Top 6 Myths

of Cloud Performance Monitoring

An eG Innovations Technical White Paper

www.eginnovations.com

eG Innovations

Top 6 Myths of Cloud Performance Monitoring

Cloud adoption is increasing at a rapid pace. IDC estimates that global end-user spending on cloud services will exceed $1.3 trillion within the next five years. As organizations increasingly rely on cloud services to operate their business, it is important to monitor and manage the performance of applications deployed on the cloud. This is where cloud performance monitoring comes in.

What is Cloud Performance Monitoring?

Cloud performance monitoring (cloud APM) is the process of tracking the performance of applications delivered from the cloud.

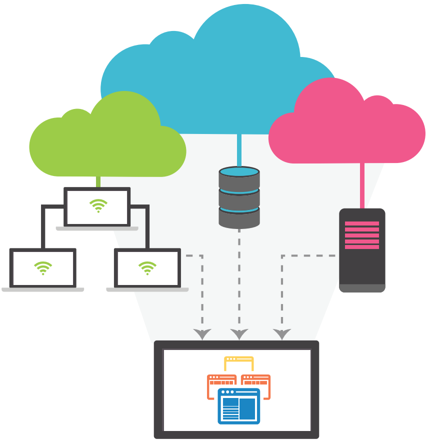

- User experience metrics are important for cloud performance monitoring as they provide an indication of how well an application is performing.
- At the same time, detailed application and infrastructure metrics are required if performance slowdowns are detected and additional diagnosis is necessary for troubleshooting.
- There are hundreds of cloud services offered by service providers like AWS, Microsoft Azure, and others. Cloud performance monitoring must collect and analyze key performance indicators for all the services used by an application – e.g., if an EC2 instance is used to host an application, is the CPU credit balance of the instance sufficient? Or, could the application be slowing down because of insufficient CPU credit balance? (Read more about why it is important to track cloud services specific metrics like CPU credit balance when monitoring applications hosted on AWS cloud in our blog post).
- Infrastructure metrics in a cloud environment usually focus on usage. At the same time, they should also include indicators that highlight whether the resources allocated are under-provisioned.
- The monitored metrics are analyzed and correlated with each other (using auto-discovered dependency mappings) to determine where the root cause of a problem lies.
- Cloud performance monitoring also includes the ability to take actions to resolve issues.

Cloud Performance Monitoring Myths

Increased agility, OPEX purchases, scaling, high availability, and outsourced managed services are some of the common reasons why organizations are moving workloads to the cloud. At the same time, there are many who believe that when applications are migrated to the cloud, it is easier to manage them and that the cloud service provider tools will help them greatly in this regard. Is this a myth or a reality?

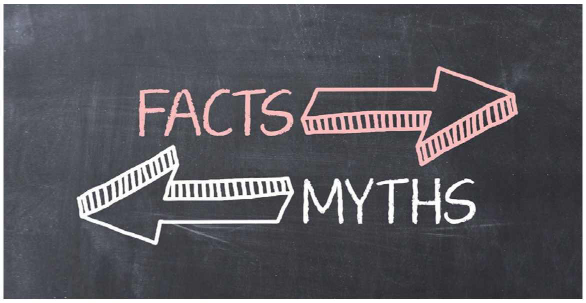

In this white paper, we discuss the six common myths about cloud performance monitoring and debunk them:

1. Migrating applications to the cloud removes the need for any performance monitoring.
2. Cloud environments support auto scaling. If I use auto scaling in cloud environments, application performance is guaranteed, and monitoring is not needed.
3. Moving to the cloud guarantees high availability because the cloud never goes down.
4. Cloud-native monitoring tools have all the required capabilities to monitor applications and infrastructure.
5. The same monitoring technologies used on-premises can be used in the cloud as well.
6. Monitoring application performance in the cloud is a lot easier than in an on-premises infrastructure.

