

Observability goes beyond reactive monitoring strategies to enable robust correlation, troubleshooting, and response to performance challenges across complex cloud environments.

Enterprise Cloud-Native Applications and Infrastructure Demand Modern Observability

October 2020

Written by: Mary Johnston Turner, Research Vice President, Cloud Management

Introduction

Enterprises around the world are adopting cloud-native application architectures built on containers, Kubernetes, microservices, and connected cloud infrastructure that encompasses on-premises private clouds and datacenters, edge computing, and multiple public clouds.

Operational complexity is rising as automated software development pipelines for continuous integration/continuous delivery (CI/CD) and agile development processes enable frequent application changes and updates. IDC's research shows that a typical cloud-native application can have from 5 to 15 dependencies on other services, resources, or application programming interfaces (APIs) — each of which generates its own set of operational risks.

Keeping applications up and running requires tracking and correlating a wide range of infrastructure and application performance metrics, traces, and logs in real time. Kubernetes and containers provide more flexible, ephemeral infrastructure to support the demands of cloud-native applications but create additional management complexity because traditional approaches to change control and governance often cannot keep up with continuous fluctuations in resource usage and availability.

Traditional monitoring tools and processes are too siloed and reactive to effectively identify and remediate root causes rapidly across these complex architectures. Traditional tools rely on IT and DevOps teams to visually correlate dependencies and manually initiate workflows to resolve conflicts. As the volume and diversity of monitoring data, formats, and sources proliferate, these operations teams need approaches driven by artificial intelligence and machine learning (AI/ML) to process data streams much faster, eliminate the noise, and focus on the most important signals as quickly as possible.

This IDC Technology Spotlight describes how modern observability solutions, designed for cloud-native applications and connected hybrid/multicloud infrastructure, are changing the way that IT and DevOps teams analyze, maintain, and troubleshoot end-to-end performance challenges.

AT A GLANCE


KEY TAKEAWAYS

- » Observability provides a common framework for collecting, normalizing, and analyzing a wide range of performance, security, cost, and capacity data.
- » Collaboration across IT and DevOps teams can speed root cause assessment and remediation using a comprehensive observability strategy to discover and correlate metrics, traces, and related data.
- » Mature observability users eliminate silos that slow responses and take full advantage of intelligent, automated operations.

Observability Overview and Benefits

IDC's research shows that 97% of global enterprises are using connected cloud strategies that depend on a diverse mix of on-premises, off-premises, hosted, edge, and public cloud infrastructure (see Figure 1). At the same time, over 80% are adopting some degree of agile DevOps methodologies, and many expect that within two years, close to 30% of production workloads will run in containers. Virtual machines (VMs) and bare metal platforms as well as public cloud infrastructure as a service (IaaS), software as a service (SaaS), and serverless solutions are also expected to be part of the mix.

FIGURE 1: **Enterprises Depend on Connected Cloud Strategies**



Source: IDC's Cloud Pulse, 1Q20

Most enterprises expect this complex blend of infrastructure and application architectures will remain in place for a number of years. Different generations of applications and infrastructure will need to be interconnected, and workloads will need to scale and migrate across clouds and different application environments for some time.

From a monitoring perspective, API-driven integrations, cross-cloud portability, rising use of SaaS and third-party application services, and rapid innovation using automated development approaches all generate ever-increasing volumes of monitoring data, including traditional logs, more modern time series data metrics, application traces, end-user experience insights, and various ad hoc notifications.

Finding the root cause of a problem or isolating an error before it creates problems can be a nearly insurmountable challenge using traditional search tools, dashboards, log-based alerts, and graphical analysis. Traditional IT monitoring tools were optimized to collect data about the state and health of individual systems and applications by collecting data from agents, synthetic traces, or network traffic flows. These logs are difficult to integrate into DevOps toolchains and require extensive ad hoc queries and analysis to drive action.

Observability solutions provide a common framework for collecting, normalizing, and analyzing this wide range of performance, security, cost, and capacity data in context to drive rapid, actionable insight across a wide variety of applications, cloud services, and dedicated infrastructure. Observability solutions allow DevOps teams to gain views of the full status and performance of the end-to-end systems across applications, on-premises infrastructure, cloud services, containers, Kubernetes, and emerging serverless functions.

Transitioning from simple, siloed monitoring strategies to more integrated observability solutions can deliver significant benefits to ITOps, DevOps, and line-of-business (LOB) teams. Benefits include:

- » Overall end-user experience improvement
- » Higher levels of application health and availability
- » Faster mean time to repair and overall reduction in operational incidents
- » More consistent configuration and access control security
- » More reliable change control and compliance reporting and automation
- » Better cost control
- » Improved development and business agility

Observability solutions can break down silos of monitoring, tracing, and metrics data and can allow teams to consolidate tools, streamline workflows, and simplify training and processes. They improve the ability to normalize, correlate, and act on data and are vital to the effective use of event-driven automation for intelligent CloudOps.

