
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

PRT027

SPONSORED BY CHAOSSEARCH

Three key considerations for deploying best-of-breed observability

Ed Walsh

CEO
ChaosSearch

aws

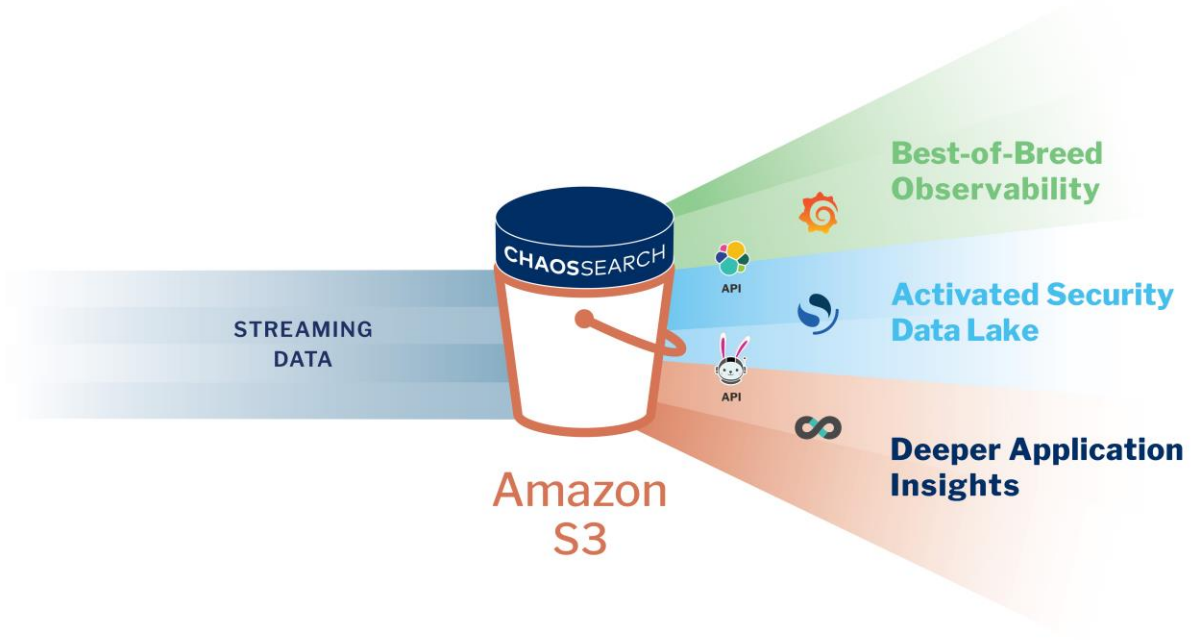
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The Cloud Data Platform for Operational Analytics at Scale

ChaosSearch was founded to solve the core challenges of analytics at scale.

ChaosSearch enables organizations to Know Better® by activating the data lake for log and event analytics and delivering a platform for observability and security data lake use cases.

The End Result: Rapid time to insights paired with simultaneous reductions in time, cost and risk.



FEATURED CUSTOMERS

EQUIFAX Blackboard bai communications REVINATE Agilence ARMOR Klarna. Digital River