MYTH #1 Migrating applications to the cloud removes the need for any performance monitoring.

Figure 1 below compares an on-premises deployment to an IaaS or PaaS deployment in the cloud. While the cloud service provider provides the platform that hosts the application, the application code is still owned and operated by you, the cloud customer. If there are issues in the application code (e.g., runaway thread, inefficient query, etc.), these issues still need to be detected and diagnosed. Performance monitoring tools are needed for this.

Changes to the underlying cloud infrastructure can also affect application performance. For example, a hot fix applied on an AWS EC2 instance can be detrimental to application performance. For fast diagnosis of the problem, you must be able to correlate that the performance issue started right around the time when the hot fix was applied.

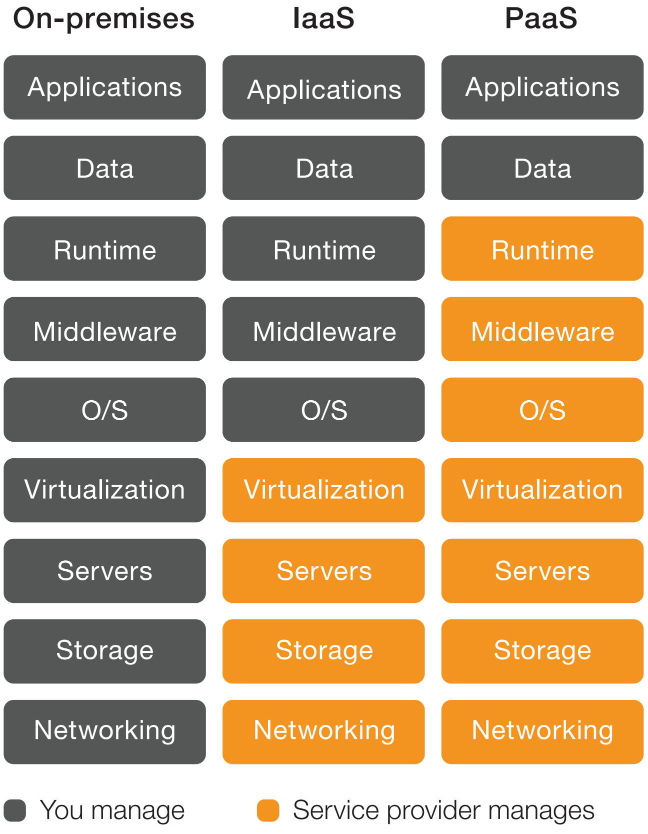

On-premises	IaaS	PaaS
Applications	Applications	Applications
Data	Data	Data
Runtime	Runtime	Runtime
Middleware	Middleware	Middleware
O/S	O/S	O/S
Virtualization	Virtualization	Virtualization
Servers	Servers	Servers
Storage	Storage	Storage
Networking	Networking	Networking

Figure 1: On-premises vs. IaaS vs. PaaS deployment

Performance monitoring is also essential to highlight whether the cloud infrastructure configured is sufficient for the workload of the application. For example, the CPU or memory configuration of an EC2 instance may be insufficient and you need performance monitoring to be able to recommend that you reconfigure the EC2 instance for improved performance.

Even if you use a managed service like AWS RDS, the availability of the RDS instance is the cloud provider's responsibility, but the performance of the instance may depend on what queries are executed on it, how fragmented the indexes are, whether queries are tuned or not, etc. All these factors are the responsibility of the customer, not the cloud provider.

Therefore, moving to the cloud doesn't guarantee all the applications will function without failure. And when failures happen, it is necessary to monitor logs, traces and metrics to determine where problems lie and fix them. Having the right monitoring in place can help you understand why a slowdown happened: is it due a recent configuration change, or due to excessive workload, or could it even relate to security attacks (e.g., too many invalid logon attempts)?

