Cribl

WHITE PAPER

2022 Observability Trends and Predictions

WHITE PAPER

2022 Observability Trends and Predictions

Introduction

We spend all year analyzing the market, exploring trends, and seeing what's on the horizon. These 2022 observability trends are based on our own analysis and what we're hearing from customers, partners, and prospects.

TREND ONE:

Observability Moves In House

There are dozens, if not hundreds, of products touting complex machine learning models aiming to capitalize on your observability data. Some describe themselves as delivering AIOps to overworked operations and security teams, while others deliver pre-packaged models and APIs for these teams to use and implement on their own.

These products have had mixed successes. Many of these products assume they're the only tool in the stack, making them hard to integrate with other tools, or with adjacent business processes. Another problem is relevance. The models and approaches used by many vendors are generic and may not be applicable to the specific problems operations teams face.

Over the next year, operations teams will shift away from these monolithic, generic automation solutions and towards more home-grown implementations built to solve the most pressing security and operational challenges. Rather than focusing on a single tool and data silo, they'll build these tools with a mix of technologies accessing data from across the enterprise.

THERE ARE DOZENS OF PRODUCTS TOUTING COMPLEX MACHINE LEARNING MODELS AIMING TO CAPITALIZE ON YOUR OBSERVABILITY DATA, WITH MIXED SUCCESS.

TREND TWO:

Security Teams Drive Observability Maturity

Much of the conversation around observability in 2021 targeted developers, with the view that developers are also the operators of their code. While this view is popular in Silicon Valley, the reality is often quite different outside the Bay Area. Developers are expensive, and having them spend time on operational tasks they're not experts in is a waste of time and effort. Instead of developers driving the observability discussion, we'll see cybersecurity teams taking the driver's seat and leading this transformation in their companies.

Security teams are heavy users of monitoring already, deploying a range of tools to uncover known threats. But these tools fall short in three ways. First, they take a one-size-fits-all to the data they ingest. They ignore varying levels of data quality and value, treating all data as the same. Second, current pricing models make broad security monitoring cost prohibitive. Ingest-based pricing penalizes users for every byte ingested, while workload-based pricing penalizes users for every search they run. These limitations hamper investigations and slow remediation. Finally, no tool or platform owns all of the data, resulting in a fragmented data picture. An observability-based security architecture must take into account all of these fragments and weave them into a coherent picture.

When it comes to security use cases, security teams are realizing that all data is in scope, not only the traditional data sources. Security-focused observability shifts the conversation from what teams already know to look for to uncovering new threats and attack patterns on the horizon while also remediating existing breaches. For security professionals, observability offers benefits around:

- *Uncovering current governance and compliance gaps.*
- *Routing data to multiple destinations for advanced analysis across a range of tools, e.g. routing the same data to Splunk ES to drive detection and case management and to Google Chronicle for threat hunting.*
- *Conduct faster and more accurate post-mortems on security events.*
- *Enriching data with additional context, allowing teams to ask better questions and support the adoption of AIOps.*
- *Filtering out low-value data, ensuring existing data is relevant and useful for analysis, thereby improving the quality of questions from SecOps teams.*

As security teams explore how observability practices benefit them, they'll demand more accessibility to observability data and better tools to manage it. These needs will become more acute as teams migrate away from legacy platforms to modern solutions.

SECURITY-FOCUSED OBSERVABILITY ENABLES TEAMS TO UNCOVER NEW THREATS AND ATTACK PATTERNS ON THE HORIZON WHILE ALSO REMEDIATING EXISTING BREACHES.

DATA VOLUMES WILL GROW AT A 23% CAGR THROUGH 2025.

WHEN INGESTING TERABYTES OR PETABYTES OF DATA, THE REQUIRED RETENTION PERIODS, DATA TRANSFER COSTS, AND API CALLS EASILY TAKE THE COST OF CLOUD-BASED OBSERVABILITY DATA CAN RUN INTO THE MILLIONS OF DOLLARS PER YEAR.

TREND THREE:

Adoption of Cloud-Based Monitoring and Observability Decelerate

We saw a massive expansion of cloud-based monitoring and observability offerings in 2021, often spurred by an overall acceleration of public cloud computing in response to pandemic-driven digital transformation efforts. These decisions were often hasty, made by teams trying to keep up with staggering demand for their services. From conversations we've had with our customers and prospects, cloud-based logging offered great promise but frequently came with high costs and unpredictable performance.

Most monitoring and observability workloads are surprisingly predictable, and the data processed grows, more or less, at a linear rate. These two predictable characteristics make the overwhelming majority of these workloads ideal for on-prem processing. There is little need for processing that can dynamically scale up or down. The workloads are constant and consistent.

