WHITE PAPER

BROADCOM

A Guide to Hybrid Cloud Monitoring

Table of Contents

Introduction 3

The Challenge of Hybrid Cloud Monitoring 3

Chapter 1: Understanding the Performance of Cloud Services: An AWS Example 3

Factors that Affect Performance 4

AWS-Provided Tools 4

Leveraging Third-Party APM Experts 5

AIOps and Proactive Monitoring with Thresholds and Alerts 6

Chapter 2: Top 4 Challenges to Monitoring a Hybrid Cloud Environment 6

Different Metrics and Tools to Track 6

Integrating the Entire Stack 6

Security 7

Cost Control 7

Shift to AIOps 7

Chapter 3: Best Practices for Monitoring a Hybrid Cloud Environment 8

Separate and Monitor with Microservices 8

Monitor Continuously Across Public and Private Clouds 8

Unify Infrastructure 9

Integrate Monitoring Tools 9

Conclusion 10

Introduction

For most of the 2010s, we were told that "the future is in the cloud." That still holds true. But today, we can be more precise: For many organizations, the future is in the hybrid cloud.

To some observers, this statement may come as a surprise. For a long time, moving workloads entirely to the public cloud—where they would enjoy unparalleled scalability—seemed to be the end-goal of many IT teams. At the same time, frameworks available for unifying public cloud services with private infrastructure in order to construct hybrid clouds, were lackluster.

Today, however, the game has changed, and hybrid cloud has come into its own. It is now recognized that problems like mounting public cloud costs and **data privacy limitations** mean that a cloud strategy based on public cloud alone is not a good fit for many workloads. And, the introduction in recent years of new hybrid cloud solutions—like Azure Stack and Google Anthos—has made it much easier to create a hybrid cloud that seamlessly weaves public cloud services into private data centers.

The Challenge of Hybrid Cloud Monitoring

What all of this means is that, going forward, IT teams are likely to be deploying and managing more workloads using a hybrid cloud architecture. In some respects, that's great, because hybrid clouds can reduce costs and provide security benefits not available from the public cloud alone.

Today, however, the game has changed, and hybrid cloud has come into its own.

Yet hybrid cloud also presents certain challenges, not the least of which are in the realm of monitoring. When you deploy workloads across public and private infrastructure, it becomes more difficult to monitor them and optimize their performance and cost than it would be if you kept all applications and data in a single public cloud. There are more management tools to juggle, more metrics to track and more variables to address when seeking to optimize reliability and performance.

With these challenges in mind, Broadcom has prepared this white paper to educate cloud admins and IT teams about hybrid cloud monitoring and optimization. In the following pages, you'll find tips on how to track performance in cloud environments, which special monitoring challenges arise in hybrid architectures, and best practices for monitoring hybrid environments.

Chapter 1: Understanding the Performance of Cloud Services: An AWS Example

Before delving into the details of monitoring hybrid clouds, let's start with an overview of what it means to monitor cloud services in general.

The most important part of providing software solutions for your customers is meeting their needs. To do this, there are two key things that you must consider: first, your software must perform the tasks that your consumers expect, and second, your software needs to perform those tasks quickly, efficiently, and accurately. In this chapter, we will explain this second component, focusing specifically on performance.

In particular, we will examine services that are deployed into the Amazon Web Services (AWS) public cloud. We'll explain what constitutes a high-performing service and why it's essential to monitor performance. Then we will explore tools for monitoring performance in the AWS environment and compare them to third-party monitoring solutions.

This chapter aims to help you understand the options that are available to you for monitoring and managing your services' performance, and also to help you get set up with the tools that will enable you to provide your consumers with the best possible user experience.

Factors that Affect Performance

When you devise a strategy for analyzing performance, you can divide what you monitor into two categories: underlying infrastructure and user experience. If the service is deployed on a virtual machine using AWS EC2, you should be aware of the following:

- Memory utilization: available vs. used memory
- CPU utilization
- Network utilization: incoming and outgoing data rates
- Disk utilization: input/output metrics and available vs. used storage

Each of these metrics will give you a window into whether or not you are using the appropriate infrastructure as well as how it is performing. You can also gain insights and implement changes if the instance is consistently operating at the upper or lower limits of its capacity. You might want to move services that operate at the upper limit to a type of instance that has more resources. In contrast, those that operate at the lower limit could be moved to a smaller instance type, which will reduce your operating costs.

