



# Finding the needle in a Haystack

**Shreya Sharma / Technical Product Manager / [shreyas@usc.edu](mailto:shreyas@usc.edu)**

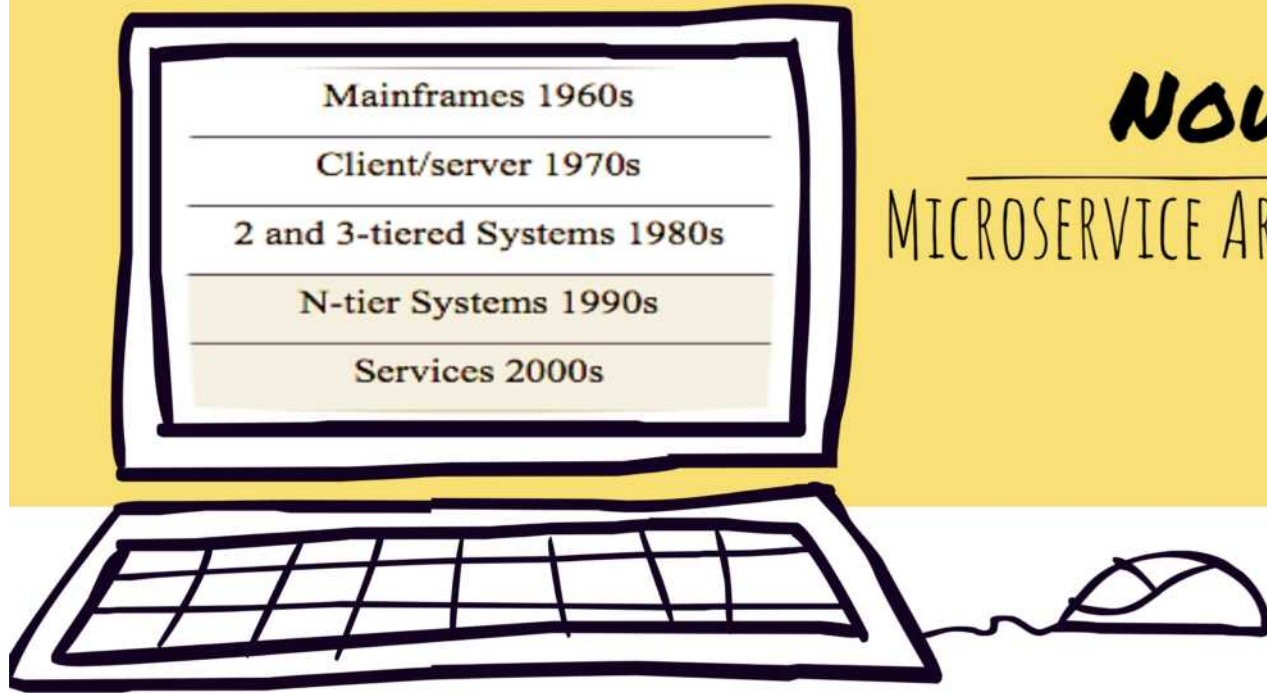
# The Who

The logo for Expedia Group, featuring the words "expedia group" in a white, lowercase, sans-serif font. A small "TM" trademark symbol is positioned to the upper right of the word "group". The text is centered within a solid blue rectangular background.

