



The Path to Multi-Cloud Success:

Why You Need Observability

By Janakiram MSV

Table of Contents

INTRODUCTION	04
CHAPTER 1: THE CHANGING IT LANDSCAPE	05
Legacy workloads	05
Private cloud	05
Public cloud	05
Containerization	06
Cloud native	06
Serverless computing	06
Edge computing	07
CHAPTER 2: THE RISE OF MULTI-CLOUD STRATEGIES	08
Hybrid cloud	08
Multi-cloud	09
CHAPTER 3: PLATFORM SUPPORT OF MULTI-CLOUD	10
Amazon Web Services	10
Microsoft Azure	10
Google Cloud Platform	11
VMware	11
IBM/Red Hat	11
CHAPTER 4: THE IMPACT OF MODERN CLOUD STRATEGIES ON IT	12
Developers	12
Operators	12
DevOps and Site Reliability Engineering (SRE)	12
The role of AIOps	13

CHAPTER 5: OBSERVABILITY FOR MULTI-CLOUD SUCCESS	14
Traditional monitoring falls short	14
Observability delivers a holistic view	14
Why teams need observability for multi-cloud environments	14
Match the scale	15
Monitor both legacy and modern workloads	15
Manage highly distributed environments	16
Making informed cloud transformation decisions	16
What observability means to key stakeholders	17
Developers	17
Operators	17
Business decision makers	17
CONCLUSION	18

About the author

Janakiram is a practicing architect, analyst, and advisor with a focus on emerging infrastructure technologies. He provides strategic advice to technology platform companies, startups, ISVs, and enterprises. As a practitioner working with a diverse enterprise customer base across the domains of AI, IoT, Cloud, Microservices, DevOps, and SaaS, Janakiram gains insight into the enterprise challenges, pitfalls, and opportunities involved in emerging technology adoption.

Janakiram is an Amazon-, Google-, and Microsoft-certified cloud architect. He is an active contributor at Gigaom Research, Forbes, and The New Stack.

Introduction

Managing enterprise IT has become complex and challenging. The ever-changing landscape of infrastructure, platforms, environments, and tools is creating unprecedented pressure on CIOs. While development teams are rapidly embracing cutting-edge technologies in the cloud, enterprise IT must ensure that traditional line-of-business systems will still run smoothly in the data center.

As more workloads shift between on-premises environments and cloud platforms, hybrid cloud models are becoming more prevalent. At the same time, enterprises are investing in multi-cloud strategies to avoid the risk of provider lock-in while taking advantage of best-of-breed services from the leading cloud vendors.

Enterprise IT teams are responsible for the availability, performance, and security of the applications they support, and in today's diverse IT environments—which span physical machines, virtual machines, ephemeral containers, stateful containerized workloads, orchestration engines, and serverless platforms—this can be a daunting task. To ensure business operations run smoothly, enterprise IT teams need visibility and insight not only into the infrastructure they manage, but also the platforms, applications, and client-side experiences that infrastructure supports.

Traditional monitoring and instrumentation techniques can't deliver what this ever-growing and dynamic IT environment needs most: end-to-end visibility of the entire stack. To get this level of visibility, enterprises are turning to an observability platform that empowers the IT team through proactive insights rather than reactive alerts.

The objective of this white paper is to help businesses understand the impact of the rapid transformation of the infrastructure and application landscape, the key players and the ecosystem behind the transformation, and the top trends that are driving the technology shift. It explains how a contemporary observability platform becomes the foundation of the modern infrastructure and application stack.

Chapter 1:

The Changing IT Landscape

Back in 2018, IDC predicted that 57% of total enterprise infrastructure investments would shift from physical infrastructure (buildings, heavy machinery, roads, etc.) to digital infrastructure (computers, systems, software, and services).¹ While the cloud is considered to be the key driver of this shift, other factors are accelerating digital transformation in enterprises.

Today, IT teams are managing an extremely diverse set of infrastructure and application platforms, from on-premises to multi-cloud environments and from legacy mainframe or midrange applications to more modern apps deployed in the cloud. This situation puts tremendous pressure on IT as internal users and lines of business expect IT to deliver a stable, highly available, and consistent experience for infrastructure and platform services.

As customers embark on their cloud and multi-cloud journeys, they face the following environments and challenges.

Legacy workloads

Large enterprises have IBM mainframes running transactional systems, x86 servers running client-server software, and desktops running older GUI applications. Many organizations rely on legacy workloads and databases, which cannot match the agility, flexibility, and scale of modern cloud-based services and applications. Because these applications are critical to the business and tightly integrated with other legacy systems, they may never migrate to the cloud. These workloads continue to run on-premises on legacy hardware and software. Nonetheless, a majority of organizations are seriously evaluating the option of modernizing these workloads.