If your application is containerized and deployed using AWS ECS or AWS EKS, then you will have fewer infrastructure metrics to monitor. You still want to be aware of CPU and memory utilization within your nodes and clusters, and you want to adjust your configurations accordingly.

The second part of your monitoring strategy should be monitoring user experience. If your users cannot reach your service, or if a load balancer is throttling their requests, you'll lose customers even if your infrastructure metrics look fantastic. To monitor consumer experience, you should be aware of the following:

- Request and response times
- Number of requests over time
- Error rates expanded by error type

Each of these metrics will help you understand the volume of calls that your service is handling, how long they are taking, and whether user requests are successful. It's essential to establish a baseline measurement for each one. Once you have a baseline, you should observe how each metric changes over time and then respond to deviations from that baseline.

AWS-Provided Tools

As a comprehensive cloud service provider, AWS provides access to Amazon CloudWatch for all of its customers. Amazon CloudWatch is a metrics repository for all AWS-hosted services. Its standard resolution is free and provides metrics at one-minute intervals. Users can also subscribe to its high-resolution offering, which provides metrics at one-second intervals. CloudWatch aggregates metrics over time to reduce storage demands, which means that your performance data will become less specific the longer that it's stored.

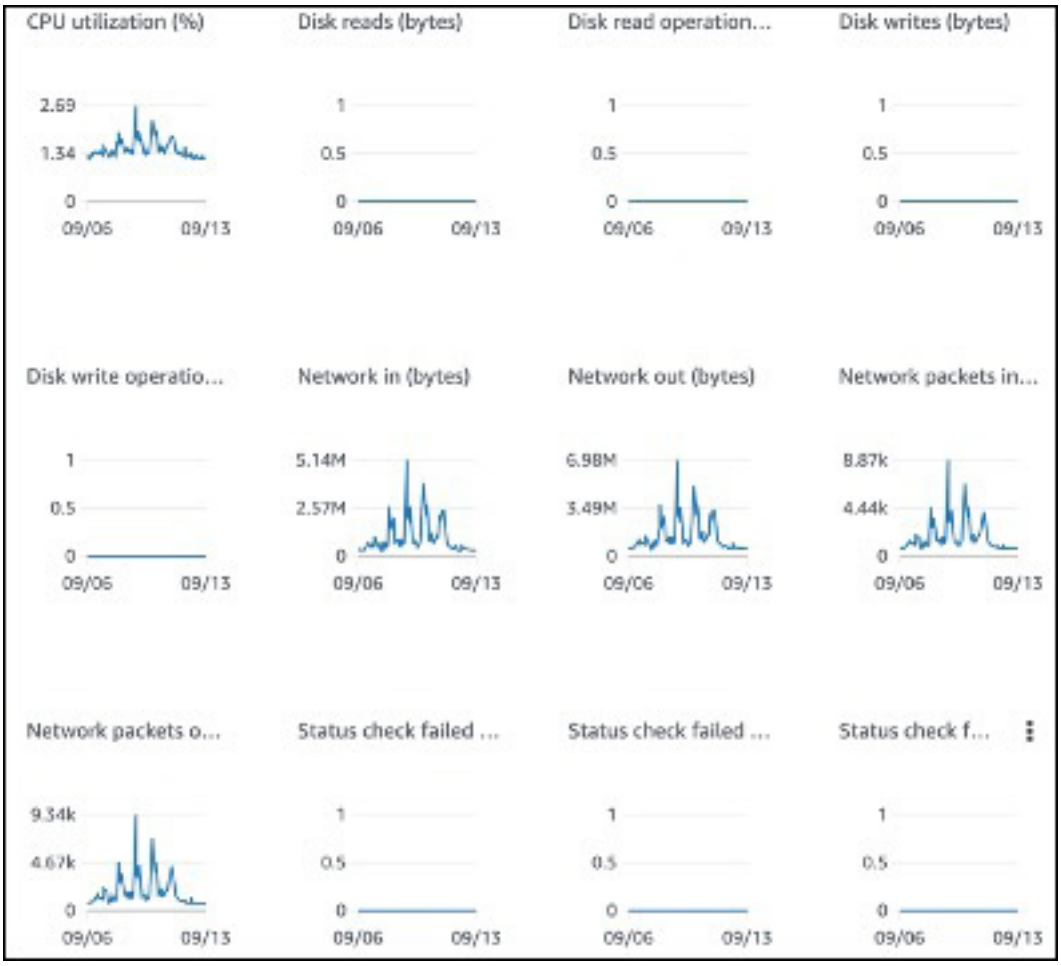

Figure A. Example of CloudWatch Metrics for an EC2 Instance

Amazon CloudWatch also allows you to set limits and alarms on your metrics. For example, you can set an alarm to trigger when the memory usage exceeds 85% on an instance for more than three minutes. You can also configure the alarm to send a message, start a predefined process, or connect to another service through a web-hook.