One counterintuitive aspect of this trend is the growth of data volumes. After all, IDC has stated that data volumes will experience a 23% CAGR through 2025. It might seem that cloud is an ideal place to house all this new observability data, but the data volumes we're talking about make cloud-based object storage inefficient when it comes to cost. Many of the companies we work with are ingesting over 40TB of data each day. Some scale up to 100TB or more. Once you add in the required retention periods, data transfer costs, and API calls, it's easy to see how the cost of cloud-based observability data can run into the millions of dollars per year.

VOLUME OF DATA CREATED AND REPLICATED WORLDWIDE

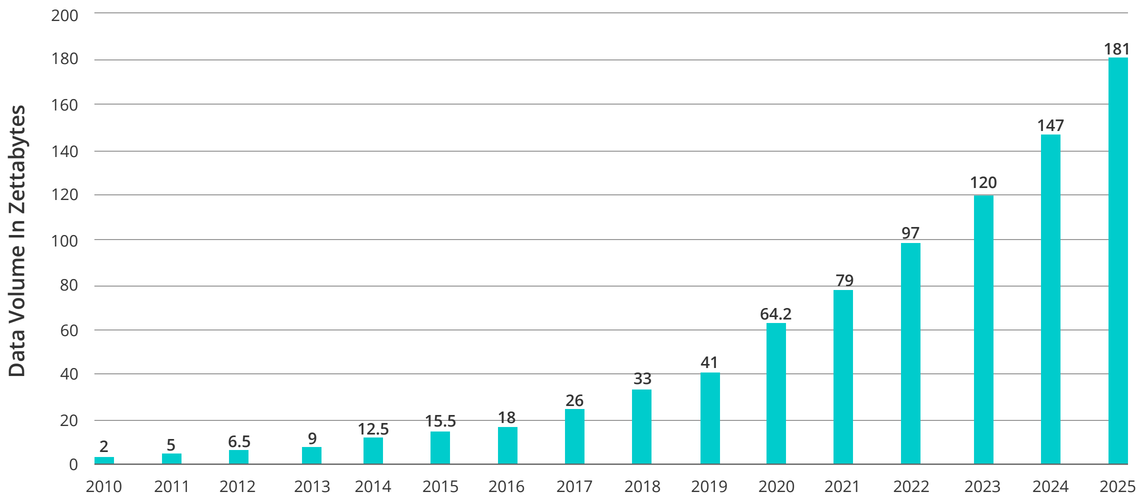

Source: IDC

AN OBSERVABILITY LAKE CENTRALIZES OBSERVABILITY DATA FOR DISPARATE TEAMS AND OFFERS THE FEATURES AND FUNCTIONALITY THEY NEED TO UNDERSTAND THEIR RAPIDLY EVOLVING ENVIRONMENTS. UNLIKE TRADITIONAL DATA LAKES, OBSERVABILITY LAKES OPTIMIZE FOR THE EXPERIENCE OF THE USER.

4

TREND FOUR:

Observability Lakes Become the New System of Discovery

As these costs continue climbing, organizations will experience invoice shock and pump the brakes on their cloud-based monitoring and observability migrations.

Today's monitoring and observability environments are fragmented across dozens of different tools and data silos. These environments haven't kept pace with the realities many overworked infrastructure, operations, and security teams deal with, like ephemeral infrastructure, highly distributed applications, and complex application architectures. Additionally, there is an operational imperative for pervasive instrumentation in modern applications. Fully instrumented applications can generate terabytes of data per day, creating acute challenges around observability data management.

An observability lake centralizes observability data for disparate teams and offers the features and functionality they need to understand their rapidly evolving environments. Unlike traditional data lakes, with their focus on data processing and analytics ranging from self-servicing canned reports to machine learning use cases, observability lakes optimize for the experience of the user. SREs, SecOps, and DevOps people aren't SQL users. They understand search and regular expressions. Observability lakes must offer robust search capabilities coupled with powerful time-based features like binning, aggregations, and ranges. Those features must work across a range of different data formats and types.

This is not to say observability lakes will be a massive, centralized element in your observability infrastructure. That type of design is too limiting, and prevents users from accessing any observability data in their environment. Instead, the observability lake looks more like a data virtualization layer, reaching into a range of data stores regardless of their physical location, including on-prem, edge, and cloud data stores.

ABOUT CRIBL

Cribl is a company built to solve customer data challenges and enable customer choice. Our solutions deliver innovative and customizable controls to route observability data where it has the most value. We call this an observability pipeline, and it helps slash costs, improve performance, and get the right data, to the right destinations, in the right formats, at the right time. Join the dozens of early adopters, including market leaders such as TransUnion and Autodesk, to take control and shape your data. Founded in 2017, Cribl is headquartered in San Francisco, CA. For more information, visit www.cribl.io or our LinkedIn, Twitter, or Slack community.