Private cloud

Private cloud has been the next logical step in data center virtualization. It enables enterprises to derive the benefits of public cloud such as self-service, elasticity, and programmability within their on-premises IT environment.

Companies that invested in virtualization and hypervisors in the mid-2000s migrated the infrastructure to a private cloud. The private cloud stack—typically running on VMware vSphere, Microsoft Hyper-V, or OpenStack—delivers the benefit of the cloud to internal users.

Public cloud

Public cloud enables enterprises to extend their infrastructure to a remote data center managed by a third-party vendor. It differs from a typical co-location facility because of the public cloud's pay-as-you-go, self-service model as well as programmability and the ability to shrink and expand the infrastructure on demand.

The past decade witnessed enterprises' mainstream adoption of the public cloud as they moved public-facing, non-mission-critical workloads to Amazon Web Services, Microsoft Azure, Google Cloud Platform, IBM Cloud, and/or other public cloud platforms. From there, many enterprises began using the cloud for business continuity, where public cloud regions became remote sites for disaster recovery. Today, enterprises have workloads running across multiple public cloud platforms.

¹"Surviving and Thriving in a Multicloud World," IDC InfoBrief sponsored by Nutanix, April 2018

Containerization

Docker made containers mainstream by making them accessible to developers. Containers are lightweight entities that isolate programs running within the same operating system. They are considered to be better than virtualization due to their speed of deployment, portability, and agility.

In 2015, Google announced **Kubernetes**, an open source project for managing containers at scale. Within a short period, Kubernetes became the de facto software to orchestrate and manage containerized workloads. Today, almost every public cloud provider offers a managed Kubernetes service—called Containers-as-a-Service (CaaS)—in their environment.

With the rise of Docker and Kubernetes, containers found their place in the enterprise. Currently, developers in individual departments are evaluating the benefits of moving to containers. According to a survey conducted by StackRox, a container security company, 29% of organizations using containers are running more than half of their workloads in production, a growth rate of 32% compared with the previous year. Hybrid deployments are used by 46% of organizations, while 35% of the respondents are running in multiple public clouds.²

Cloud native

Kubernetes has emerged as the foundation of cloud native infrastructure and applications. It is available in multiple flavors, ranging from the managed container platform in the public cloud to a Platform-as-a-Service (PaaS) running within the enterprise data center.

Kubernetes delivered what was lacking in Infrastructure-as-a-Service (IaaS) offerings: standardization and consistent infrastructure. IaaS was built on the foundation of virtualization, which varied between each implementation. Each IaaS provider chose a different hypervisor, the technology that enabled

virtualization, resulting in vendor lock-in and lack of portability. Migrating customer workloads running in one IaaS environment to another is complex and expensive.

Because Kubernetes is the lowest common denominator of the infrastructure—both on-premises and the public cloud—customers can seamlessly move cloud native workloads across environments. Kubernetes and related cloud native technologies are delivering a standardized, consistent platform, making workload portability a reality.

Serverless computing

Applications are increasingly event-driven, where they respond to triggers generated by internal and external systems. For example, a shopping cart service may be invoked only when an order completes. The service will consume system resources only during the invocation and may get terminated when the execution completes. This design pattern, where no pre-provisioning of resources is required for running services, is called **serverless computing**.

This is fundamentally different from the standard approach of continually running processes that wait for a request while continually consuming resources. Instead, serverless computing delivers event-driven computing environments that can run software based on external events such as an order confirmation or a change in inventory.

Microservices rely heavily on this on-demand, event-driven execution environment for orchestrating business processes. AWS Lambda in the public cloud, OpenFaaS, Apache OpenWhisk, and Knative for on-premises containerized platforms deliver serverless computing. A significant part of the application code is deployed in these environments, along with virtual machines and containers.

² "The State of Container and Kubernetes Security," StackRox, 2020.

Edge computing

Today's public cloud is similar to a mainframe where the compute resources are centralized in one location. This model results in increased latency and bandwidth costs when sending massive amounts of data to the cloud for processing. Especially in industrial environments where connectivity is intermittent, running everything in the public cloud is not ideal. Edge computing offers a middle ground by moving a subset of cloud resources to a local environment, which can be accessed even in offline scenarios.