expedia group™

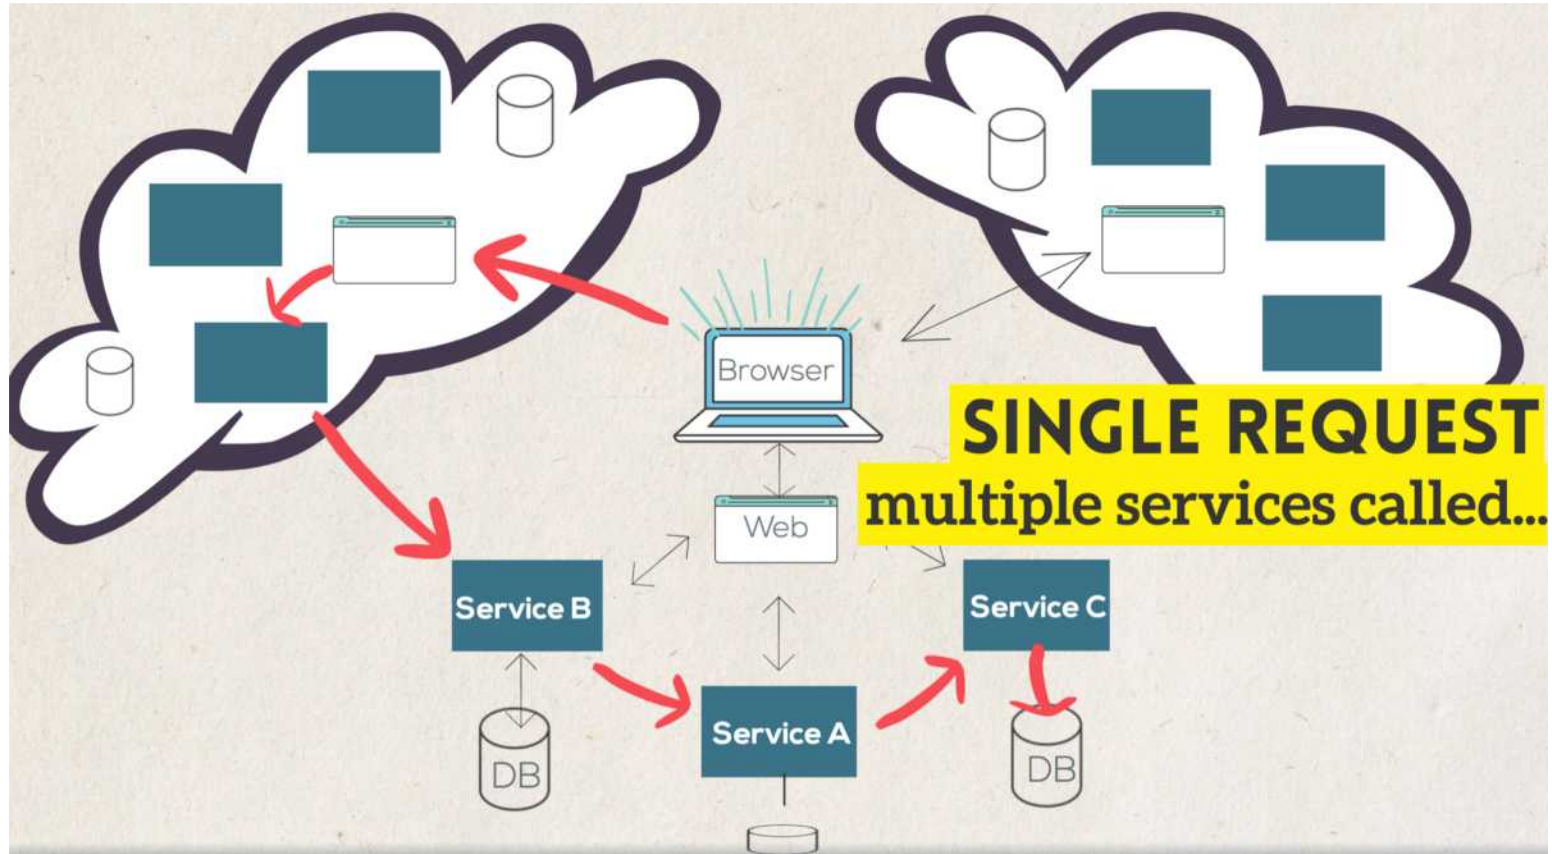
Bringing the world within reach

# The Why




**NOW**  
MICROSERVICE ARCHITECTURE

# The Why continued



# The final Why





# Finding the needle in a Haystack



Source: hikingartist @ FLICKR

# The What



# Haystack

A resilient, scalable tracing and analysis system

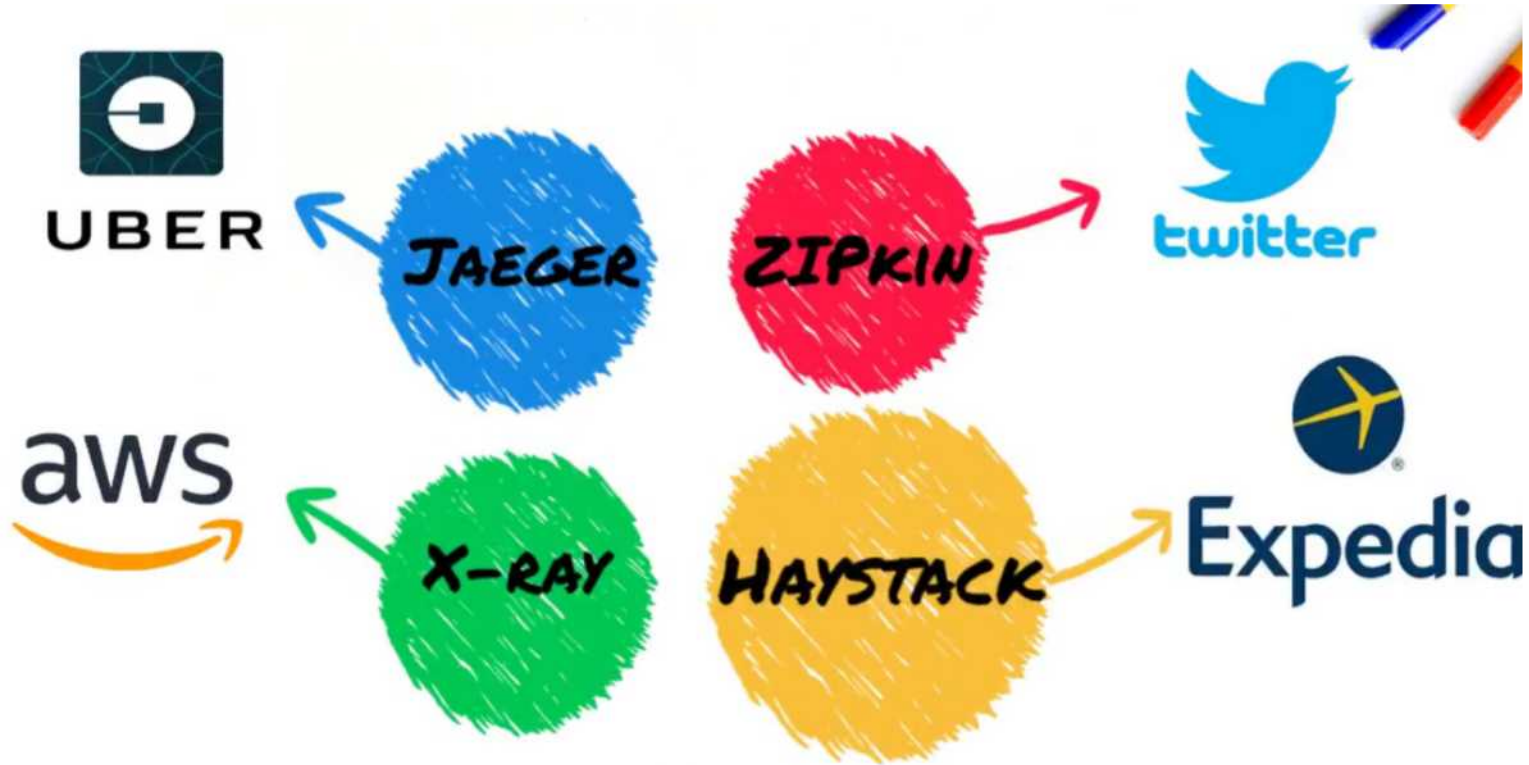