Also note that cloud applications may also rely on external, third-party services beyond your own cloud infrastructure. For example, eCommerce applications leverage payment gateways from suppliers such as Visa or Mastercard. A slowdown leading to abandoned shopping carts and transactions. Performance monitoring is needed to identify such issues so you can get them rectified at the earliest opportunity.

Fact: Even after moving to the cloud, performance monitoring is required to ensure the optimal functioning of the application and to assist with diagnosis when problems occur.

MYTH #2 If I use auto scaling in cloud environments, application performance is guaranteed, and monitoring is not needed.

Auto scaling is the process of automatically increasing or decreasing the computational resources delivered to an application based on demand. The primary benefit of auto scaling is that your workload gets exactly the cloud computational resources it requires (and no more or no less) at any given time. You pay only for the resources you need when you need them. Most popular cloud providers offer auto scaling as a service.

Performance monitoring is needed to make sure the application and the supporting infrastructure are performing well. It can also keep a tab on how auto scaling is working and can alert you if scale-outs are repeatedly happening. With performance monitoring in place, you can get insights for you to understand why a scale-out happened.

Fact: Performance monitoring is required to keep tabs of how auto scaling is working and to alert IT operations teams if excessive scale-outs are happening or if a scale-out is triggered by a problem with an application.

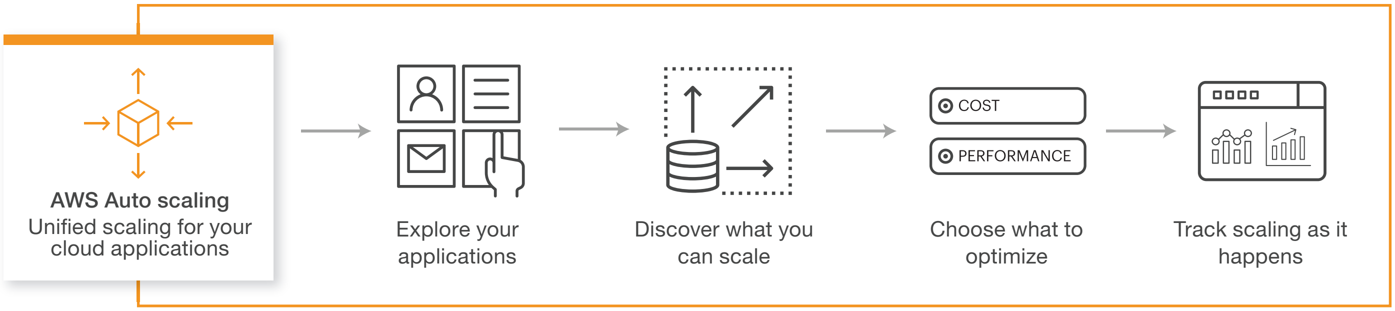
AWS Auto scaling
Unified scaling for your cloud applications
Explore your applications
Discover what you can scale
COST
PERFORMANCE
Choose what to optimize
Track scaling as it happens

Figure 2: How AWS Auto scaling works

While auto scaling helps to save cost and efficiently use resources, scale-in and scale-out is based on a set of metrics. Sometimes a malfunction of an application may make it appear as if it is under pressure. If scale out keeps happening, you could be incurring a high cost which is due to an issue in your application.

> *Technologies such as Kubernetes and containers are inherently architected to auto scale. 48% of respondents to a recent survey reported that they were using containers in the cloud. This is 50% higher than the number of organizations deploying containers on-premises.*
>
> **Source:** eG Innovations and DevOps Institute APM Survey 2021

MYTH #3 Moving to the cloud guarantees high availability because the cloud never goes down.

Cloud providers market high availability for their services. So, there are many who believe that by just moving to the cloud, they can make their applications highly available.

It is true that cloud service providers have invested a lot of technology and resources to ensure that their key services are highly available, especially as their reputation is on the line.

- Moving to the cloud will also protect you from normal hardware failures and other unexpected outages within an availability zone (Amazon) or fault domain (Azure) because the cloud provider has in-built mechanisms to withstand such failures.