In environments where industrial automation is critical, enterprises are investing in edge computing, which brings the compute layer closer to the data layer. Operational technology (OT) and information technology (IT) collaborate in deploying, managing, and maintaining edge computing. Stream analytics, complex event processing, and machine learning models are some of the popular workloads that run at the edge. Microsoft Azure IoT Edge, AWS IoT Greengrass, and VMware Pulse IoT Center are some of the main-stream edge computing platforms.

Chapter 2:

The Rise of Multi-Cloud Strategies

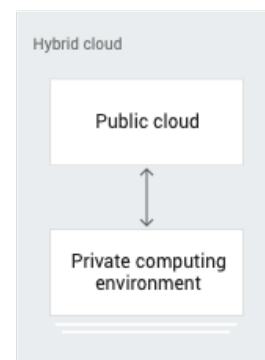
According to research and advisory firm Gartner, the worldwide public cloud services market is forecast to grow 17% in 2020 to a total of \$266.4 billion, up from \$227.8 billion in 2019. Sid Nag, research vice president at Gartner, says, “At this point, cloud adoption is mainstream. ... Building, implementing, and maturing cloud strategies will continue to be a top priority for years to come.”³

These strategies include hybrid cloud, multi-cloud, and the combination of the two. Most enterprises are adopting a multi-cloud strategy for a range of different reasons, including reducing risk in their investment, avoiding provider lock-in, and taking advantage of best-of-breed services from the leading cloud vendors to address a company’s use cases.

Hybrid cloud

An enterprise data center is not a silo. Mission-critical systems that form the digital backbone of an enterprise continue to run within the data center. Modern applications that are built for and deployed in the cloud need seamless access to the systems running on-premises.

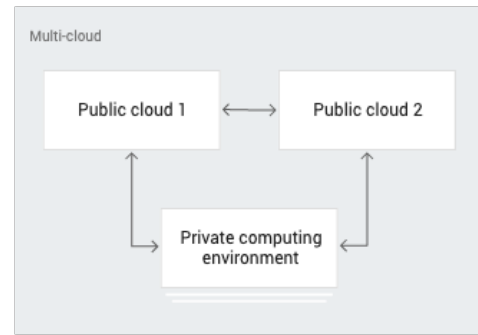
Large enterprises such as **Johnson & Johnson**, **Comcast**, and **Pacific Life Insurance** that are investing in the cloud have a well-defined hybrid cloud strategy. These companies are securely extending their data centers to connect existing assets with modern applications running in the cloud. With a hybrid cloud approach, enterprise IT continues to enforce the same policies and governance model that spans on-premises infrastructure and the public cloud.



³"Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17% in 2020," Gartner press release, November 2019.

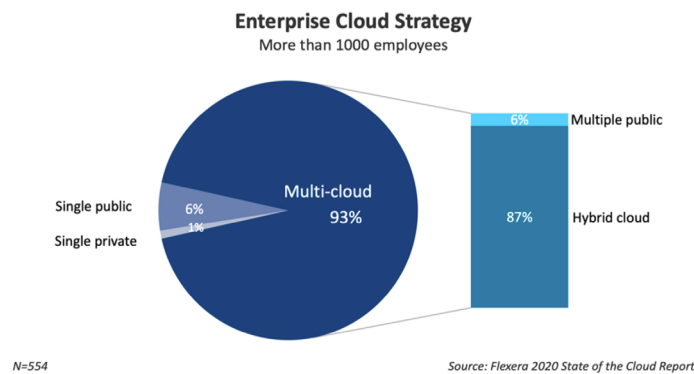
Multi-cloud

The public cloud has won the confidence of enterprise decision makers. Unlike startups that are born in the cloud, enterprises are taking a different approach to adopting the cloud by investing in more than one cloud platform to deploy their workloads. In fact, a 2020 survey reports that 93% percent of enterprises (companies with more than 1,000 employees) have a multi-cloud strategy and are using an average of 3.4 public cloud platforms.⁴



With Kubernetes becoming the universal control plane to manage both legacy and modern infrastructure environments, customers are inclined toward multi-cloud and hybrid cloud platforms based on Kubernetes. Thanks to Kubernetes, hybrid cloud platforms not only enable workload portability but also deliver the ability to scale the workloads across disparate environments.

Some of the recent offerings from cloud vendors manage clusters running on-premises and in their own cloud platforms as well as any Kubernetes cluster, including those that are deployed in other cloud environments.



Two important factors are driving the adoption of multi-cloud: 1) choosing best-of-breed cloud services and 2) workload portability delivered by cloud native platforms based on Kubernetes.