Speed and data are essential to compete in today's economy

Softwarization of the economy + **Cloud** + **Faster dev cycles**



Usage growth





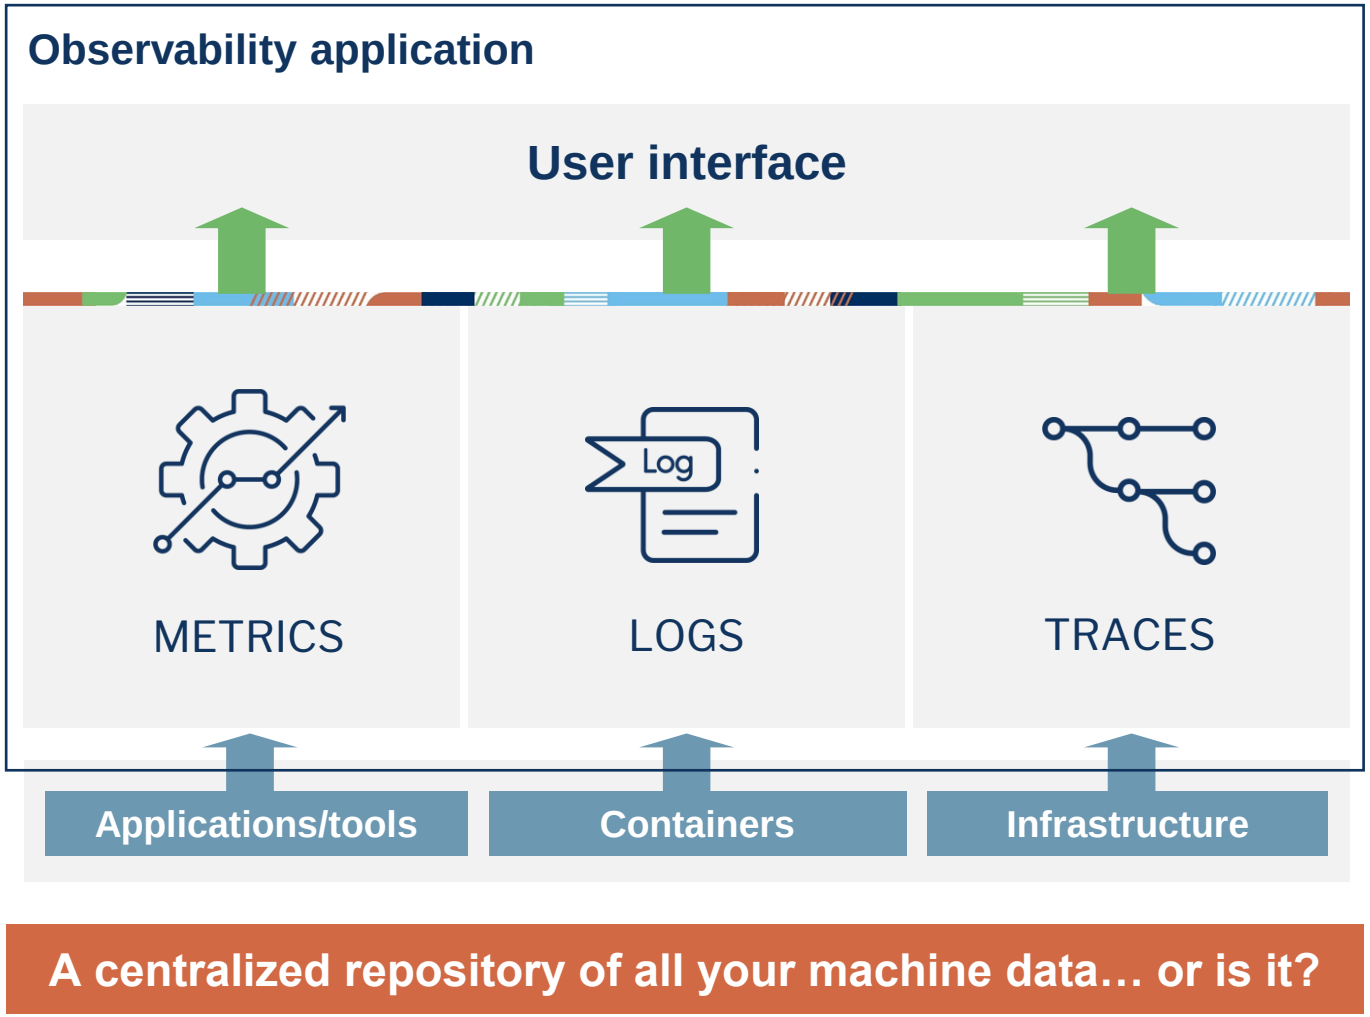
EXPONENTIAL DATA GROWTH

INFRASTRUCTURE FRAGMENTATION

SEMI-STRUCTURED and CHANGING DATA

Development, speed, and scalability are essential, but result in production chaos

You need centralization – The solution: Observability



A centralized repository of all your machine data… or is it?