- You can also choose to use a managed service like Amazon RDS or Azure Database to reduce the burden on your team and minimize the chances of a failure impacting your applications

Failure of your application – e.g., a crash, may also impact application availability. The cloud service provider is not responsible for any faults in your application logic.

While migration to the cloud usually offers better resilience than an on-premises deployment, it does not guarantee high availability. Like any other large organization, cloud providers also experience downtime and regional disasters.

- To ensure high availability, you must design your application so that its architecture does not have any single point of failure. Failover mechanisms for both stateless and stateful components must be considered. Setting up application components in a cluster with load balancing can be considered.
- To protect against regional disasters, you should consider infrastructure design patterns that ensure high availability, redundancy, and failover. For instance, AWS or Azure offer geo-redundant storage to ensure high availability even if a regional outage occurs. They allow for geo-redundant replication to a secondary region that is hundreds of miles away. These services not only cost you money, but they must be considered when you design and deploy your application in the cloud.

Cloud providers too have outages

Azure experienced a six-hour outage in March 2020 due to cooling system failures, which hampered the performance of network devices, rendering compute and storage instances inaccessible.

Amazon suffered three outages in December 2021 that affected everything from airline reservations and auto dealerships to payment apps and video streaming services to Amazon's own e-commerce operation. The issues were blamed on power outages in a single data center, in the US-EAST-1 Region. Customers with critical systems only available in that region were especially affected.

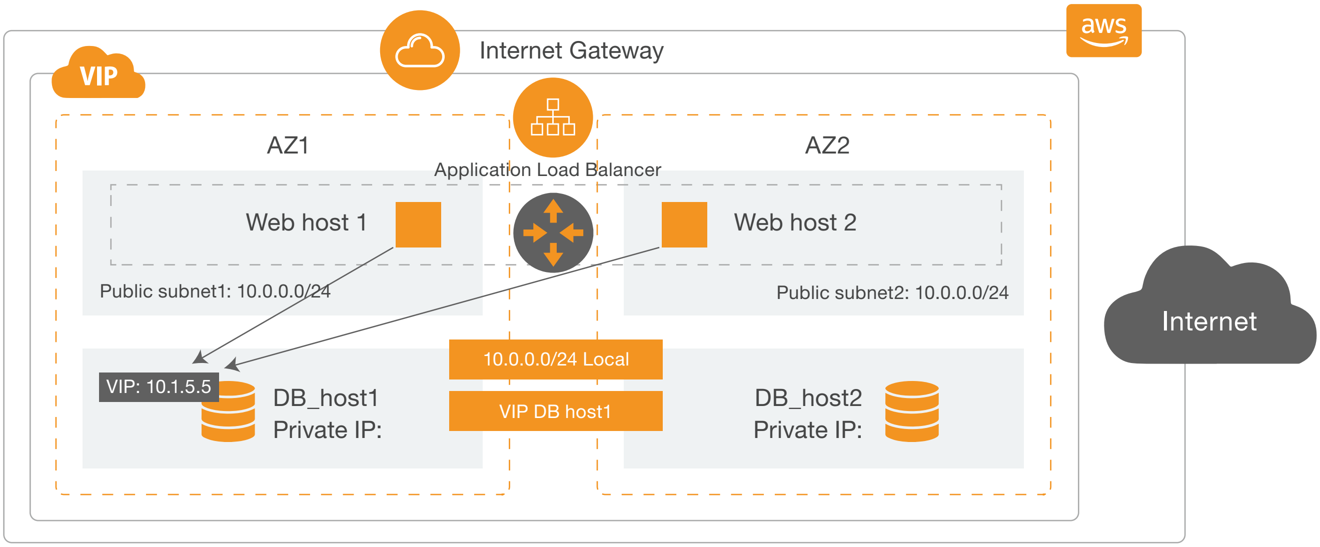

Figure 3: An application spanning multiple availability zones

Fact: Just deploying your applications in the cloud does not guarantee high availability. You need to incorporate additional mechanisms to ensure high availability. Depending on the criticality of the application, your SLAs, and your cost budget, you can choose the appropriate HA configuration to ensure high availability.