Enterprise customers are investing in best-of-breed cloud services aligned with business use cases. For example, a storefront application is deployed in the public cloud. It streams logs and aggregated transactional data to a data warehouse running on a different cloud platform, chosen for its analytics strengths. Similarly, a customer might choose to run a legacy content management system (CMS) running on Microsoft Windows deployed in an IaaS environment that is optimized for that workload, while maintaining content backups in a different cloud storage service.

⁴ "Flexera 2020 State of the Cloud Report," Flexera, 2020.

Chapter 3:

Platform Support of Multi-Cloud

As enterprises embrace hybrid and multi-cloud strategies, cloud providers are responding with new services and offerings that help customers connect on-premises with infrastructure across multiple cloud platforms.

Amazon Web Services

Amazon Web Services (AWS) is one of the first public cloud platforms to offer extensive hybrid capabilities. From dedicated connectivity to running Amazon EC2 instances on-premises, AWS has a wide range of services that enable the hybrid cloud.

AWS Outposts brings the elasticity and programmability of Amazon EC2 to the data center. It is a fully managed service where the physical infrastructure is delivered and installed by AWS, operated and monitored by AWS, and automatically updated and patched as part of being connected to an AWS Region. Customers using AWS Outposts can run a wide range of managed services within their data center, including Amazon ECS and Amazon EKS clusters for container-based applications, Amazon EMR clusters for data analytics, and Amazon RDS instances for relational database services.

Apart from AWS Outposts, Amazon customers can also choose to deploy **AWS Snowball Edge** on-premises. Snowball Edge is a data transfer device with on-board storage and compute power that provides select AWS services. Snowball Edge comes in two options, Storage Optimized and Compute Optimized, to support local data processing and collection in disconnected environments such as ships, windmills, and remote factories.

Microsoft Azure

Microsoft's **Azure Arc** is a hybrid and multi-cloud platform that extends Azure to on-premises and other public cloud platforms. With Azure Arc, Microsoft is taking a different approach to hybrid cloud. Customers can register bare metal servers, virtual machines, and Kubernetes clusters with Azure Arc to manage them from a centralized control plane. Azure Arc enables customers to run managed databases, Kubernetes clusters, and servers across a variety of environments, including data centers, public cloud, and the edge.

Azure Kubernetes Service (AKS), Microsoft's managed Kubernetes service, plays an important role in Azure Arc. As a managed service in the public and hybrid cloud, it is responsible for running containerized workloads. Customers can mix and match legacy workloads running on Linux and Windows virtual machines with modern cloud native applications running on Kubernetes.

Apart from Azure Arc, Microsoft offers the **Azure Stack** family of products that extend Azure services and capabilities to the data center and edge. Azure Stack appliances bring compute, storage, and networking closer to the origin of data.

Google Cloud Platform

Google is betting big on modern applications built as microservices running in Kubernetes, basing its hybrid and multi-cloud strategy on this approach.

Anthos is a hybrid and multi-cloud application platform that enables GCP customers to modernize existing applications, build new ones, and deploy them in multiple environments. Anthos is built on the strong foundation of Google Kubernetes Engine (GKE), a managed Kubernetes service on Google Cloud Platform.

Customers can seamlessly manage GKE clusters deployed in the cloud or on-premises through Anthos. Anthos Service Mesh, Google Cloud's fully managed service mesh, enables management of complex microservices architectures through mesh telemetry, traffic management, and declarative security policies.

Anthos comes with a tool to migrate virtual machines to containers that can run in Kubernetes clusters. It can migrate Amazon EC2 instances, Azure virtual machines, and VMware virtual machines to Anthos.

VMware

VMware is on a mission to make vSphere the best platform to run both virtual machines and containers. To that end, it has acquired Heptio and Pivotal Software to complement the container strategy.

VMware announced **vSphere 7**, a unified platform to manage virtual machines and containers. VMware vSphere 7 is a convergence of vSphere and Kubernetes that exposes both APIs. To developers, the platform looks like a Kubernetes cluster where they can use Kubernetes declarative syntax to manage cloud resources such as virtual machines, disks, and networks. To IT administrators, vSphere 7 looks like vSphere, but with the new ability to manage an entire application instead of the individual virtual machines that make it up.

While vSphere 7 is meant to run both virtual machines and Kubernetes clusters, **VMware Tanzu Mission Control** is designed to act as a single pane of glass

to manage Kubernetes clusters deployed anywhere. Similar to Google Anthos, Tanzu Mission Control can register Kubernetes clusters, centrally manage policies and configuration, and enable developer self-service.

vSphere 7 and Tanzu Mission Control can run in any VMware environment that includes enterprise data centers and public cloud platforms. With its partnership with cloud vendors, VMware is available on mainstream public clouds, including AWS, Azure, IBM, and GCP.