It does centralize but…

This approach works at smaller scale, but as data volume and infrastructure complexity grow, costs get quickly out of hand.

Why?

- Observability solutions are built so they can monitor and alert quickly.
- But very fast queries (second time-to-glass, millisecond queries) are not a friend of data at scale or cost-effectiveness.

An example: Log analytics

Log analytics costs	Previous workarounds	Today's reality
Typical costs • 80¢ to $3 per GB of raw data • Cost increases with retention limits, and anything beyond 30 days is not typical	• Reduce log volume by prioritizing logs • Reduce retention	• Harder to reduce volume: Increased usage and microservices • Need for longer retention: Lingering issues, regulation, and additional users

Log analytics is a key culprit of runaway observability costs because observability tools were built to monitor, not hunt

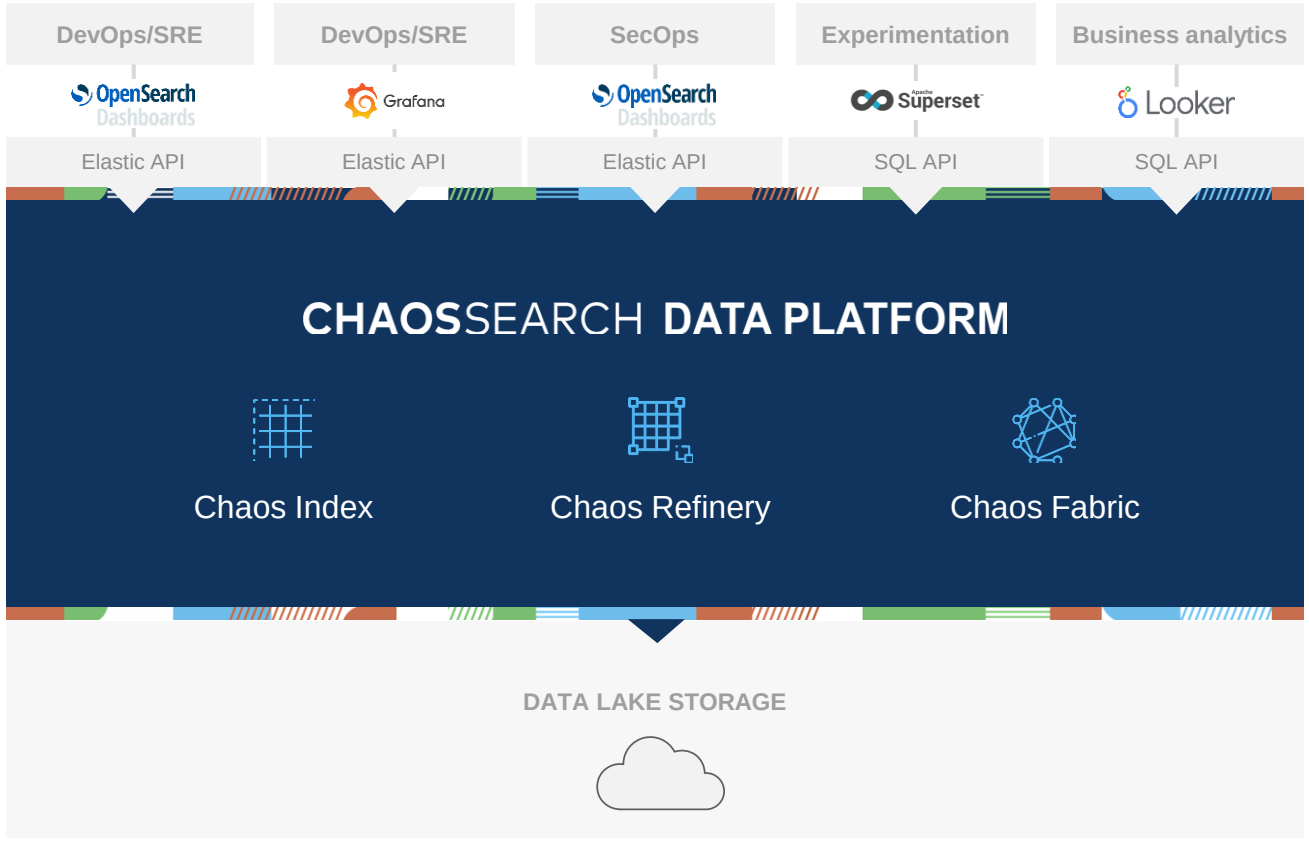
Three key considerations for observability