# The How





# The Where and when



# The Why, again: Build our custom solution

## UUID Support

- Existing solutions did not support UUID's as identifiers for Span

## Extensible Platform

- Ability to build more systems to make use of this immensely useful data stream

## Open Tracing

- Open Tracing Compliance so the organization is not coupled to a single implementation

## Open Source

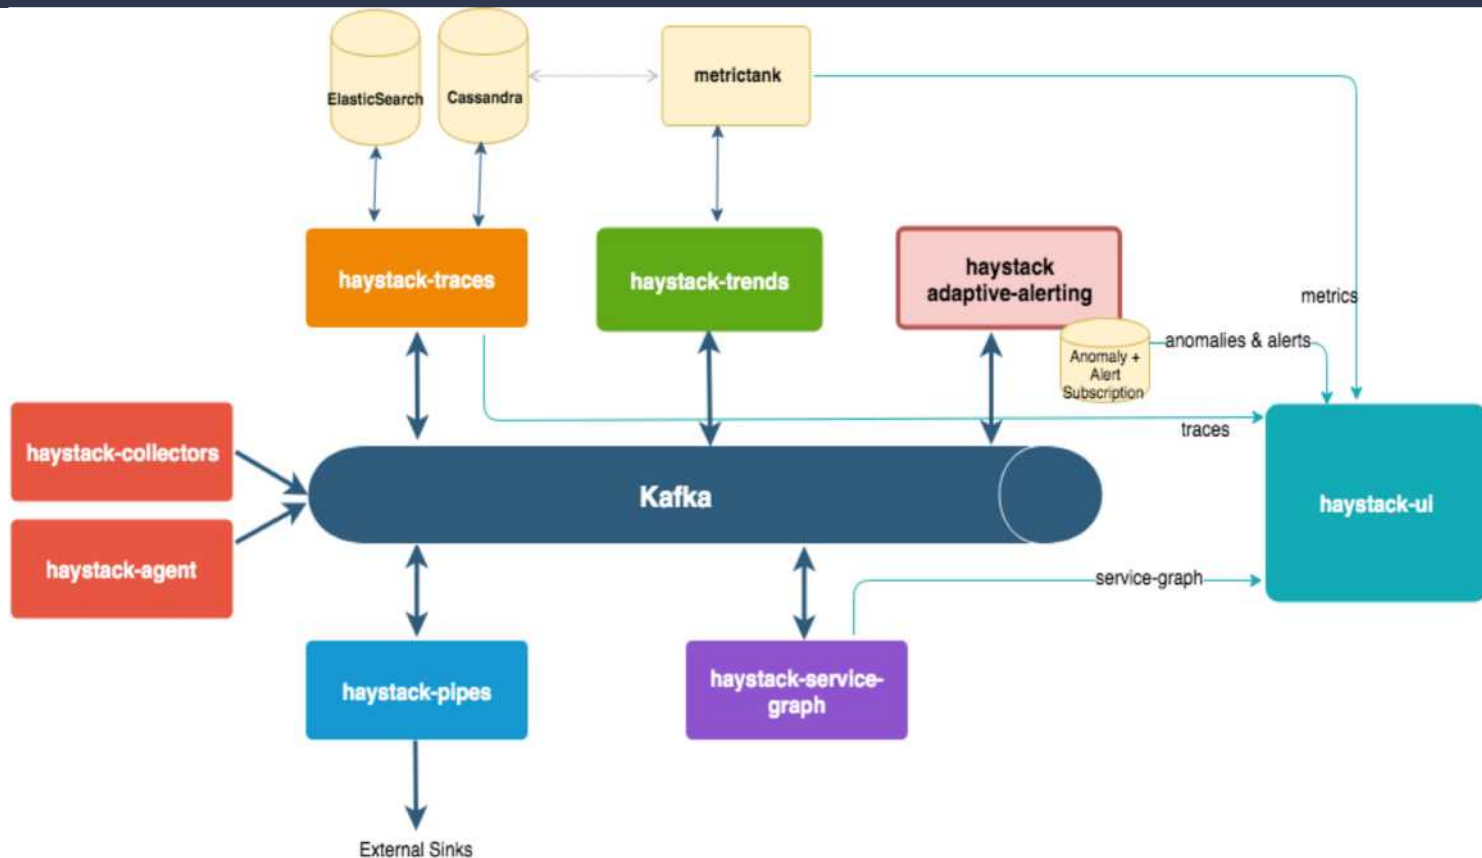
- Opportunity to contribute to the open-source community