IBM/Red Hat

With the acquisition of Red Hat, IBM is heavily relying on OpenShift for its hybrid and multi-cloud strategy.

IBM Cloud Paks, one of the hybrid cloud offerings from IBM, delivers a containerized software solution to build contemporary cloud native applications as well as to modernize traditional applications. Enterprises can choose from a variety of Cloud Paks to run their workloads. At the bottom of the stack, there is Red Hat Enterprise Linux. The next layer is powered by Red Hat OpenShift Container Platform, one of the few enterprise-ready container orchestration platforms available today. The OpenShift Container Platform manages a set of containers deployed as a part of Cloud Pak.

Similar to Azure Arc and Google Anthos, **IBM Cloud Pak for Multicloud Management** is a single dashboard for managing multiple Kubernetes clusters across public and private clouds. It provides centralized control, consistent security policies across environments, and advanced automation.

Red Hat OpenShift is one of the fastest-growing enterprise PaaS offerings on the market. Because it runs in both public and private clouds, customers prefer OpenShift for a consistent developer experience. Combined with IBM MultiCloud Manager and IBM Cloud Paks, OpenShift is transforming into a hybrid and multi-cloud PaaS.

Chapter 4:

The Impact of Modern Cloud Strategies on IT

As enterprises accelerate cloud adoption, internal stakeholders—developers, operators, DevOps teams, and site reliability engineering (SRE) teams—face increased pressure to ensure application performance and availability in complex, distributed environments.

Developers

Applications are moving from being monoliths to microservices, a design pattern to build modern applications composed of independent services. Because containers are smaller when compared with virtual machines, microservices are built as a collection of containers and deployed on platforms such as Kubernetes or Red Hat OpenShift. These modern applications, also known as cloud native applications, are portable and built to take advantage of key attributes of the cloud, such as elasticity, scalability, and programmability.

Developers building microservices-based cloud native applications use best-of-breed languages, libraries, runtimes, and tools. Unlike traditional applications running within virtual machines, cloud native applications translate to tens of thousands of instances at runtime. Developers need to debug these distributed applications running across multiple containers, clusters, machines, and environments. They also need to ensure that the application generates key telemetry data that will help the operations team assess the health of the application at runtime.

Operators

Operators are expected to manage the increasingly disparate infrastructure and platform services spanning on-premises, public cloud, PaaS, and CaaS environments. Internal business teams rely on IT and operations teams to deliver against service level agreements (SLAs) for uptime, availability, and scalability.

Operators need to gain insight into the infrastructure and applications to reduce mean-time-between-failures and mean time to resolution (MTTR). They must handle the deluge of data originating from the infrastructure, storage, network, and application platforms to identify and remediate issues quickly.

The operations team needs a way to aggregate, process, analyze, and visualize metrics, events, logs, and tracing information generated by a variety of workloads deployed across private and public clouds.

DevOps and Site Reliability Engineering (SRE)

DevOps has emerged as a culture and a set of best practices that encourages collaboration between the software development and the software operations teams. However, the DevOps movement does not explicitly define the best approach to implementing key principles for demonstrating efficiency.

SRE ensures that the core principles of DevOps are implemented in the right way. SRE—which originated at Google independently of the DevOps movement to meet Google’s internal needs—embodies the core philosophy of DevOps, but has a much more prescriptive way of measuring and achieving reliability through engineering and operations work.

SRE ensures that all stakeholders agree on how to measure availability and what to do when availability falls out of specification. As complexity increases due to disparate environments and diverse workloads, SRE becomes critical to the success of enterprise IT.

The role of AIOps

Artificial intelligence (AI) is going to have a significant impact on system operations and administration, redefining the way infrastructure—both enterprise data center and cloud—is managed through AIOps, the convergence of AI, and traditional IT operations.

IT infrastructure generates a lot of data. From the temperature of the chassis to the latency rate of an outbound API call, it is possible to acquire data from disparate layers of the stack. When aggregated, normalized, and analyzed, this data becomes a rich source from which to derive insights.

By taking an algorithmic approach to IT operations, DevOps and SRE teams can take advantage of intelligent capacity planning, optimal resource utilization, instant outlier detection, and faster remediation. With AI, workloads can be mapped to the right configuration of servers and virtual machines. After running the workload in its peak state, AIOps can recommend the correct instance family type, storage choices, network configuration, and even the IO throughput of storage.