MYTH #4 Cloud-native monitoring tools have all the required capabilities you need for performance monitoring.

Most cloud environments have some level of monitoring built in. AWS has CloudWatch to monitor your cloud services while Microsoft Azure has Azure Monitor. Many organizations migrating to the cloud are under the impression that the built-in cloud native tool has all the capabilities you might need for performance monitoring in the cloud.

AWS CloudWatch

Azure Monitor

Google Cloud Monitor

71% In a recent survey of IT professionals by eG Innovations and DevOps Institute, 71% of respondents indicated that the built-in cloud monitoring tools were not sufficient for their needs.

Those who have sufficient experience of deploying and maintaining applications in the cloud know that this is far from the truth. Cloud-native monitoring tools like AWS CloudWatch have several disadvantages when compared with specialized monitoring tools like eG v:

- Deploying the cloud-native tool is often time-consuming. For example, AWS has a CloudWatch agent that can be used to monitor the operating system of your EC2 instances. However, setting up the monitoring is an elaborate process. Setting up the rules and configuring the metrics and any associated alert thresholds that you need collected and monitor is a manual process. There are no pre-defined templates or models that make your job simpler.

25% of users of the top 3 cloud providers feel that the functionality of the native monitoring services is basic.

Source: eG Innovations DevOps Institute Cloud survey

- The pricing of cloud-native monitoring tools is also not straightforward. Every metric you collect costs money. The more frequent the metric collection, the greater the cost. Even applying thresholds for metrics can cost you money! It is, therefore, no surprise that 30% of organizations using AWS see AWS CloudWatch as being expensive (Source: eG Innovations and DevOps Institute APM Survey). Third-party monitoring tools like eG v are simple to set up and their licensing is straightforward and not based on the number of metrics or frequency of metrics collection.

32% of users see gaps in capabilities with the cloud provider's monitoring service that they need to fill with 3rd party solutions.

Source: eG Innovations DevOps Institute Cloud survey

- While tools like AWS CloudWatch provide sufficient details into the utilization and performance of AWS services, the depth of insights into applications deployed on the cloud can be limited. 38% of organizations using AWS indicated that there are gaps in AWS CloudWatch's monitoring capabilities (Source: eG Innovations and DevOps Institute APM Survey).

Fact: Built-in cloud monitoring tools require elaborate set-up and are time-consuming to operate. They also lack all the capabilities of a complete application and infrastructure monitoring tool.