# Haystack Architecture

COMPONENTIZED

RESILIENT

SCALABLE



# Haystack Components

Traces

Trends

Service Graph

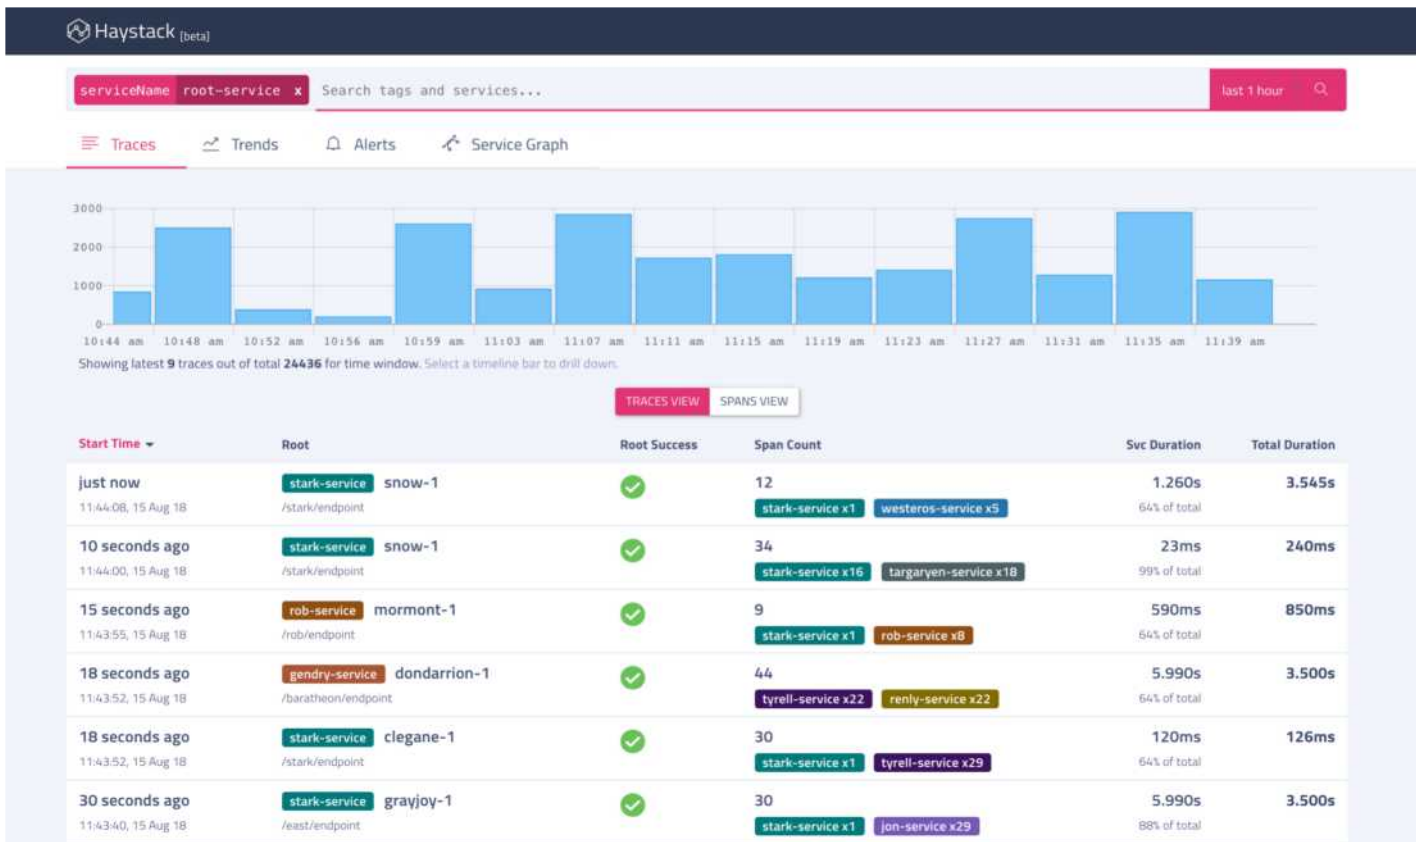
Latency Cost

Alerts

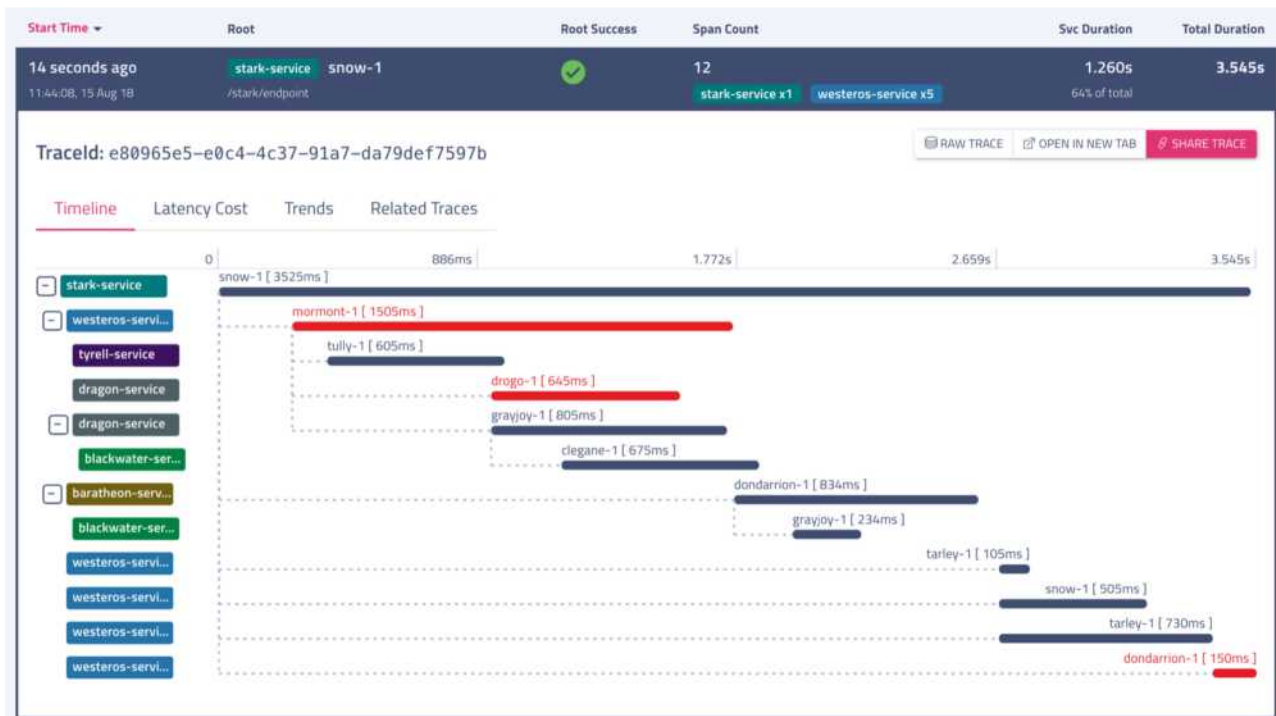
Pipes



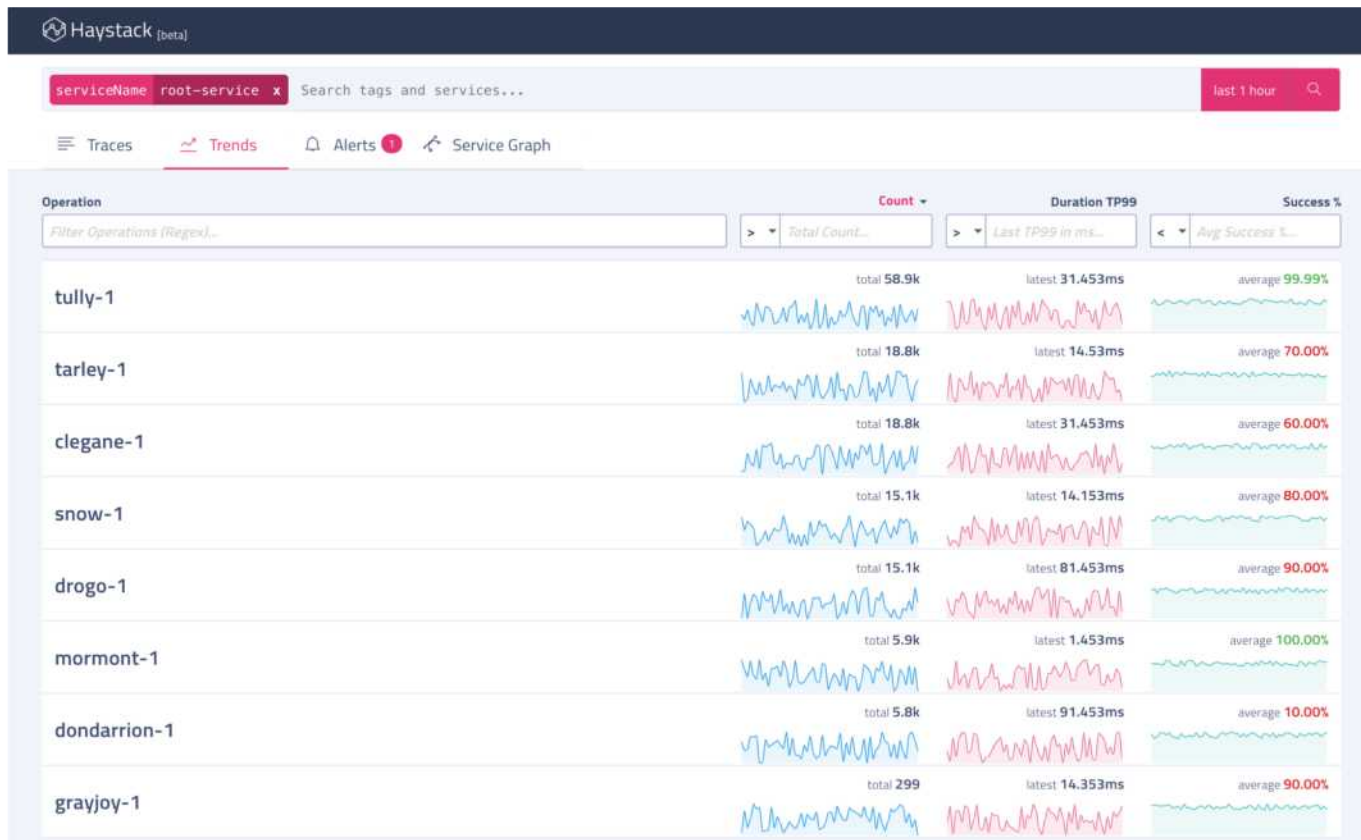