With AIOps, administrators can rely on predictive scaling in which the infrastructure intelligently adjusts itself based on historical data. It learns how to reconfigure itself based on current and anticipated utilization patterns. AIOps can detect anomalies in near real-time from the data originating from various layers of the technology stack, preventing potential outages and disruptions.

Chapter 5:

Observability for Multi-Cloud Success

As organizations move to microservices running on modern infrastructure based on hybrid and multi-cloud environments, they need a robust mechanism to deliver insight and track the full estate, from infrastructure through applications and the digital customer experience.

Traditional monitoring falls short

Traditionally, infrastructure and application monitoring tools provided baseline capabilities for monitoring the health of operating systems as well as workloads such as databases, web servers, messaging services, and others. The user community of these tools built a variety of plugins to monitor and report the state of systems to support engineers.

Monitoring systems provide an outside-in view, with no understanding of the internals of the software they are monitoring. The key expectation from these tools is that they periodically perform a check and report the state to a central location, which is a portal or dashboard.

Monitoring is a passive, reactive mechanism of tracking the infrastructure and applications. It reports only that there is a change in the state instead of identifying what influenced the change in the state. In other words, monitoring is primarily focused on what is happening in the environment, not why it is happening, which is where observability comes into the picture.

Observability delivers a holistic view

Observability delivers the inside-out view, which is rich, meaningful, and contextual. It delivers on the promise of continuous monitoring, a key element of DevOps, by providing contextual insights not just from the infrastructure but from the entire stack, including end users.

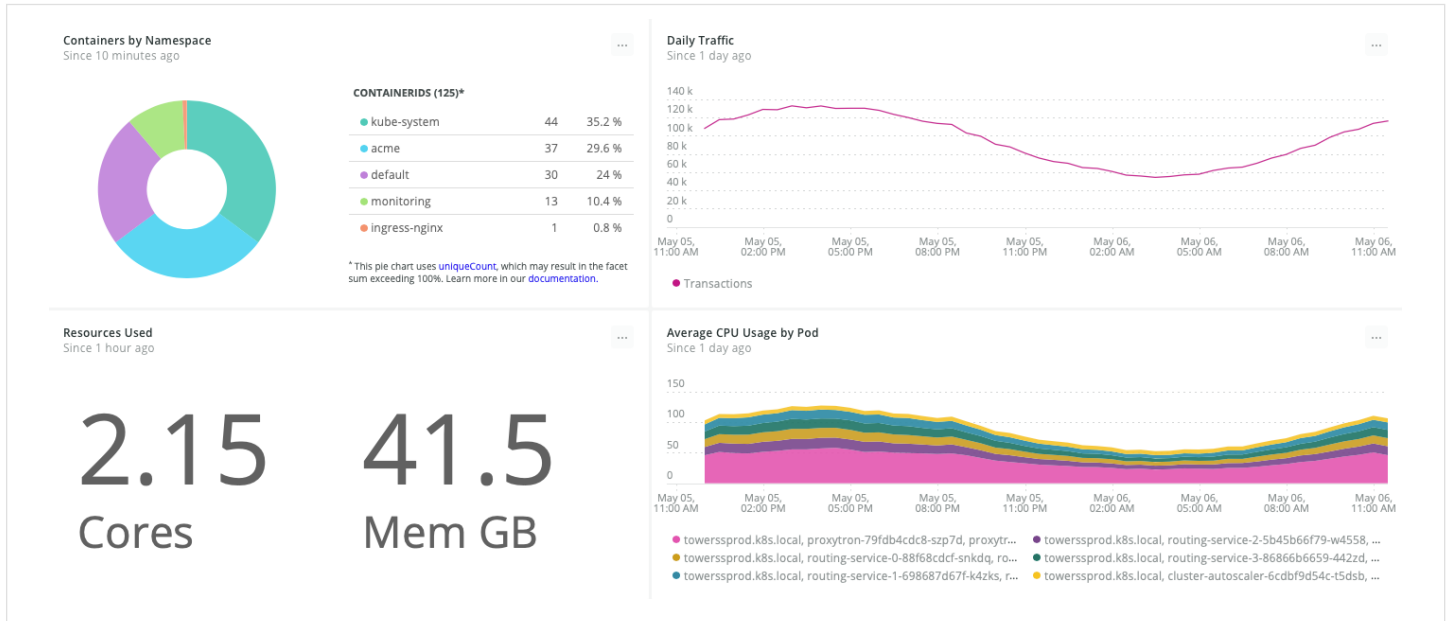
Instead of relying on external agents that periodically check the state and collect the superficial metrics, observability becomes an all-encompassing layer that provides insights across the stack—network, hardware, operating systems, and applications, as well as internal and external services. It infers the internal health of systems from external metrics and data points and is proactive in reporting the change in state.