1. **Retention and scale** – How much retention and scale do I need to fulfill my different use cases?

2. **Toil** – How much toil can my current solution generate? What's the toil impact of reducing coverage or retention?

3. **Tools** – Do I have the right tool(s) for the job given the considerations above? How can I minimize tool sprawl while optimizing cost, scale, and toil?

Search + SQL database for best-of-breed observability



Simply stream your data into your S3 and get instant insights out

- ✓ **Unified data platform for Operational Data at scale**
  - ✓ Simultaneously support multiple data consumers and/or use cases (search, SQL, and ML* API access)
  - ✓ ELK/OpenSearch replacement for log analytics
  - ✓ Security data lake to complement your XDR/SOAR/SIEM
  - ✓ Centralized log analytics without retention limits for modern observability stack
- ✓ **Game-changing simplicity**
  - ✓ Liberate valuable resources from data pipeline creation, constant maintenance, and troubleshooting
  - ✓ All while using the same set of analytic tools
- ✓ **Superior cost economics and scale**
  - ✓ Unlimited data ingest and retention
  - ✓ Painlessly analyze at petabyte scale, while reducing costs by 50–80%

* Note: This is Roadmap and subject to change

Centralize logs for an operational data lake to complement observability

Detection/response

APPDYNAMICS

DATADOG

dynatrace

new relic.

Observability application

Centralize operational data with short retention for fast identification and response to issues without breaking the bank

