Monitoring Public Cloud Infrastructure and Services

Key requirements for optimizing performance,
value and the user experience

ca
technologies

Table of Contents

Executive Summary

Lured by opportunities to speed innovation, reduce costs and enhance agility, many business executives are opting to move their organization's applications into public cloud environments. Whether or not organizations realize these advantages to the fullest extent possible will in part be dictated by the monitoring capabilities in place. This paper offers a detailed look at the monitoring challenges cloud environments can pose, and it examines the key approaches organizations need to take to foster maximum service levels, efficiency and agility in the cloud.

Introduction

As they look to compete in the application economy, many organizations are growing increasingly reliant upon public cloud services. It's now practical for organizations to reliably leverage cloud offerings for a range of infrastructure services, including computing, storage and applications. Following are a couple of the reasons cloud services have emerged as such compelling alternatives.

Agility to fuel innovation

"Customers and end users need more robust features and a better experience, and they want it now, though they would have preferred it yesterday."

In the application economy, application innovation can't come too quickly. Customers and end users need more robust features and a better experience, and they want it now, though they would have preferred it yesterday.

As business leaders seek to speed their time to market for innovative new services, they have to move away from their traditional IT approaches. When they need to expand capacity, they can't afford to wait for lengthy procurement, testing and deployment cycles.

By running workloads in public cloud environments, organizations can realize significant gains in agility. They can much more quickly and easily expand their capacity, so they can immediately adapt to support new applications and services. If demand drops, they can also instantly reduce their footprint accordingly.

New servers or services can be provisioned on demand, which means staff time isn't eaten up doing procurement and installation. By leveraging cloud environments, internal staff can spend more time and effort on bringing new application-centric innovations to market—and they can get these innovations to market faster.

Operational efficiency and budgetary flexibility

As they compete in increasingly dynamic markets, organizations face heightened pressure to use resources in the most efficient manner possible. They can't afford the periodic, often massive outlay of capital expenditures that's needed to expand their infrastructure capacity. Nor can they afford the cost and disruption associated with re-allocating infrastructure resources to address changing market conditions.

ca
technologies

Cloud-based infrastructure services offer a compelling way to alleviate this challenge. By leveraging on-demand, elastic cloud services, organizations can move away from the massive upfront capital expenditures associated with on-premises infrastructure, and instead adopt pay-as-you go services that come out of operational budgets. This also allows organizations to reprioritize IT resources as business objectives change.

Key Monitoring Challenges for Public Clouds

As organizations move increasingly critical business services to public cloud environments, ensuring customers and users receive an optimized experience is vital. However, tracking progress and managing performance and service levels when running applications in the public cloud can present a number of challenges for IT organizations.

Lack of holistic, actionable insights into cloud services and the processes running on them

When organizations move workloads into cloud environments, IT teams often start by leveraging monitoring tools from their cloud service providers. Typically, cloud vendors' monitoring offerings aren't purpose-built monitoring solutions, however. These tools lack many of the capabilities administrators need to proactively manage performance and service levels of workloads and processes running in the cloud.

When employing a public cloud service, an organization will usually sign up for specific infrastructure offerings, such as computing, storage and so on. The cloud provider's tool would only deliver monitoring metrics for these resources, resulting in very narrow visibility. For example, a tool would tell you that CPU utilization is 60 percent, but it wouldn't provide any details about which processes or services are responsible for this utilization rate.

Further, many of these tools limit the extent to which IT teams can do historical performance analysis and future capacity planning. These tools typically only provide rudimentary reporting, and they only enable customers to retain monitoring data for a short period of time. Consequently, a retailer relying on monitoring tools from a cloud service provider would have a difficult time trying to track seasonal trends because they'd have no way of assessing year-over-year statistics.

Fundamentally, tools from cloud vendors lack capabilities for effectively monitoring service levels. As a result, it is hard to effectively measure the performance of a cloud service against SLAs, and to hold the provider accountable for the service levels delivered.

Limited insights into migration lifecycles

As organizations move applications and workloads to the cloud, they need to ensure that these migrations happen reliably. To do so, it is important to effectively track performance of workloads in development and production. By doing so, staff can most effectively ensure that no errors or performance issues arise. Eventually, they need to able to compare pre and post-production performance metrics so they can continue to optimize service levels and realize maximum benefits from the cloud.