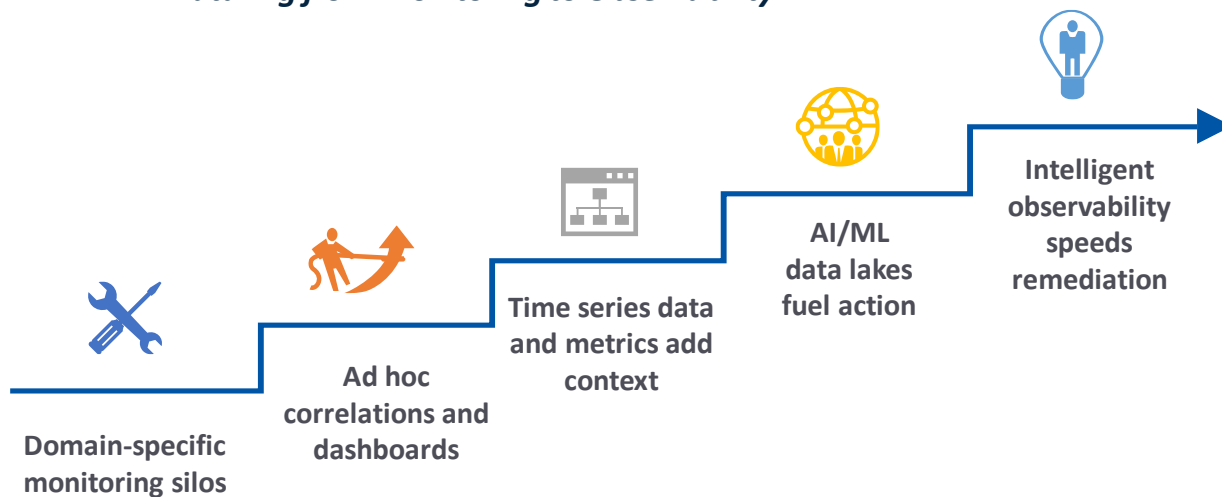
Enterprise Strategies for Evolving from Monitoring to Observability

IDC forecasts that enterprises will continue to accelerate the adoption of cloud-native application architectures and connected hybrid and multicloud infrastructure. VM-based workloads will be lifted and shifted into containers to improve security and portability and to take advantage of Kubernetes-based orchestration and scaling. Edge locations will make extensive use of containers and cloud-native workloads as well.

At the same time, most established organizations will continue to rely on mission-critical VM and bare metal workloads running on traditional infrastructure platforms and datacenters. In many cases, systems of record and systems of engagement will need to share data and unified workflows via APIs and scheduled batch processes. Regardless of the physical resource supporting the workload or the extent to which the application relies on microservices and containers versus more traditional application architectures, the enterprise will need to ensure the end-to-end health of applications and services that run across this connected environment. It will also need to maintain service levels even as individual applications and platforms are updated or replaced.

To optimize service-level agreements (SLAs), costs, and security across these complex environments, enterprises will need to evolve and mature the ways they consume and act on the full range of monitoring and performance data generated across their infrastructure and application footprints. IDC expects that most organizations will need to move away from primary reliance on system- and service-specific log tracking silos. Many organizations have built performance monitoring dashboards that provide some amount of visualization and correlation but frequently offer little in the way of real-time alerting and predictive analysis across systems. Evolving to a more mature operational environment depends on broadening the type of data consumed, particularly time series metrics such as those generated by Kubernetes systems and services. It also requires faster analysis and more event-driven autonomous operations. As shown in Figure 2, the evolution from monitoring silos to integrated intelligent observability provides faster analytics and greater levels of business context to speed remediation and maintain the end-user experience.

FIGURE 2: **Maturing from Monitoring to Observability**



Source: IDC, 2020

Fundamental to this shift from reliance on silos to integrated intelligent observability is the implementation of a data lake strategy that can rapidly ingest, normalize, analyze, and share insights from rapid AI/ML assessments and queries. To achieve this type of environment, developers and infrastructure operations teams will need to rely on API integrations and or agent-based data collectors across development and operations environments. The most mature environments will significantly improve operational efficiency and service levels by linking observability assessments to automated actions and remediations.

The evolution from fragmented monitoring to intelligent observability will often result in a consolidation of tools and workflows across ITOps and development teams. For many organizations, the introduction of cloud-native applications and container-based infrastructure creates a tipping point where the need for modern observability solutions becomes inescapable as a result of the rapidly changing, ephemeral nature of these cloud-native environments.