Troubleshooting/exploration



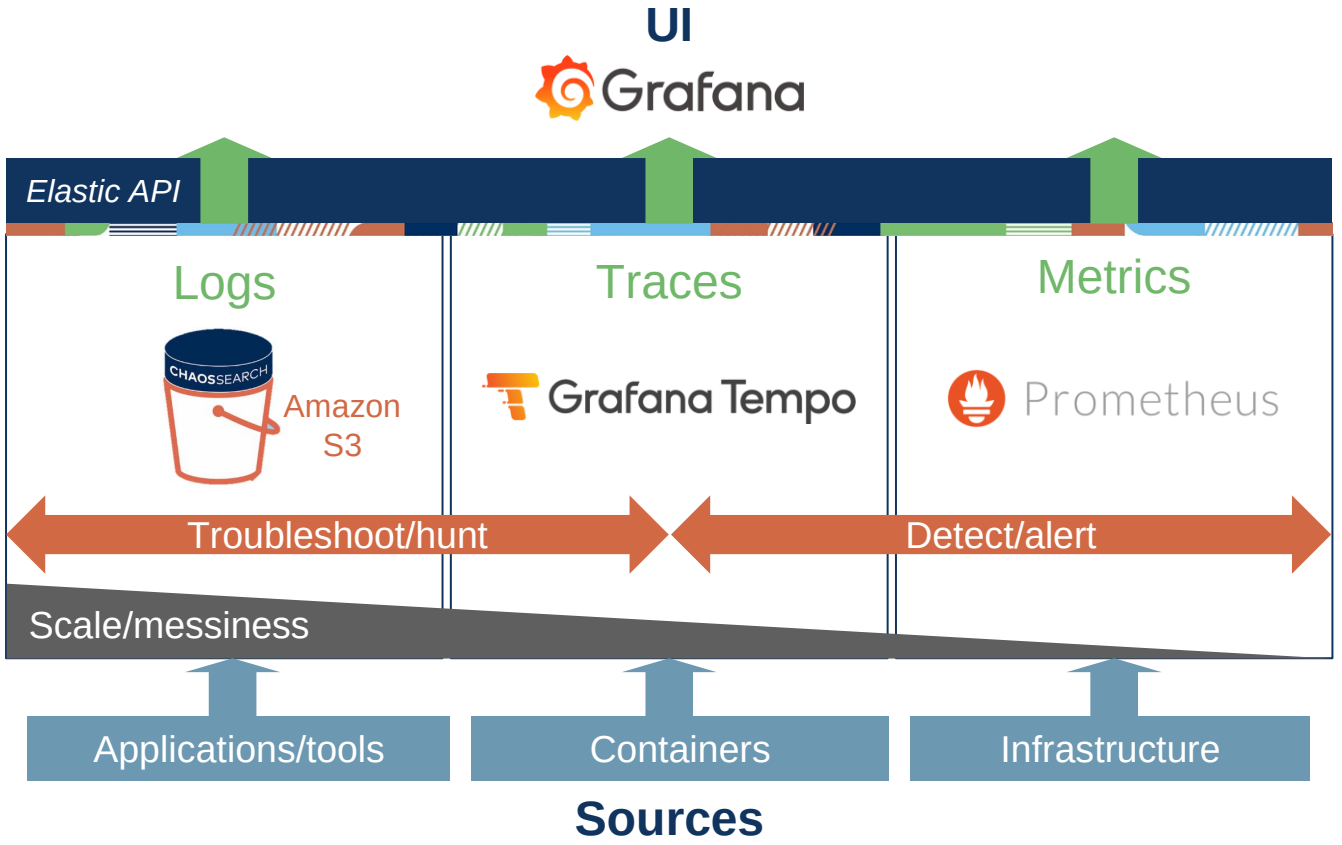
Unified operational data lake

Centralize logs and other operational data with unlimited retention for cost-effective troubleshooting/exploration without rehydration. Flexible and efficient ingest. From 50–80% savings.

Complement your observability app with ChaosSearch to manage today's complex microservices architectures

Modern modular observability stack

The best of both worlds for observability: Single UI, best-of-breed under the hood



Unified user interface and best-of-breed database approach

Superior user experience and analytics capabilities. No data retention limits. Fraction of the cost.
Superior visibility over the ever more complex infrastructures across users.

Best-of-breed observability at Equifax

- 90% cost savings vs. prior solution
- Fast time to value and reduced toil
- Best-of-breed observability with ChaosSearch and Grafana
- SQL access to log data for unified operational analytics platform

Link here: https://www.chaossearch.io/resources/customer-stories/equifax

EQUIFAX®

"With ChaosSearch, we no longer have to move or transform our data. That means we don't have to think about how we're going to use the data before we access it. ***It's just there for us to query****. That helps us deliver new opportunities to the business or discover usage patterns in our systems."*