Figure A.

To manage their cloud and hybrid IT environments, many organizations have relied on multiple point monitoring tools
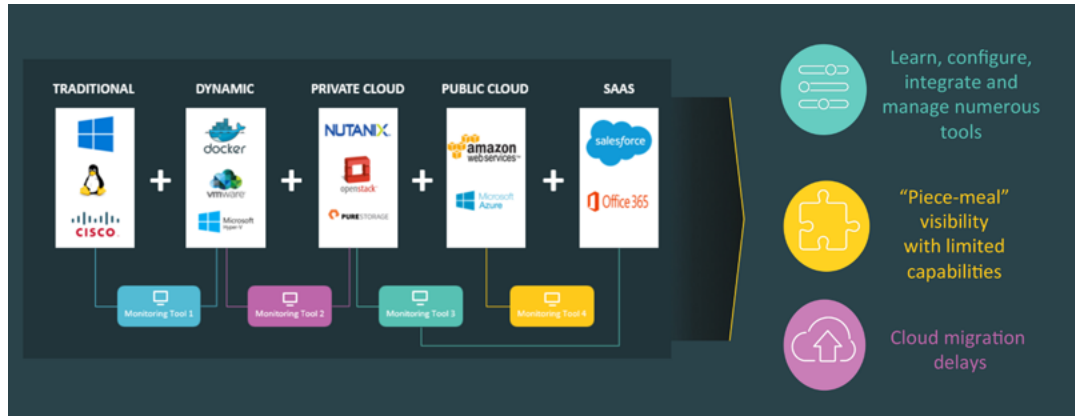
Exacerbated complexity of multiple monitoring tools for cloud and hybrid IT environments

In spite of the rapid growth in cloud services adoption, the reality is that most organizations today are taking a hybrid approach, running some workloads in the cloud and others in internal on-premises environments. When enterprise IT teams rely on cloud vendor-specific monitoring technologies, the reality is that the number of tools and associated challenges will increase. IT teams will need to continue to use their existing on-premises tools, which they've already invested time in procuring, configuring and managing. The cloud-specific tools will represent additional elements that staff needs to learn, use and manage.

Exacerbating matters further is the fact that most organizations will ultimately leverage cloud services from multiple cloud vendors. The reality is that each application has its distinct set of requirements, and each cloud service provider's offering will have its unique strengths and weaknesses. Consequently, for each cloud service, an organization would potentially have to add a distinct monitoring tool to the mix. The end result is that an organization may end up needing to work with monitoring tools from several cloud providers, in addition to the dozens of tools they may already have in place to monitor their on-premises technologies.

As the number of tools employed proliferates, so do the administrative burdens and costs. Further, when an issue arises, staff will have to spend a lot of time moving from tool to tool and participating in cross-functional team meetings in order to isolate where in the distributed environment the cause originated.

Limited utilization insights

When leveraging cloud services, businesses pay for capacity as they go. However, with cumbersome, limited monitoring tools, it can be difficult to track and fully understand current and ongoing resource utilization. As a result, organizations run the risk of spending on capacity they don't need, which diminishes some of the potential returns that can be realized by moving to the cloud. In addition, organizations need to analyze historical data to better plan for future capacity and budgets and to provide infrastructure perspectives to development teams so they can improve application performance.

Lack of end-to-end user experience monitoring

With limited point tools, IT teams lack insights into the one aspect that matters the most: the quality of the end user's experience. Point tools deliver monitoring of specific infrastructure elements but they don't provide capabilities for tracking performance and availability from the user's perspective, or for measuring the end-to-end response times of a transaction that spans a number of distributed infrastructures and services.

In the end, IT teams may see that the monitoring metrics available indicate that various infrastructure elements are performing well, but users may still be receiving slow performance or they may not be able to complete the transactions they need at all. This lack of visibility leaves businesses exposed to lengthy and costly performance and availability issues.

Key Ingredients for Optimizing Public Cloud Infrastructure Performance

Gain actionable insights for tracking performance

Quite often, point tools or tools from cloud vendors will provide a lot of metrics, but they are difficult to understand, provide piece-meal insights and are not normalized. They'll lack actionable insights needed to optimize performance and resolve issues faster. Be sure to use actionable metrics that provide real insights into your cloud infrastructure and how it's being utilized.