For optimal observability, DevOps teams need to identify and ingest four essential types of telemetry data—**metrics, events, logs, and traces (MELT)**—from code, infrastructure, applications, end-user interactions, and external systems.

Why teams need observability for multi-cloud environments

The emerging hybrid cloud and multi-cloud platforms enable businesses to deploy applications spanning the data center and diverse public cloud environments. One of the key benefits of modern, multi-cloud platforms is the optimal utilization of infrastructure resources. Workloads can seamlessly move between on-premises and cloud infrastructure with minimal effort and almost no disruption.

But this approach puts additional pressure on developers and operators managing the application life cycle. They face the following challenges in the new multi-cloud environment.



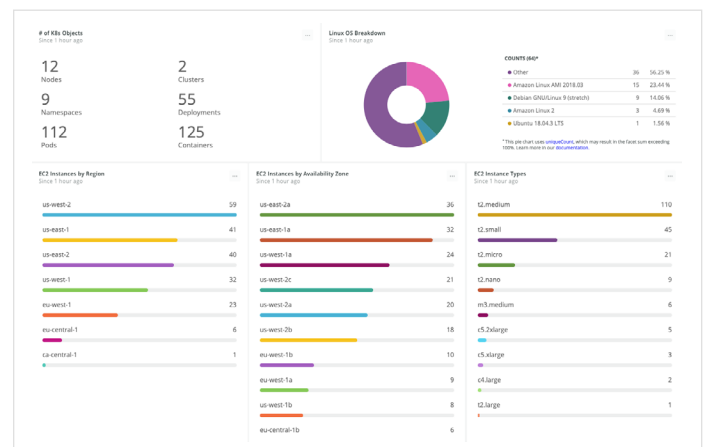
Observability correlates customer traffic against backend resources.

Match the scale

Kubernetes is designed to scale exponentially. Workloads can dynamically grow and shrink based on external factors such as incoming traffic, resource availability, and custom policies. DevOps teams not only need to monitor the availability of workloads but also gain insight into the environment. They need visibility into tens of thousands of containers running across hundreds of nodes.

Monitor both legacy and modern workloads

While contemporary applications target containers and Kubernetes, enterprises will continue to have legacy workloads that run on physical servers, virtual machines, and private cloud. DevOps teams need an effective strategy and tooling to monitor and manage workloads running in a hybrid environment. Organizations need a unified approach to observability, performance monitoring, and application life cycle management.

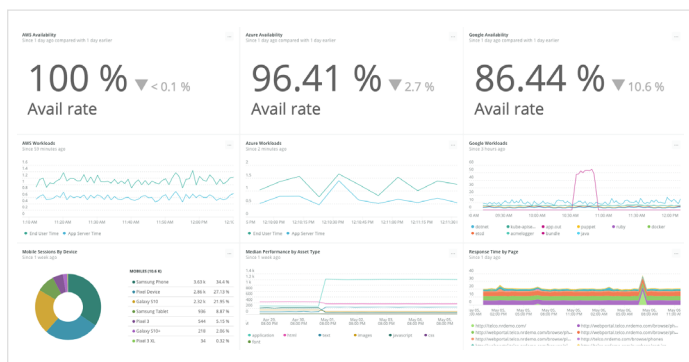


Visibility across container and virtual machine infrastructure

Manage highly distributed environments

Modern hybrid and multi-cloud platforms are based on a fabric that connects on-premises data centers with multiple public clouds. While these platforms provide centralized management and federation, observability is left to customers to solve.

DevOps teams need a unified, centralized observability platform aligned with the multi-cloud design and architecture implemented by platform vendors. Unless developers and operators have complete visibility into every workload, server, virtual machine, cluster, pod, container, and cloud resource, they can't deliver on the promise of high availability and maximum uptime.