Jeff Kinsherf
SVP, Engineering Services and SRE

ChaosSearch bridges operational and business analytics

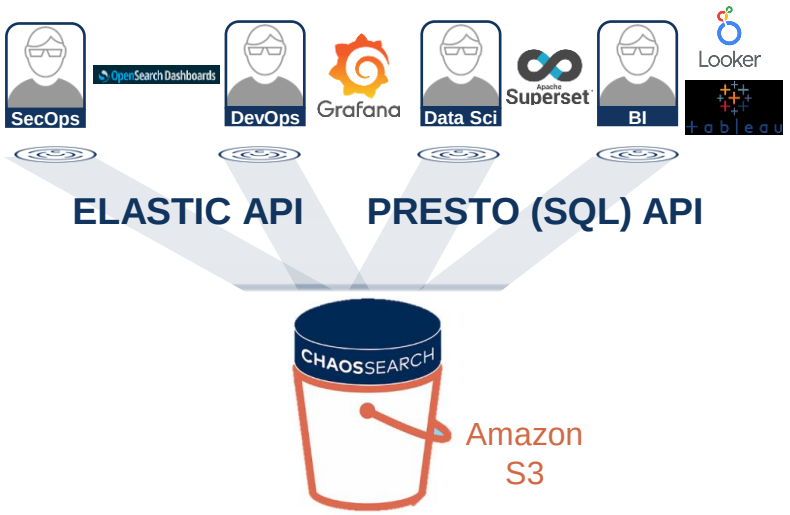
Centralize all your machine-generated data without retention limits to unleash the value of IT data without the cost

SIEM/XDR/SOAR

splunk> · CROWDSTRIKE · CORTEX · DEVO · TANIUM · sumo logic · SentinelOne

Observability

DATADOG · new relic. · Grafana · dynatrace · APPDYNAMICS · CloudWatch



All "raw" machine-generated data

(reference tables and other data in Amazon S3)