Figure B.

Monitoring tools should provide unified visibility of cloud infrastructure as well as services running in these environments
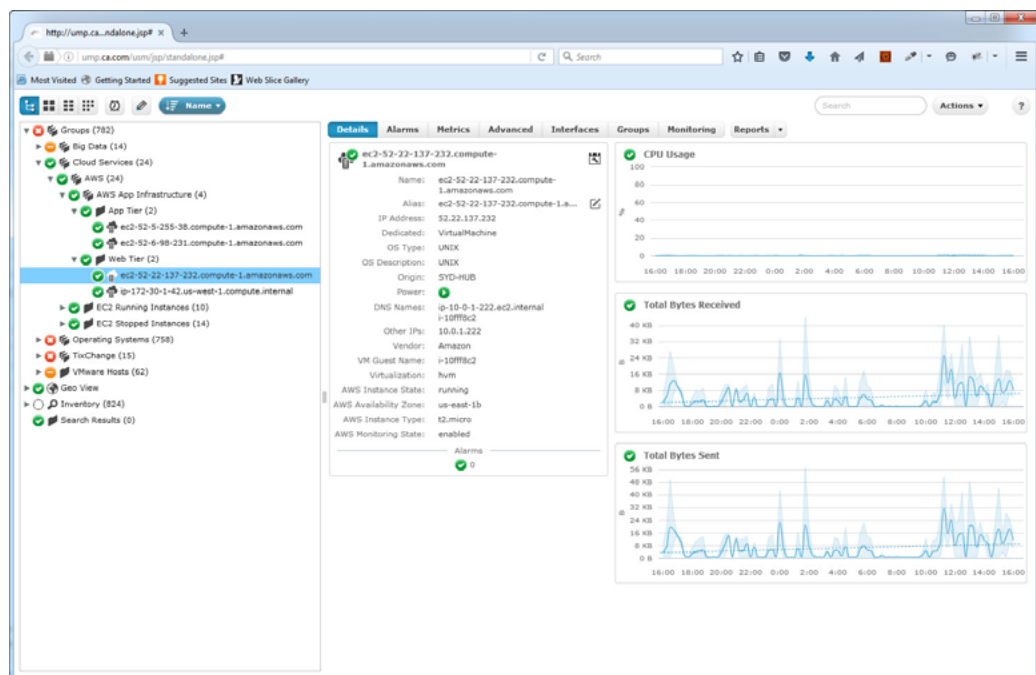
Monitor application and services running in the cloud to detect issues faster

When an organization signs up for a public cloud service, they're often subscribing to specific set of systems or infrastructure services, such as servers, storage, databases and so on. While tools from cloud service providers will provide monitoring details for these infrastructure elements, they don't offer any visibility into the performance of the applications or processes running on top of these elements. These tools might indicate CPU utilization is 60 percent, but not which process is responsible for the workload. Teams need this visibility to understand and optimize performance in a holistic fashion.

### Establish a unified view of cloud and on-premises infrastructure to speed mean time to repair

The reality is that most organizations are running workloads on a mix of on-premises and cloud-based environments. Having a unified view when managing an internally sourced IT environment is critical. While public cloud services offer significant differences in terms of the degree of hands-on control over servers and other infrastructure elements, the reality is that having a unified view is nevertheless just as vital. When IT teams get a holistic view of all their infrastructure capacity and utilization, they can resolve issues faster.

Track end-to-end user experience to ensure service reliability

Today, when a banking customer logs into a mobile application to check her balance, a number of distributed environments may support the interaction. The user may log into a Web server hosted in one public cloud provider's environment, then the credentials may be verified by an on-premises server and then a call may be made to a back-end system housed in a second cloud provider's facility in order to access account details. It's critical that IT teams can track these multistep transactions from end to end and measure exactly what kind of experience the user is receiving.

Figure C.