An unfortunate downside of Amazon CloudWatch is that, while it does provide access to a wide variety of metrics for each of its services, you need to know which metrics you're looking for and how to combine them to provide actionable insights. At its core, CloudWatch is just a collection and reporting service, so when it comes to anomaly detection and monitoring intelligently, you'll need something more.

Leveraging Third-Party APM Experts

Ideally, you'll want your engineers to devote their time to adding new features and improving your software's performance. One of the benefits of using standardized services and hosting them in the cloud is that you can leverage the expertise of those whose sole focus is on application performance management (or APM). You can add an agent to your instances or a sidecar application to your container environment that will gather essential metrics and transmit them to an APM provider.

APM providers typically provide standard dashboards and monitoring as part of their product offering. In most cases, you can enable these systems quickly and begin monitoring them intelligently with just a few hours of work. Some of these providers have recently started offering AIOps, which is an exciting option that adds exceptional value to your performance monitoring strategy.

AIOps and Proactive Monitoring with Thresholds and Alerts

Artificial Intelligence for IT Operations, or AIOps, combines machine learning and data science with a performance monitoring solution. It provides automated remediation capabilities, enables you to detect problems sooner, and ultimately improves your consumers' experience. AIOps can help you improve performance as well as identify new ways to increase your efficiency and performance.

If you would like to learn more about AIOps, how it works, and the potential benefits of using it, **The Definitive Guide to AIOps** is an excellent place to start. This white paper defines AIOps in more detail, explores the underlying principles and technologies, and explains how you can apply it to your organization. You can also download the **AIOps from Broadcom solution brief**, which provides specific details about Broadcom's AIOps product offering.

Providing a high-performing user experience is essential for meeting your users' needs. Fortunately, partnering with experts at organizations like Broadcom makes it easy to achieve this goal, ensuring that your software is reliable and that your customers can access it easily.

Chapter 2: Top 4 Challenges to Monitoring a Hybrid Cloud Environment

As noted in the introduction to this eBook, special monitoring challenges arise when you shift from monitoring a single-cloud environment to monitoring a hybrid architecture. This chapter offers a look at four critical challenges, and offers tips on addressing them.

Different Metrics and Tools to Track

The growing diversity of environments across private and public clouds makes monitoring more complex. Performance metrics for one environment differs from another. One environment may report metrics in seconds, while the other in one-minute intervals. Though tracking the same metric, their names and labeling differ and need to be correlated to be useful.

The tooling is also different for each platform. While organizations have their legacy monitoring tools like Nagios, they now also have cloud vendor monitoring tools like AWS CloudWatch and open source monitoring tools like Prometheus. There is some overlap between the metrics of each of these tools, while some metrics are unique to each tool.

The challenge is to unify all these metrics and attain a unified view of the hybrid system, end-to-end. This "single pane of glass" view is the holy grail of hybrid cloud monitoring. None of the purpose-built monitoring tools can deliver an end-to-end view. That requires a separate tool that can integrate all metrics from all tools and make them available in a way that is meaningful and usable.

Integrating the Entire Stack

The private and public clouds need to be integrated with each other at all levels—infrastructure, data, networking, and application. At the infrastructure layer, instances need to be spun up and destroyed between the private and public cloud environments as workloads are shifted between the two.

At the data layer, storage and transfer of data needs to be seamless between the multiple environments. Additionally, some requests could require data across environments to be processed.

At the networking layer, things like load balancing and service discovery should cover all environments. Also, during these times of remote work, VPN access has taken center stage in most IT organizations. Finally, applications should be integrated with API-based integrations. These APIs should be compatible across the board.

With so many moving parts at every layer of the stack, it's easy to see why things can go wrong with hybrid cloud. SLAs aren't uniform as there are multiple vendors to be managed, which brings more responsibility in-house to the organization itself. When these failures happen, it disrupts the end user experience.

Security

As the stack expands, so does the attack surface. With additional components and services to secure, security monitoring is of key importance.

For a data center, security practices start with securing the physical facility and hardware. Then, there are the network and device security measures like firewalls and anti-virus software. At the application level, user access needs to be configured via SSO or LDAP. Finally, data needs to be secured for data loss or disaster recovery.