# Traces component



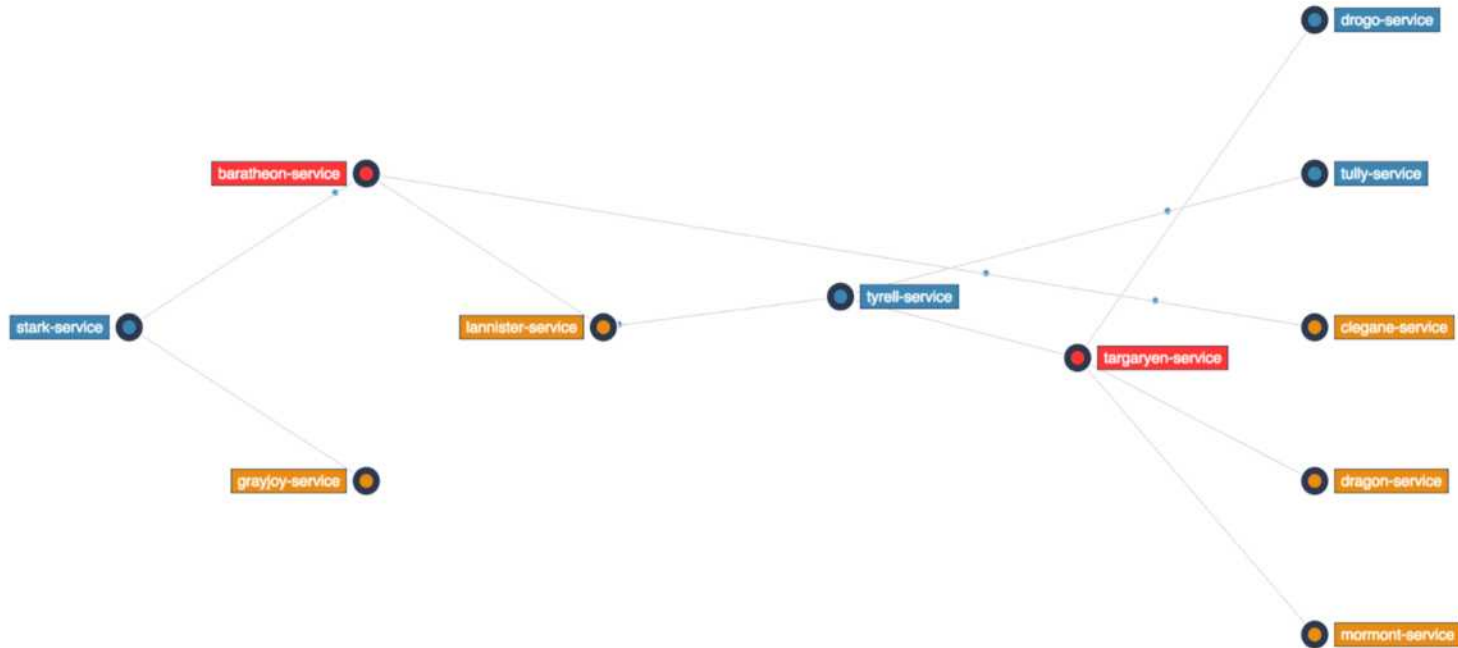
# Traces Waterfall View



# Trends component



# Service Graph component





# Network Latency Cost component

TraceId: e80965e5-e0c4-4c37-91a7-da79def7597b

[RAW TRACE](#)[OPEN IN NEW TAB](#)[SHARE TRACE](#)[Timeline](#)[Latency Cost](#)[Trends](#)[Related Traces](#)[SINGLE TRACE LATENCY](#)[AGGREGATED LATENCY](#)

Network time **996ms** (10 measured out of 11 calls)

Network time cross datacenters **518ms** (5 measured out of 6 calls)