IT teams need holistic views of instances running in cloud environments
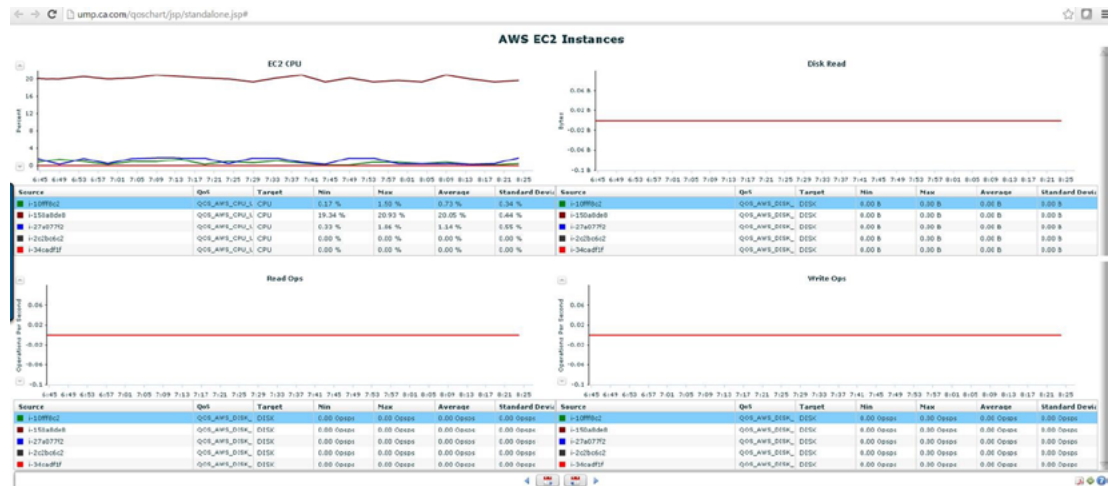
Harness intelligent alerting that keeps teams informed—not overwhelmed

Cloud environments are highly elastic, with computing resources being turned up and turned down on a continuous basis, which can wreak havoc for teams relying on basic monitoring tools.

IT teams need to be able to set intelligent, dynamically updated thresholds so they can ensure they get apprised of critical issues when they need to be—without getting deluged by false and redundant alarms. By leveraging sophisticated time-over-threshold analytics, IT teams can identify real, persistent performance issues and eliminate false alarms associated with occasional spikes. Time-to-threshold analytics can automatically identify threats of potential performance degradation and issue an early warning so administrators can respond—before internal and external users are affected.

Advanced platforms can provide a prioritized list of problems that represent situations administrators should watch. They can also place these items in order of time remaining before a problem is expected to occur, bringing attention to immediate problems first. In addition, monitoring platforms need to offer integration with service desk solutions, so that, when an issue arises, IT teams can automatically have a service ticket opened that includes all relevant details on the issue.

Establish rapid, template-based monitoring deployment

One of the principal advantages of cloud environments is their agility, however, the dynamic, elastic nature of cloud environments also presents IT teams with some of the most difficult challenges. As virtualized, on-demand resources are constantly being turned up and turned down to accommodate changing workloads, monitoring also needs to be started and stopped on a continuous basis.

To establish effective and efficient monitoring of these environments, IT teams need to minimize or eliminate manual efforts wherever possible. To realize these objectives, IT teams need to establish templates for monitoring specific categories of technologies, and leverage automation to the fullest extent possible to apply these templates.

## Leverage monitoring throughout the migration lifecycle

As organizations migrate applications off on-premises infrastructures and into cloud environments, monitoring will need to be employed in order to ensure that no disruptions in service levels are introduced. To most effectively manage service levels over the course of the transition, IT teams should employ the following approaches:

- Establish monitoring of services immediately as they are migrated into a new cloud environment. This is vital, given this is the point in which configuration and other migration-related issues are most likely to arise.

- Develop dashboards and reports that provide granular, side-by-side comparisons of prior implementations and new deployments, which is vital in spotting trends that may indicate the emergence of performance and availability issues.

- Conduct synthetic transaction monitoring of services, both as they run on original on-premises implementations and after they are migrated to cloud environments. These monitoring metrics provide an invaluable, before-and-after comparison of end-user service levels, so IT teams can most effectively ensure a cloud migration doesn't introduce any degradation of performance.

Figure D.