Cloud data warehouse

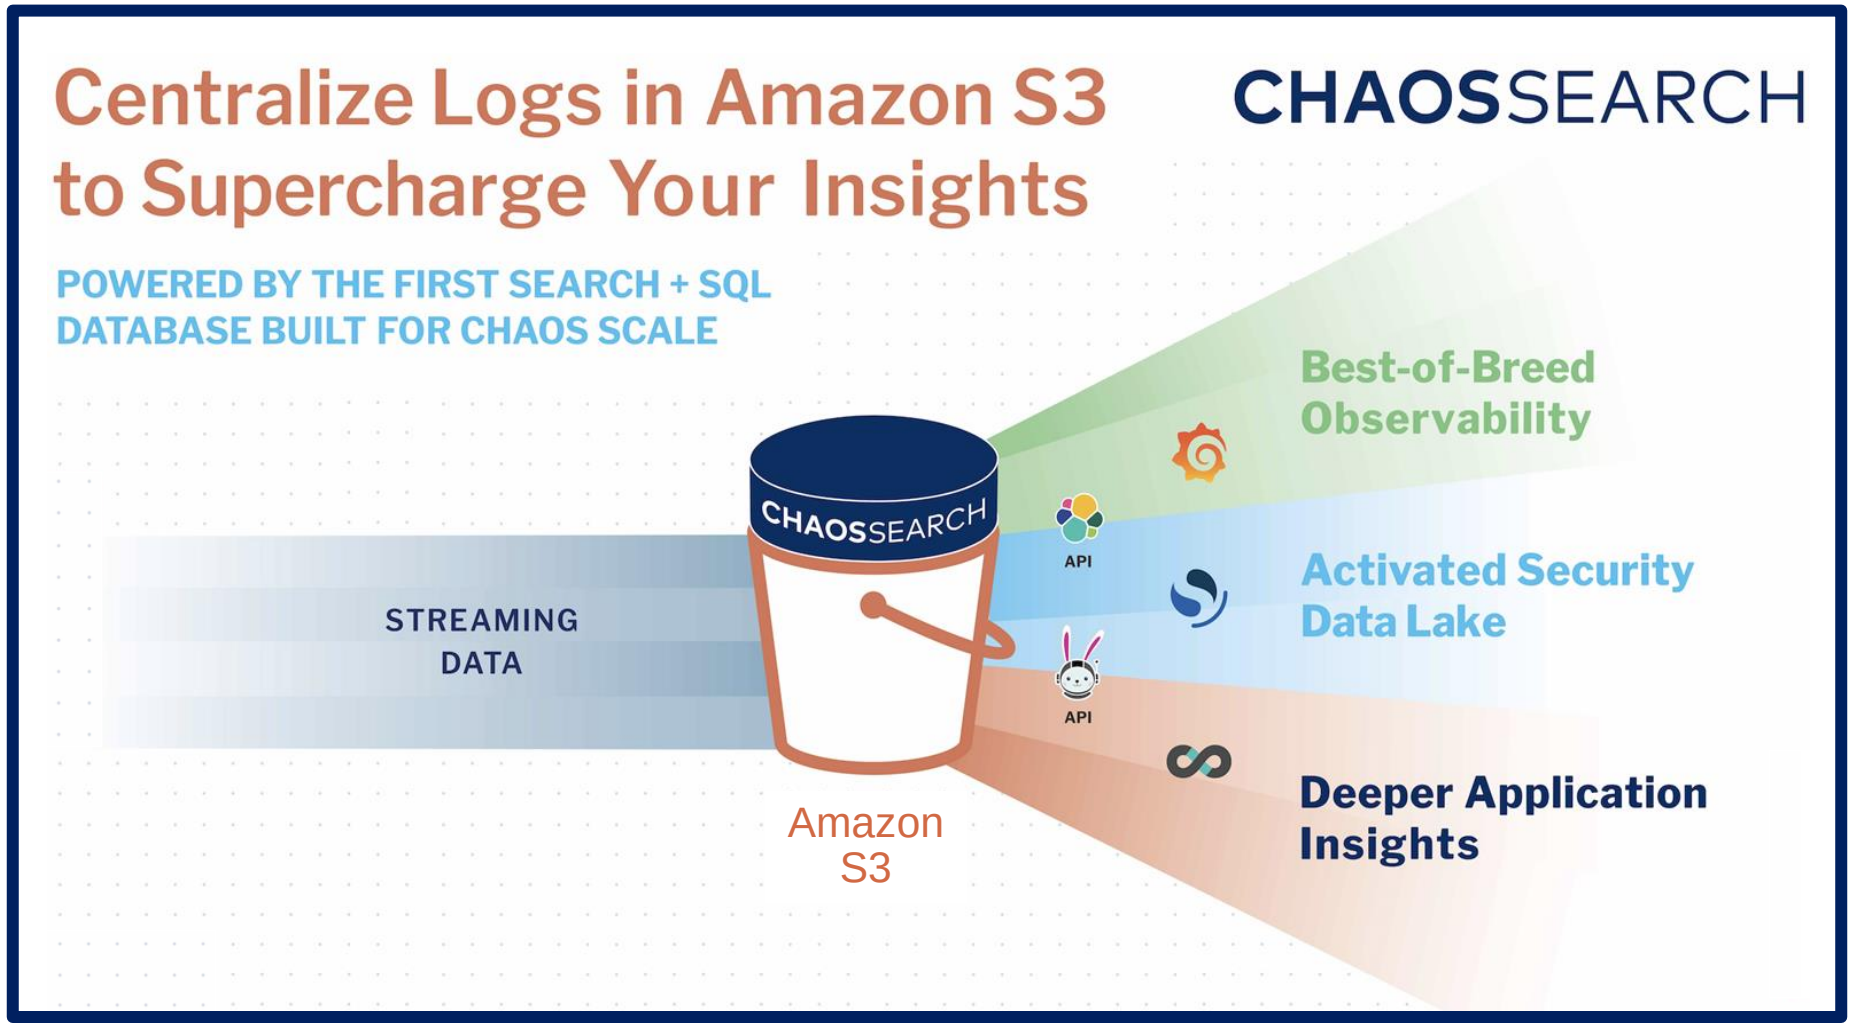
Amazon Redshift · snowflake · SQL

Business apps

fullstory · productboard · HubSpot · Amplitude · Jira · Marketo · Gainsight · salesforce · intuit quickbooks

Complement your apps and existing analytics databases for analytics at scale at a fraction of the cost

Getting started is easy. Stop by Booth 3040 to see how.



Thank you!



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.