Some of these practices like security of physical premises are rendered moot in a cloud platform. But some, like data backup, need to be continued even in the cloud.

The cloud operates on a shared responsibility model where the cloud vendor handles security of the platform; whereas the organization would still be responsible for their security 'in' the cloud platform. Cloud security involves a completely different approach to IAM, and new tools for data encryption and key management.

This makes compliance and governance all the more challenging as it needs to span both private and public clouds. Finally, throw in threat monitoring that is essential to monitor for phishing, DDoS attacks, and downloading of vulnerable container images—and you have a security nightmare.

Cost Control

More resources drives up the TCO (total cost of ownership) quickly. If unused resources were a drawback with on-prem, that problem is easily exacerbated in the cloud. The cloud is cheap at the start, but as the traffic volume grows, and the number of cloud services being used increases, it's easy to inadvertently run into sticker shock.

Monitoring is essential to prevent this. It requires keeping track of resource utilization at the infrastructure level. Monitoring done right should yield opportunities to reduce costs with hybrid cloud without compromising on performance. Additionally, it requires alerting whenever usage crosses a threshold.

Shift to AIOps

To counter these challenges, organizations need a completely different monitoring practice; something that leverages machine learning and artificial intelligence to augment humans and monitoring tools. AIOps (Artificial Intelligence in IT Operations) is the answer to this challenge. AIOps combines monitoring for all the purposes listed above and provides a "single pane of glass" view of hybrid cloud.

CIOs looking to make the transition to a modern and agile cloud system should leverage the power of AIOps, which can help them meet the demands of hybrid cloud. AIOps will also help make this transition seamless as it builds confidence when running and managing a newly-set up hybrid cloud.

Chapter 3: Best Practices for Monitoring a Hybrid Cloud Environment

According to Keysight Technologies, monitoring is a top concern for organizations operating hybrid cloud systems.

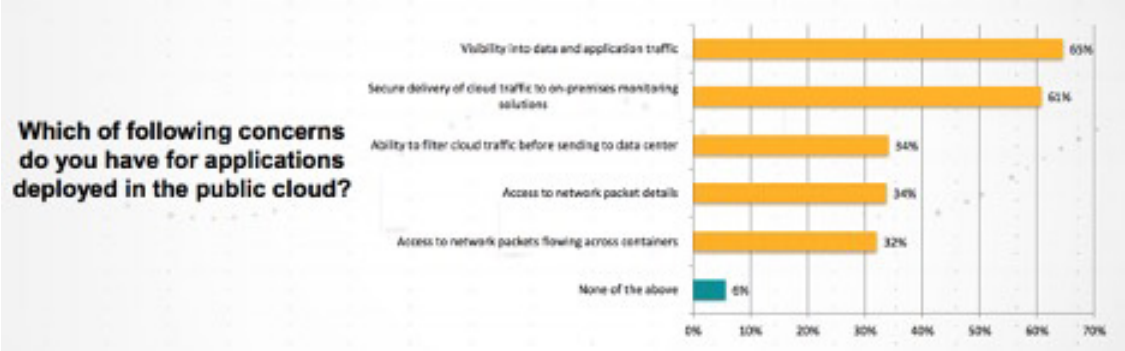

Source: **Keysight Technologies**

Visibility into traffic data, and delivery of data to on-premises monitoring solutions were the top concerns. Indeed, in the rush to go cloud, many organizations find that their cloud monitoring tools are only compatible with one cloud vendor, or with a couple of public cloud platforms, but not with their private cloud or on-prem environments. This poses a big problem for organizations looking to gain an end-to-end view of the system.

Separate and Monitor with Microservices

To make the most of hybrid cloud, your apps would need to be nimble enough to allow for services to be moved between clouds. In terms of monitoring, microservices brings deeper visibility into every layer of the applications stack. It allows you to view services independently, and in combination with other services. It separates applications from infrastructure, and from networking, and so on.

Once an issue is identified, microservices makes it easier to pinpoint the root cause, and even rectify it. For example, an issue that affects users in one region can be resolved without affecting users in other regions.

Monitor Continuously Across Public and Private Clouds

The reason for hybrid cloud is that some workloads require the security and exclusivity of a private cloud and some are best run in the public cloud. The key is to know which to run where.

A private cloud is ideal for applications that have specific requirements for compliance, performance, security or security that public clouds are unable to meet. The public cloud is great for apps being built from scratch, and for apps that need cloud burst capabilities.