To manage their cloud and hybrid IT environments, many organizations have relied on multiple point monitoring tools
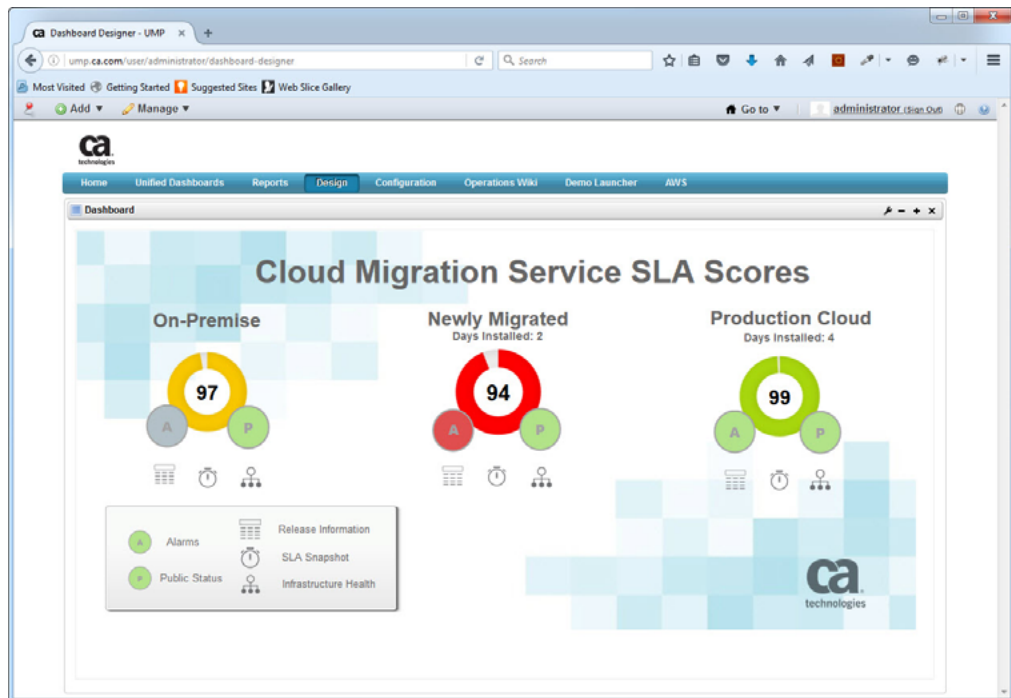
Proactively monitor cloud utilization

Many decision makers ultimately choose to move to cloud environments precisely because of the cost savings and budgetary flexibility that can be gained. However, once the move into the cloud is made, IT teams need to proactively manage capacity in order to ensure maximum cost savings are realized, both initially and over the long term. To achieve these objectives, look to institute:

- **Historical trending reports and dashboards**. With these views, analysts can track longer-term utilization trends and evaluate year-over-year comparisons. These views are particularly vital for organizations that have to accommodate seasonal and dramatically fluctuating usage trends.

- **Intelligent thresholds for cost and utilization data**. By establishing intelligent, time-to-threshold alerting, IT teams can ensure they are notified before resources are over provisioned, so they can more proactively manage service levels and capacity. Further, through intelligent thresholds, IT teams can constantly ensure they are apprised of under-utilized resources, so they can more proactively reduce ongoing expenses.

Conclusion

While the benefits of cloud implementations can be massive, these returns can be elusive for many organizations. Monitoring capabilities will play a vital role in whether an organization can realize maximum business benefits from their cloud deployments. It is only through powerful, unified monitoring that organizations can effectively track and manage the level of service users receive, and enable the business to most fully capitalize on the cost savings and agility that cloud offerings can provide.

Today, many organizations leverage CA Unified Infrastructure Management (CA UIM) to optimize the service levels and business benefits of their entire IT ecosystem, including cloud-based and on-premises environments. To learn more about CA UIM, visit **ca.com/cloud-monitoring**.

Connect with CA Technologies at ca.com

CA Technologies (NASDAQ: CA) creates software that fuels transformation for companies and enables them to seize the opportunities of the application economy. Software is at the heart of every business, in every industry. From planning to development to management and security, CA is working with companies worldwide to change the way we live, transact and communicate—across mobile, private and public cloud, distributed and mainframe environments. Learn more at **ca.com**.