Cribl

WHITE PAPER

# 2021 Observability Predictions

WHITE PAPER

# 2021 Observability Predictions

## Introduction

The pandemic accelerated digital transformation in many companies, driving investment in infrastructure, developers, and updated tooling to deploy applications faster across diverse environments. This resulted in a massive increase in data volumes flowing into logging analytics systems, like security information and event management (SIEM) platforms and application performance management (APM) tools.

Accelerated digitalization hasn't been orderly or well planned in most enterprises. Companies will spend 2021 reinforcing and reestablishing their governance and security programs, getting a handle on costs and customer expectations.

On the following pages are our predictions on the impact of that restructuring.

PREDICTION ONE:

**Three-quarters of container-based deployments will exceed infrastructure budgets by 200% due to container blindness.**

**Businesses are coping with uncertainty by rearchitecting for real-time adaptability and resilience.** Instead of crafting monolithic applications on depreciating legacy hardware, they're building modular applications on flexible, cloud-based infrastructure. These architectures and deployment possibilities allow enterprises to compose new solutions to capture rapidly changing market conditions and application loads.

Containers are the core technology enabling this agility. Containers provide consistent, portable runtime environments for applications. According to a survey from the Cloud Native Computing Foundation (CNCF), production use of containers increased from 23% in 2016 to 84% in 2019. Nearly 20% of those survey respondents report deploying over 5000 containers in production!

The challenge for infrastructure and operations (I&O) leaders is observing what all these containers are doing. With over 100 container management tools spread across cloud and on-premises environments, aggregating logs, metrics and traces to understand infrastructure costs is, at best, daunting. Without an observability pipeline, it's impossible to know how your containers are performing.

That same observability pipeline is *also* essential for cost management.

According to a survey from Datadog, **49% of containers use less than 30% of requested CPU**. Similar problems occur for memory, with **45% of containers using less than 30% of requested memory**. Of course, you're paying for 100% of those resources, regardless of what you use. The amount of money I&O leaders waste on over-provisioned container deployments is staggering. It's no surprise many I&O leaders simply give up rather than try to make sense of their container infrastructure environment.

Some I&O teams lean on existing application performance monitoring (APM) systems for insight into their container deployments. Traditional APM has its own challenges with expensive per-application or per-container pricing, driving up costs even further. These price pressures force many enterprises to only install APM on a fraction of applications, making it unworkable in today's dynamic application environment.

All of these factors, from environment complexity to exorbitant monitoring costs, contribute to container blindness.

PRODUCTION USE OF CONTAINERS INCREASED FROM 23% IN 2016 TO 84% IN 2019.

49% OF CONTAINERS USE LESS THAN 30% OF REQUESTED CPU.

45% OF CONTAINERS USING LESS THAN 30% OF REQUESTED MEMORY.

**EVIDENCE:**

**CNCF Survey:** Deployments are getting larger as cloud native adoption becomes mainstream

**Datadog:** 11 facts about real-world container use

PREDICTION TWO:

**Log anxiety results in 40% of security-related breaches going undetected as teams disable log collection to stay within budgets.**

**According to a survey from Stanford, 42% of U.S. workers are now working from home full-time due to pandemic-related restrictions.** That staggering number represents more than two-thirds of economic activity. This shift has been so successful that, according to Gartner, 74% of CFOs plan to make the shift to remote work permanent.

The rapid transition to working from home has created new challenges for cybersecurity and SecOps teams. As the edge moves out of the traditional walls of the corporate data center or cloud and into households, cybersecurity risks are on the rise.

According to a new report from the Center for Strategic and International Studies and McAfee, losses from cybercrime in 2020 are projected to be just under a billion dollars. Foreign criminal enterprises targeted industries while the coronavirus pandemic spurred online scams targeting workers. Those workers, pushed into work from home scenarios, became their own I.T. and security support as corporate policies failed to adapt.

Addressing this challenge requires multiple responses. Consulting firm McKinsey recommends companies expand web-facing threat intelligence and security information and event management (SIEM) programs to compensate for expanded work from home risks. That's good advice, but the challenges these staffers face are often more practical.

Collecting more log data from firewalls, VPNs and multi-factor authorization (MFA) is essential for effective threat intelligence, but that flood of data may burst budgets for SIEM and logging analytics platforms. Many of those products charge based on daily ingest rates, resulting in skyrocketing costs. Worse, much of the data stored is routine traffic, meaning companies are paying a premium for worthless data.