Considering VMware Tanzu Observability by Wavefront

VMware Tanzu Observability by Wavefront is a SaaS-enabled solution that provides unified enterprise observability for ITOps teams, site reliability engineering (SRE) teams, and DevOps practitioners. VMware Tanzu Observability offers a modern platform to discover and map dependencies across cloud-native infrastructure and applications as well as a growing set of integrations with third-party automation and monitoring tools.

The Tanzu Observability service can consume, pool, and analyze time series data generated by applications/microservices, public or private cloud container services, and Kubernetes as well as data from other tools such as Prometheus. Recent updates have extended Tanzu Observability to support open distributed traces (e.g., OpenTelemetry, Spring Cloud Sleuth, Jaeger, and Zipkin) as well as cloud security metrics and application performance management (APM) data. The addition of OpenTelemetry support is particularly important in that customers will now be able to send application metrics, traces, and histograms to Tanzu Observability to allow for deep insights and correlations across applications and infrastructure without needing to rely on proprietary agents.

Tanzu Observability Application Maps provide an intuitive visual representation of dependencies across last-mile end nodes as well as core back-end systems and microservices. This provides insights into microservices communication bottlenecks and dependencies that can impact performance. Combined with the ability to set dynamic thresholds and automated alarming, this capability can help rapidly identify root cause. Baseline dashboards available as part of the core service enable both DevOps and ITOps teams to quickly gain insights and take action.

A recently announced partnership with Catchpoint provides business context by associating observability data with user satisfaction information obtained via end-user and endpoint monitoring. Integrations with VMware Tanzu Mission Control allow for a unified holistic management view and consumption experience across multiple Kubernetes clusters. Terraform integrations with standardized API support allow CloudOps and DevOps teams to define and code automated responses and remediations to errors and alerts. Spring Cloud Sleuth integrations support Tanzu Observability integrations with development pipelines.

Integrations with the VMware vRealize product line allow for increased collaboration between ITOps teams managing their own software-defined datacenters and DevOps/LOB teams developing applications. For example, Tanzu Observability can alert on and visualize data captured by vRealize Log Insight and enable in-context root cause analysis to machine logs, serving as a first pane of glass for IT operators and application developers to expedite problem troubleshooting and reduce mean time to repair (MTTR).

Challenges

For many enterprise customers, the organizational and policy issues related to implementing observability operations may be more challenging than the technical considerations. Many organizations will start small with DevOps teams focusing on the specifics of emerging applications and infrastructure with an emphasis on time series metrics such as those generated by Prometheus on Kubernetes clusters. Scaling up consistent, widespread use across the enterprise, including older applications and infrastructure, will require broader collaboration across more teams. The ability to correlate and respond quickly will be critical in terms of supporting large-scale multicluster production applications.

Tooling that is appealing to both platform site reliability engineers and developers is necessary to democratize monitoring and observability and to enable collaboration across the enterprise. VMware will need to help customers bring together diverse decision makers representing developer, ITOps, CloudOps, and security personas and provide solid integrations across traditional log management systems as well as modern cloud-native systems. As more VM-based applications are deployed into cloud-native environments, the integration of these solutions will be critical.

Conclusion

The success of many cloud-native applications and dynamic containerized infrastructure environments will be enabled by modern observability solutions that power faster discovery, assessment, and action across a growing array of on-premises, hosted, edge, and public cloud infrastructure and applications. IDC believes observability solutions deployed consistently across organizations will be important to optimizing end-to-end SLAs, security, and costs. To the extent that VMware can help customers enact collaborative, actionable best practices and decision making powered by comprehensive observability technology, the company will be positioned as an important innovation partner.

The success of many cloud-native applications and dynamic containerized infrastructure initiatives will be built on modern observability solutions.

About the Analyst



Mary Johnston Turner, Research Vice President, Cloud Management

Mary Johnston Turner is Research Vice President for Cloud Management, part of IDC's Infrastructure and Operations Management software research team. Her research focuses on emerging software and solutions for cloud, container and DevOps IT operations, cost optimization, automation, performance, and analytics. She contributes to vendor analysis as well as enterprise IT buyer advisory and custom consulting activities.

MESSAGE FROM THE SPONSOR

VMware Tanzu Observability by Wavefront

VMware Tanzu Observability by Wavefront delivers enterprise observability and analytics to optimize modern cloud-native applications by providing real-time operational insights from millions of data points per second. DevOps engineers and developers can interrogate real-time telemetry streams to resolve production issues faster, identify resource bottlenecks sooner, and understanding business impact. To learn more about Tanzu Observability or start a free trial, visit <https://tanzu.vmware.com/observability>.



The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
5 Speen Street
Framingham, MA 01701, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2020 IDC. Reproduction without written permission is completely forbidden.