Datcenters involved

6

aws us-east-1

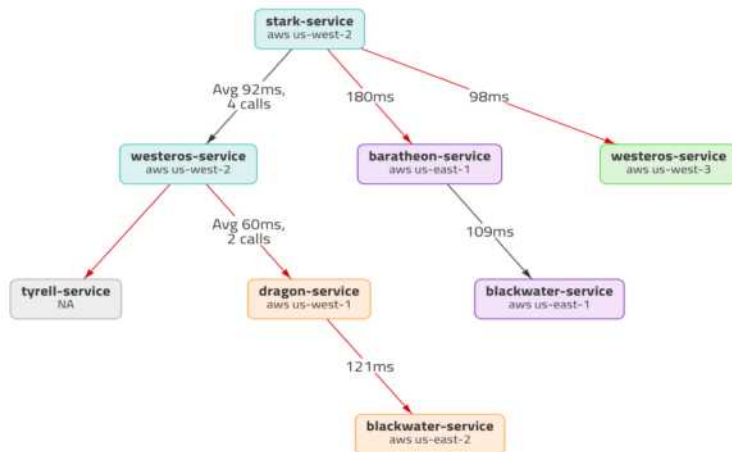
aws us-east-2

aws us-west-1

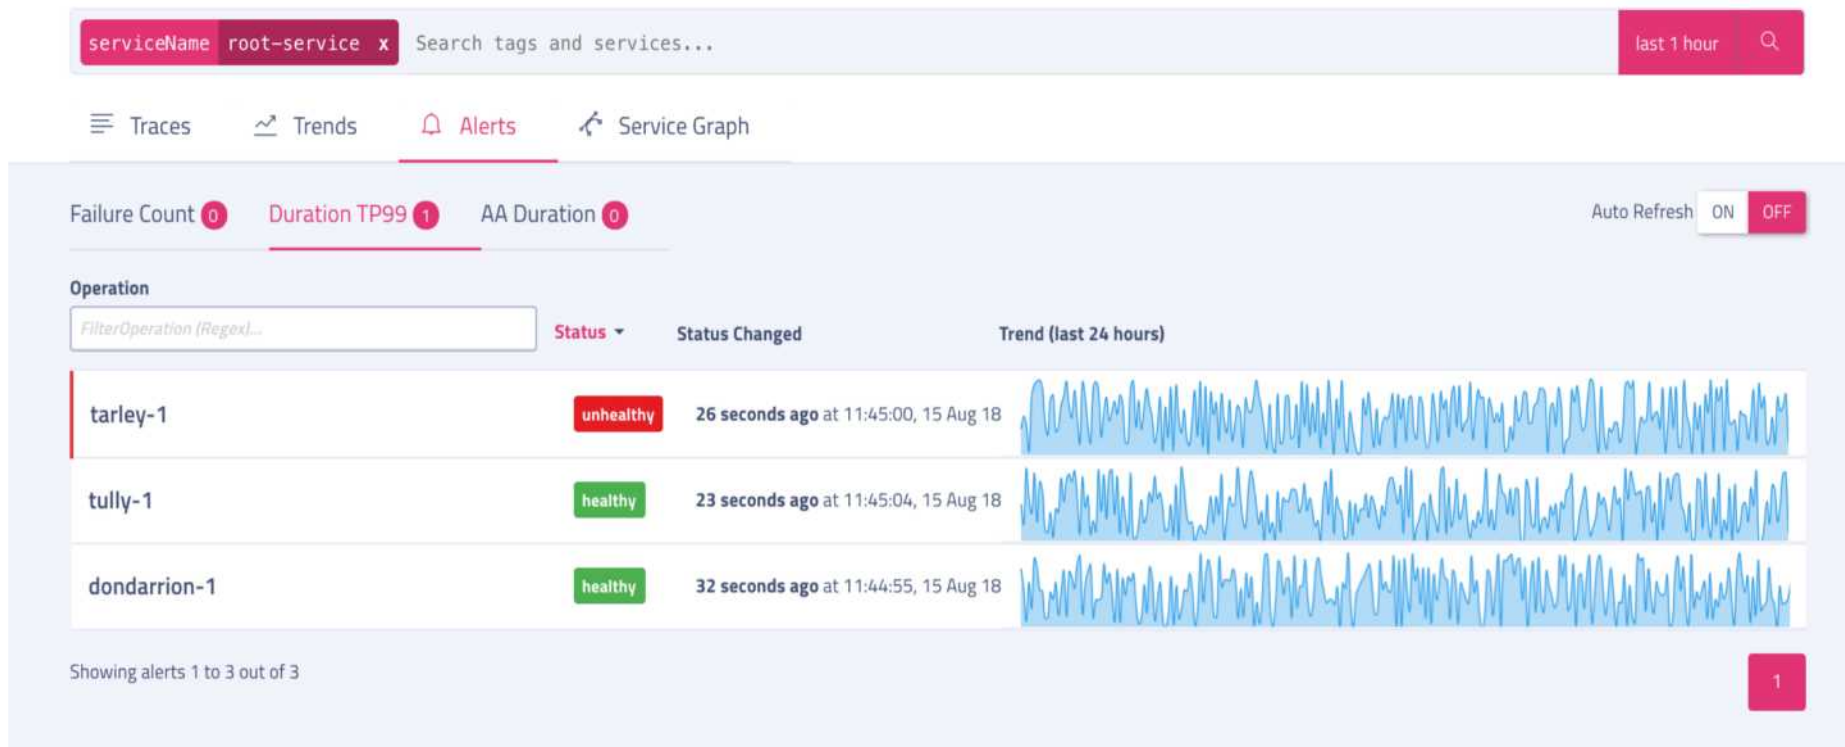
aws us-west-2

aws us-west-3

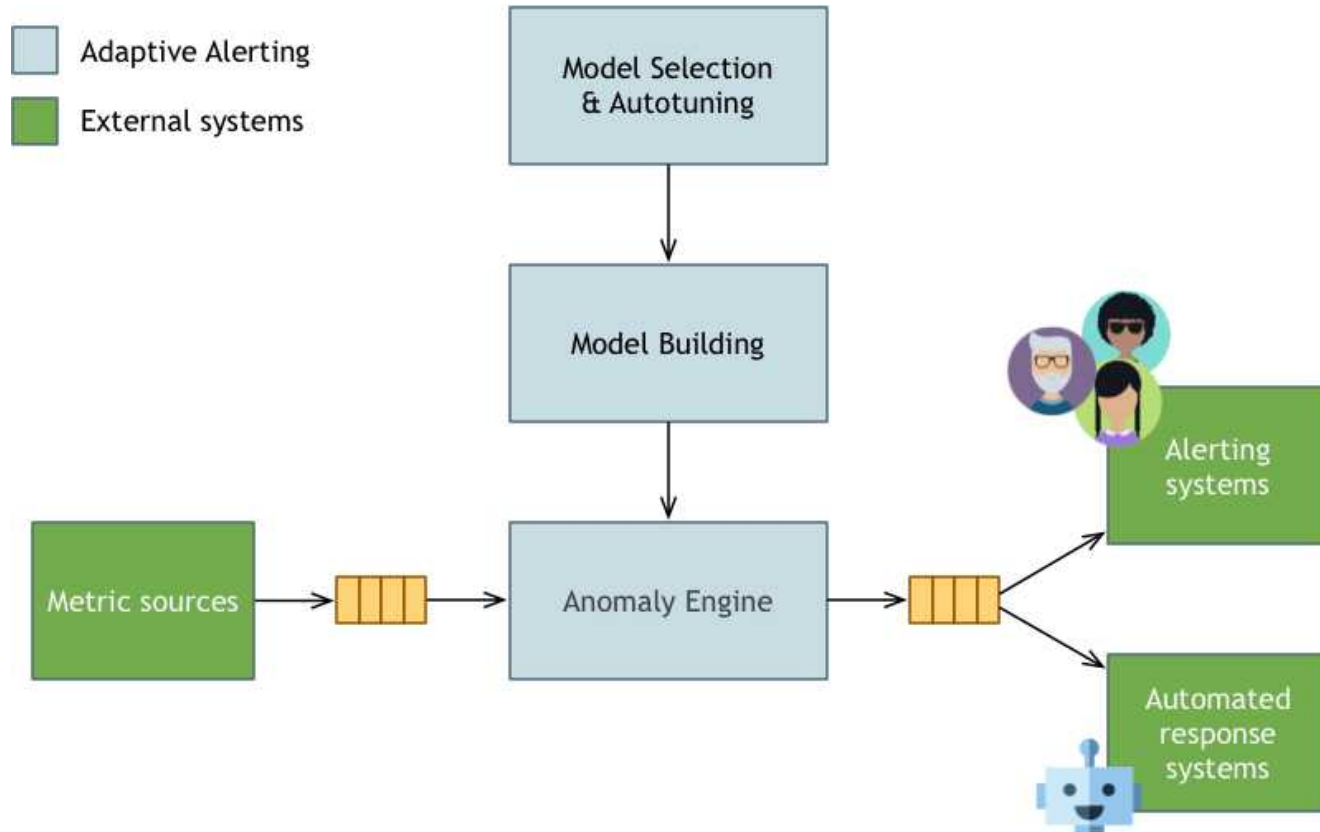
NA



# Alerts component



# Next steps: Adaptive Alerting



# Haystack @ Expedia Group

- **Multiple brands**
- **More than five hundred services**
- **500k/sec spans ingestion**
- **150+ node c5.xlarge k8s cluster, 40 node Kafka cluster, 50 node c5.xlarge Cassandra, 16 ES node cluster**
- **Support OpenTracing clients in Java, NodeJS, Go & Python.**
- **Request/Response logs in open source**



# Key Takeaways

- **Observability is a must for modern day complex systems**
- **You can sample the logged traces according to your needs to keep costs in check**
- **If you do not wish to sample so as to not lose any interesting cases, rethink the data collection and analysis architecture**
- **A heterogeneous system means more work and time on implementing observability**

# Welcoming contributions



# References

<https://github.com/ExpediaDotCom/haystack>

<https://expediadotcom.github.io/haystack/>

[haystack@expedia.com](mailto:haystack@expedia.com)

# Thanks