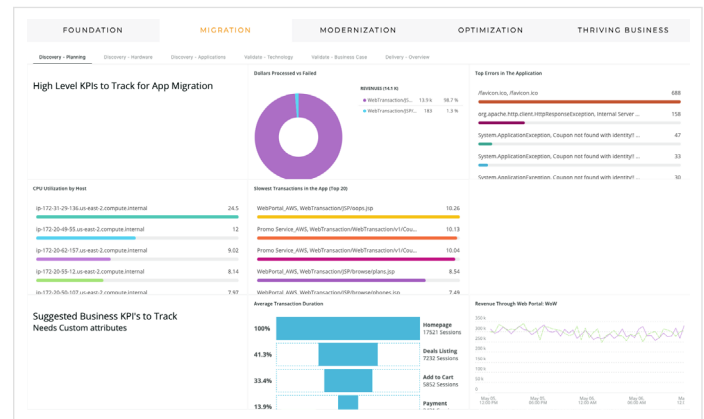


Multi-cloud visibility correlated against frontend experiences

Making informed cloud transformation decisions

To minimize risk and measure the results of change from the customer perspective, organizations need an observability platform throughout their cloud journey. Observability allows DevOps teams to make informed decisions before, during, and after changes, including architecture, cloud migration, modernization, and optimization efforts:

- **Before:** Use baseline metrics to understand the current state and make an informed decision on the future state.
- **During:** Rely on real-time data to reduce risk and MTTR if issues arise during the change.
- **After:** Validate the change by comparing results to the baseline and mapping to business KPIs to demonstrate the value of the change.



Information about the planning phase of a cloud journey

What observability means to key stakeholders

A comprehensive and mature observability stack positively impacts the stakeholders within an organization.

Developers

As developers increase the velocity of delivering software through continuous integration and continuous deployment, observability provides insights into performance, code-level issues, and usability, while delivering efficiency in managing the life cycle of microservices deployed in a multi-cloud environment. They can compare and contrast the key metrics that influence application performance across multiple environments, such as development, testing, staging, and production.

Operators

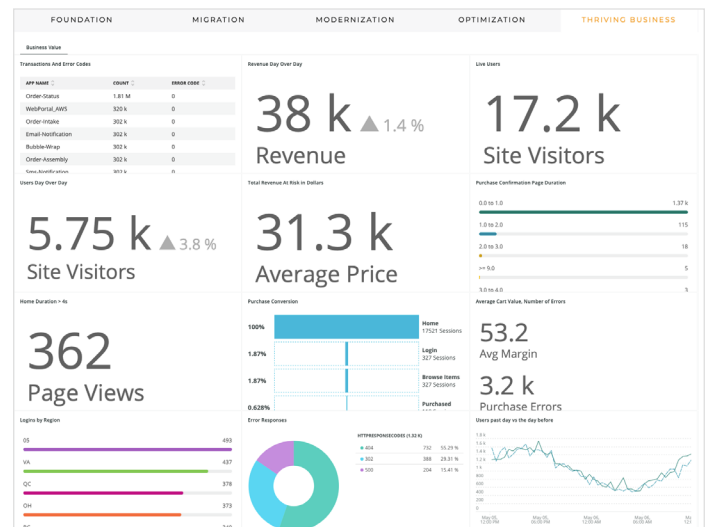
IT Ops and DevOps teams suffer from tool sprawl—a phenomenon used to explain the complexity of using a plethora of tools—which is aggravated in a multi-cloud environment. An observability platform lets operators see all of their observability data in one place, reducing risk and complexity and alleviating tool sprawl.

A well-defined observability strategy to monitor and track disparate environments is invaluable for operators. By focusing on key metrics that are curated and highlighted by an observability platform, operators can identify and remediate issues faster.

Business decision makers

The move to multi-cloud and emerging technologies introduces challenges for business decision makers who need a way to justify the return on investment from modernizing existing applications and demonstrate the value of investments made in developing greenfield applications.

A robust observability platform enables business decision makers to gain visibility into applications, infrastructure, and the customer experience. The key metrics generated by line-of-business applications can be mapped and transformed to match business KPIs. When integrated with existing business intelligence and analytics, these KPIs showcase the business impact of the investments in cloud native and multi-cloud approaches.



Business KPIs provide information for executive decisions.

Conclusion

There's no doubt that cloud computing is now main-stream. Enterprises are embracing sophisticated, cloud native technology and services to introduce innovation at speed, deliver differentiating digital customer experiences, and drive exceptional business outcomes.

But enterprises can't simply abandon their existing systems running on-premises. To that end, hybrid environments bridge the on-premises and public cloud infrastructure. Hybrid approaches are often multi-cloud in nature, with multiple public cloud platforms included.

The downside of these approaches is the complexity of managing and operating distributed applications across multiple containers, virtual machines, servers, serverless functions, and more. That's why observability is crucial for multi-cloud success. It's the key to achieving the promises and goals of a multi-cloud strategy and approach.

To learn more about observability and what's required to achieve it, read our ebook [The Age of Observability](#).