	AWS CloudWatch	eG Enterprise Monitoring Solution
Monitoring of key AWS services	Pick and choose what metrics to collect. Cost is based on API calls made.	Integrates with AWS CloudWatch. Can control the frequency of polling to reduce cost of monitoring.
Synthetic monitoring of user experience	Supported using AWS Synthetics.	Supports synthetic monitoring for web applications, client server, thin client applications, etc.
Detailed insights into EC2 instances, databases, applications, etc.	Possible but separate tools are needed for each. AWS X-ray is needed for application transaction tracing, AWS Service Lens is needed to integrate traces with logs and metrics. AWS CloudWatch only provides basic metrics for database monitoring. AWS Database Performance Insights is needed for in-depth monitoring.	Converged application and infrastructure monitoring tools like eG Enterprise provide a single pane of glass to monitor user experience, trace transactions, deep dive into application components and using this detail, determine the root cause of performance issues. Additional add-on products are not required.
Monitoring of other cloud environments, on-premises environments	Only supports AWS services. Separate monitoring needed for other cloud providers. Not straightforward to monitor.	Supports monitoring of multiple cloud environments and hybrid cloud environments from a common console.
Customized dashboards and reports tailored for different stakeholders	Have to spend time and effort creating dashboards and reports. Time consuming.	Has extensive pre-defined and readymade, domain-specific dashboards and reports tailored for the different stakeholders.
Ease of deployment	Basic metrics are easy to obtain. There is no template for advanced metrics. So, administrators have to spend time picking and choosing what they want to monitor.	Auto-discovers what to monitor. Pre-defined templates specify what to monitor for each cloud service and application in the cloud. Very little manual/tedious work is required for deployment and configuration.
Cost of monitoring	The overall cost of monitoring is difficult to predict. The cost varies based on metrics collected – more the metrics, greater the cost. Cost is also based on frequency of monitoring – greater the frequency, higher the cost. Type of baselining mechanism also changes the cost.	Monitoring cost is straightforward to compute. Licensing is based on number of target application/OS to be monitored. Not licensed by amount of metrics collected, frequency of monitoring, etc.
Root cause diagnosis and analytics	Often left to other tools that take feeds from these tools.	All-in-one solution includes AIOps and data analytics capabilities to pinpoint the root cause of problems. Supports trend analysis and forecasting of metrics.

Table 1: How built-in cloud monitoring tools like AWS CloudWatch compare with specialized monitoring tools like eG Enterprise

MYTH #5 The same monitoring technologies used on-premises can be used in the cloud as well.

Most organizations are looking to reduce the toolsets they have to use. Furthermore, every time a new tool is introduced, there is a learning curve for the IT operations team. As a result, one of the first things that organizations look to do in the cloud is to use as much of the tooling they are using on-premises in the cloud as well.

To some extent, this approach works. There may be some communication restrictions between applications in the cloud, but beyond that the on-premises tools that monitor the application stack can be used in the cloud as well, as long as the right permissions are configured. For example, you can use the same database monitoring tool to monitor the health of your Oracle or Microsoft SQL database instances in the cloud.

There are several other factors to consider as well:

- Complexity arises when you are using cloud-native technologies. For instance, if you are using a cloud-native database like AWS DynamoDB, the on-premises tool may not be able to monitor it.
- Many cloud-native applications make extensive use of capabilities like auto scaling. Monitoring tools need to be aware of such technologies and must track all such activities. Furthermore, the deployment of the monitoring tools for cloud environments must be automatic. Configuration of the monitoring also must be automated and should be possible to do without any human intervention. Tools designed for on-premises environments do not support cloud-native capabilities like auto scaling and may not be designed for automated deployment.
- Security is another key consideration. Many on-premises tools can operate with lower levels of security. Several on-premises tools employ an agentless monitoring technique: a central data collector connects to the target systems to obtain performance metrics. Often, this creates a central point of vulnerability – one system that can access all other systems. Agentless monitoring in the cloud must be configured carefully so that it does not compromise the security of the systems and applications being monitored. If agent-based monitoring is used, the agents on systems in the cloud should not listen in on TCP ports as open TCP ports can often provide a foothold for security attacks.
- Another key requirement for cloud environments is monitoring of utilization levels and billing. If you are using burstable instances and you are continuously exceeding the usage limit, the cloud provider may throttle the workload. Hence, application users may experience slowness. Tools designed for on-premises environments cannot track utilization against permissible limits and alert in advance.
- While over-provisioning in on-premises environments does not have a severe impact, the same in a cloud environment can cost your organization a lot of money. Hence, capabilities such as detecting instance sprawl, underutilization of provisioned resources, etc. are more important in a cloud environment.

Fact: It may not be possible to directly reuse your on-premises monitoring tool in a cloud environment. Look for monitoring tools that have capabilities that are well integrated with cloud environments.

MYTH #6 Monitoring application performance in the cloud is a lot easier than on-premises.

