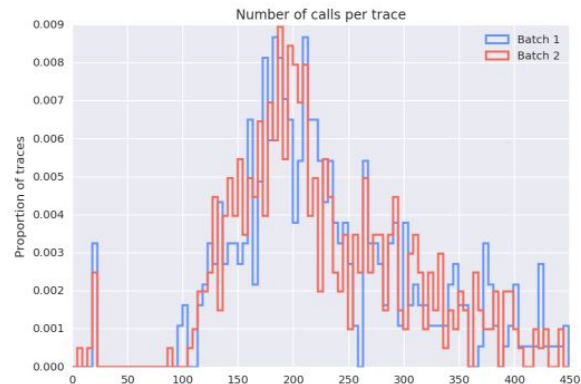
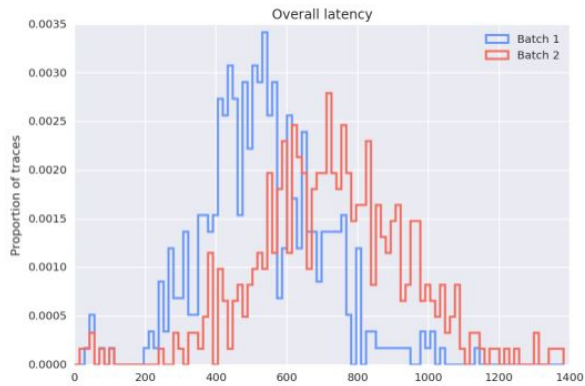



Automated Root Cause Analysis



Automated Root Cause Analysis

Trace Analyzer



Automated Root Cause Analysis

Trace Analyzer

Downstream Services Latency

	First traces	Second traces	Difference
smartfeedservice	99.3975	118.445	19.0478
pinacle_p2p	88.9196	92.2861	3.36652
smartfeedgenerator	67.0345	89.5526	22.5181
tenansix	65.3212	52.185	-13.1362
instantpfyservice	10.4688	10.6408	0.172027
pinlaterservice	10.4142	10.8673	0.453135
usercontextservice	9.80483	10.8275	1.02263
pinlinks	9.07438	8.88087	-0.193505
metatron	7.60317	8.72428	1.12111

Downstream Services Calls

	First traces	Second traces	Difference
metatron	764	1350	586
pinandboardservice	5526	4891	-635
pinacle_p2p	1924	1723	-201
terrapinthrift	596	400	-196
pinacle2	530	236	-294
usercontextservice	79	65	-14
anticlimaxservice	34	29	-5
asterix	32	26	-6

Automation Roadmap

- **Anomaly detection**
- **Auto remediation**
- **Error budgets (SLO/SLI)**

Lessons Learned



Lessons Learned

- **Avoid tool fragmentation**
- **Observability is challenging**
- **Observability is expensive**
- **Great ROI**


Observability Team



Brian
Overstreet




Colin
Probasco



Dai Nguyen



Humsheen
Geo



Peter Kim



Wei Zhu

Acknowledgements

- **Storage and Caching team (HBase, Goku)**
- **Logging team (Kafka)**
- **BDP team (Compute platform for Spark)**

We're hiring! Come work with us!



Scan me

hiring-srecon@pinterest.com



Questions