Some apps can be run on a combination of public and private cloud infrastructure. For example, you could use the public cloud to run the application, as it can easily scale to meet demand; and use the private cloud for data storage for that added layer of security.

When it comes to monitoring these workloads, the last type of applications would be the most challenging as its workloads would span both private and public cloud. What can help make managing these apps easier is to report on them in a single unified AIOps monitoring tool. Further, as workloads are shifted between private and public clouds, the AIOps monitoring tool should be aware of these changes and be able to bring continuity to monitoring.

Unify Infrastructure

While hybrid cloud is inherently distributed, how distributed it is depends on the business requirements and the preferred architecture. This has a bearing on how many cloud vendors make up the spread of hybrid cloud environments.

One organization may prefer to use a mix of many cloud vendors and services. This is the best- of-breed approach where you stop at nothing in your quest for the perfect cloud experience. For example, you may want to opt for a smaller cloud vendor to leverage their expertise with OpenStack; or a particular vendor's serverless platform for its ease of use. You may value the ability to change vendors in the future, and may not want to commit too much to a single vendor. What can help in this case is an AIOps monitoring tool that can adapt to your changing infrastructure requirements.

The other approach is to centralize on a single cloud vendor. The big three cloud vendors—AWS, Azure, and Google Cloud—would be an ideal choice for this setup. They have been making conscious efforts to strengthen their hybrid cloud offerings of late. They each have announced a major focus on hybrid cloud solutions with their services like Google Anthos, Azure Arc, and AWS Outposts. These services enable the creation and management of resources across their cloud platform and on premises. However, they would be able to provide only basic monitoring; beyond which, you'll need a more robust AIOps monitoring tool that can provide in-depth and advanced monitoring capabilities.

Integrate Monitoring Tools

The hybrid cloud model is made up of a mix of containers, VMs, and physical servers. There is a mix of traditional IT and modern cloud-native networking and storage components to be monitored alongside each other. They are best managed via policies rather than manual configuration. Policies enable you to achieve SLOs while maintaining the scale and diversity of the hybrid cloud.

Monitoring is key to implementing and maintaining policies across the stack. It takes an AIOps monitoring tool to gather all meaningful monitoring data from various disparate sources, bring them together, and extract meaningful insight from them.

By integrating monitoring tools using AIOps, it simplifies ad-hoc analysis that is essential for troubleshooting and day-to-day monitoring operations. Practically, this allows users to throw different metrics into the AIOps tool and view them in a single dashboard.

Importantly, you should be able to do this in a matter of seconds, or at most, minutes. This speed is critical when incidents occur. Having the right monitoring data, and being able to analyze it to identify the root cause is where the rubber meets the road. A modern AIOps monitoring tool would excel at separating the signal from the noise and suppressing false alarms.

Further, it should have a mature alerting and notifications system so that the right people are informed at the right time. An AIOps tool helps to bring together diverse teams in one place and enables them to 'speak the same language.'

By making sense of the complexity of hybrid cloud environments AIOps is ideally positioned to be the default solution for hybrid cloud monitoring. As you make the journey from on-premises to hybrid cloud, ensure you have a capable AIOps monitoring tool to rely on at every step.

Conclusion

Hybrid cloud presents a range of opportunities, but those come coupled with new challenges. To monitor hybrid environments effectively, IT teams must move beyond a reliance on tools that work only with a particular public cloud, and build a more integrated, cloud-agnostic monitoring tools set. At the same time, they must master new levels of complexity and take advantage of methodologies such as AIOps to interpret the complex, multi-layered metrics that hybrid environments generate.

For teams faced with these challenges, **AIOps from Broadcom** can help. By providing a cloud-agnostic, AIOps-enabled solution for monitoring and optimizing even workloads that span public clouds and private infrastructure, AIOps from Broadcom allows teams to thrive in the fast-changing hybrid cloud landscape.

Learn more by visiting **www.broadcom.com/aiops**

About Broadcom
Broadcom Inc. (NASDAQ: AVGO) is a global technology leader that designs, develops and supplies a broad range of semiconductor and infrastructure software solutions. Broadcom's category-leading product portfolio serves critical markets including data center, networking, enterprise software, broadband, wireless, storage and industrial. Our solutions include data center networking and storage, enterprise and mainframe software focused on automation, monitoring and security, smartphone components, telecoms and factory automation. For more information, go to www.broadcom.com.

0121