Many CFOs, facing 10x increases in logging analytics budgets, will likely accept higher cybersecurity risks to stay within budget. That trade-off can be tough to make. According to IBM, the average cost of a data breach for U.S. companies in 2020 is $8.64M. That's a staggering sum. CFOs may believe the increased risk is worth disabling increased log collection, imaging that a data breach won't occur. However, that same IBM survey reports that the average chance a company will experience a data breach at 27.7%.

42% OF U.S. WORKERS ARE NOW WORKING FROM HOME FULL-TIME DUE TO PANDEMIC-RELATED RESTRICTIONS.

74% OF CFOS PLAN TO MAKE THE SHIFT TO REMOTE WORK PERMANENT.

**EVIDENCE:**

**Stanford Institute for Economic Policy Research:** How Working from Home Works Out

**Gartner Newsroom:** Gartner CFO Survey Reveals 74% Intend to Shift Some Employees to Remote Work Permanently

**McKinsey & Company:** Cybersecurity Tactics for the Coronavirus Pandemic

**IBM:** Cost of a Data Breach Report 2020

PREDICTION THREE:

**Enterprises implementing an end-to-end observability pipeline will lower infrastructure costs by 30% and resolve issues four times faster than competitors, improving customer satisfaction and increasing spend by 15%.**

Modern applications are composed of hundreds or thousands of services that are developed and tested independently by teams that may not communicate with each other. Software deployments may be automated, occurring several times a day across production environments.

Increasing the complexity, each service may also have its own database and data model, which is also independently managed. Add in short-lived containers and dynamic scaling, and it's easy to understand why the only time companies can test their applications is when they're deployed in front of customers. Customers have become unwitting acceptance testers.

Traditionally, infrastructure and operations teams would deploy monitoring for visibility into their environments. The challenge is monitoring hasn't kept pace with modern complexity for three reasons:

- *Exorbitant costs forces teams to compromise on what they're monitoring. Forced into decisions about which logs, metrics, and traces to keep to stay within budget, teams simply can't store everything they need to observe their environment.*

- *As we've pointed out above, pre-built dashboards and alerts don't reflect today's infrastructure reality. Systems scale dynamically, and DevOps teams may deploy code across thousands of containers dozens of times each day. The static views offered by traditional monitoring systems don't reflect this reality.*

- *Monitoring is a point solution, targeting a single application or service. A failure in one service cascades to others, and unraveling those errors is well beyond the scope of monitoring applications.*

ITOps and SecOps teams must evolve past monitoring into observability. Observability is the characteristic of software and systems allowing them to be "seen," and to answer questions about their behavior. Unlike monitoring, which relies on static views of static resources, observable systems invite investigation by unlocking data from siloed log analytics applications.

Implementing observability requires a way to collect and integrate data from complex systems, which is where the observability pipeline comes in. An observability pipeline decouples the sources of data from their destinations. This decoupling allows teams to enrich, redact, reduce and route data to the right place for the right audience. The observability pipeline gets you past what data to send and lets you focus on what you want to do with it.

An observability pipeline makes debugging faster by allowing you to ask "what if" questions of the environment, rather than the pre-calculated views prevalent in monitoring solutions. Faster debugging and root cause analysis means fewer customers experiencing errors in production, which drives up sales.

FORCED INTO DECISIONS ABOUT WHICH LOGS, METRICS, AND TRACES TO KEEP TO STAY WITHIN BUDGET, TEAMS SIMPLY CAN'T STORE EVERYTHING THEY NEED TO OBSERVE THEIR ENVIRONMENT.

Another benefit of an observability pipeline is rationalizing infrastructure costs. Often, the team deploying infrastructure isn't the team paying for it, resulting in over-provisioned infrastructure. Collecting performance data, even for transient infrastructure like containers, gives ITOps teams visibility into how many resources are actually being consumed and where optimizations are possible.

**EVIDENCE:**

**The Software House:** State of Microservices 2020

## Conclusion

The rush to digitalize business operations and customer experiences has resulted in enterprises absorbing complexity across their technology landscape; 2021 is the year IT leaders will reflect on those changes and work to correct them. From cost overruns related to container deployments and logging analytics, to understanding dynamic infrastructure landscapes, companies are forced to grapple with decisions made quickly to cope with the pandemic.