One of the key challenges you will encounter in a cloud environment is a lack of complete access to and visibility of the infrastructure. Even if you are using just the compute platform like AWS EC2, you only have access to the compute instance. You cannot see which hypervisor it is hosted on and what the performance of the hypervisor is.

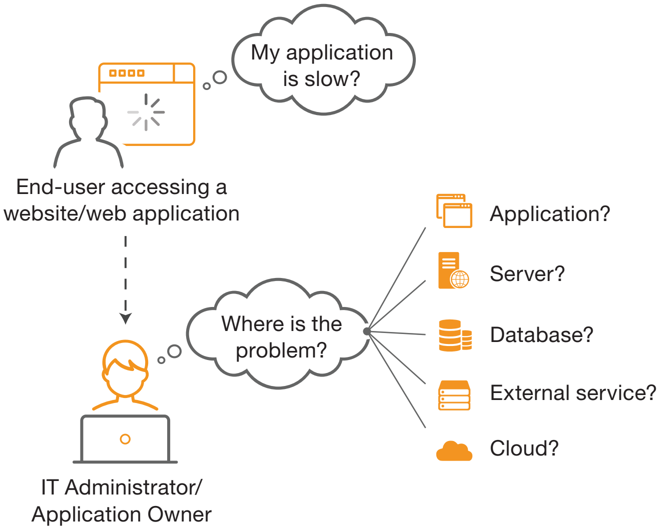

Figure 4: Diagnosing why an application in the cloud is slow is an IT admin's toughest challenge

Cloud environments are often multi-tenant and share resources across tenants. You will not be able to see the usage of other tenants. So, there could be situations where one tenant's usage of resources can affect the resources available to you. Without complete insight into the underlying infrastructure, it is not always possible to prove that any performance issues you are seeing are because of the cloud provider's infrastructure.

The separation of domains of control and restricted access make performance problems harder to diagnose in a cloud infrastructure. Often, you have to rule out that the problems are not in your application and its components, before you can point fingers at the cloud service provider.

In an on-premises infrastructure, it is easier to get complete visibility – of the application stack and the underlying infrastructure, from one console. Hence, diagnosing performance issues is easier.

Fact: Cloud applications are often harder to diagnose and troubleshoot than on-premises applications.

Conclusion

Many organizations are adopting cloud computing for its many benefits. At the same time, many are not aware of the challenges of cloud computing. As we discussed in this white paper, because of its dynamic and multi-domain characteristics, cloud computing does make secure monitoring and diagnosis of application availability and performance more challenging.

In this white paper, we have debunked the 6 common myths regarding application deployments in the cloud. Organizations that are well prepared as they embark on their cloud journey are likely to be more successful than ones that are not.

Learn More

Choosing the right observability and monitoring tool is key to your short and long-term business agility and cloud operations success. For more details, see:

- AWS Monitoring with eG Enterprise - https://www.eginnovations.com/aws-monitoring
- Azure Monitoring with eG Enterprise - https://www.eginnovations.com/azure-monitoring

Next Steps

✉ | To contact eG Innovations sales team: sales@eginnovations.com

🌐 | Get a free trial of eG Enterprise: www.eginnovations.com/FreeTrial

✉ | For support queries and feature requests: support@eginnovations.com

About eG Innovations

eG Innovations provides the world's leading enterprise-class performance management solution that enables organizations to reliably deliver mission-critical business services across complex cloud, virtual, and physical IT environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations' award-winning solutions are trusted by the world's most demanding companies to ensure end user productivity, deliver return on transformational IT investments, and keep business services up and running. Customers include Anthem, Humana, Staples, T-Mobile, Cox Communications, eBay, Denver Health, AXA, Aviva, Southern California Edison, Samsung, and many more. To learn more visit www.eginnovations.com.

References

AIOps Solutions and Strategies for IT Management | eG Innovations

AIOps Tools – 8 Proactive Monitoring Tips | eG Innovations

Service and Help Desk Automation Strategies | eG Innovations