Top 10 Requirements for Performance Monitoring of Cloud Applications and Infrastructures

An eG Innovations Technical White Paper

www.eginnovations.com

eG Innovations

Introduction

Cloud is no longer a buzzword. Organizations of all sizes are adopting cloud technologies at a rapid pace. IDC forecasts that worldwide spending on cloud technologies will surpass $1.3 trillion by 2025, with a compounded annual growth rate of 16.9%.

Cloud adoption is being driven by the need to be agile, the desire to consume a resource only when it is needed, i.e., on-demand, and by the ease of management – i.e., no need to deploy software components and manage and patch them yourself.

At the same time, everything is not simple with the cloud:

- Cloud migration is a complex activity. While the focus is on ensuring that your applications are functional, **you also have to make sure that the performance of the application doesn't suffer as a result of the migration**. Performance of the application post-migration may depend on several factors including the cloud infrastructure configuration and performance.
- Once your applications are operational in the cloud, slowdowns could occur at any time. Patching of the systems, automatic updates, change in configuration, changes to the workload, etc. can all introduce slowness. **Your cloud provider is not responsible for the performance of your applications**. Their SLAs focus on system and service uptime, not on performance.

You will still be responsible for the performance of applications used by your employees and customers. If those applications are unavailable, unreliable, or deliver poor user experience, your business will suffer. Hence, you will be required to triage problems quickly, determine what caused an issue and resolve it.

81% of organizations have applications deployed in the cloud.

21% of organizations are exclusively using cloud technologies.

Source: eG Innovations and DevOps Institute APM Survey 2021

If you suspect that a slowdown is due to the cloud service provider, you will have to prove that it was their issue. In fact, because cloud environments involve multiple domains of control and limited visibility across domains (you do not have much insight into the cloud provider's infrastructure – e.g., a slow disk or an overcommitted hypervisor or a slow network connection), **performance monitoring and management is even more challenging in cloud environments than it was in on-premises infrastructures**.

In this white paper, we will explore the key requirements that organizations adopting cloud services are going to have for performance monitoring of their applications and infrastructure.

Adoption of Cloud Technologies - A Golden Chance to Rethink Monitoring Strategies

Adoption of cloud technologies is an inflection point in an organization's IT evolution. Most on-premises environments already have a plethora of monitoring tools. These tools may have been acquired over the years for a variety of reasons. While most organizations have wanted to move to a single pane of glass for efficiency and efficacy, the challenges have been:

- Support for legacy technologies
- Finding a common toolset that addressed all the needs of the organization
- Desire to get payback for all the investments already in place
- Getting acceptance across all the different IT teams for a common monitoring tool

Many legacy, on-premises monitoring tools use protocols like SNMP that do not work in cloud environments. Not only is the management protocol used for cloud services and components different, many cloud services have unique capabilities not seen in on-premises infrastructures. For example, services like S3 bucket storage, elastic load balancing, serverless computing, etc., on AWS cloud do not have a direct equivalent in an on-premises infrastructure. Auto-scaling is also a challenge – monitoring tools can no longer operate based on manual configurations done by IT admins.

Organizations using cloud services have an opportunity to start with a clean slate. They have an opportunity to understand what the key monitoring requirements will be as they adopt cloud services and frame their monitoring strategy accordingly.

In the following sections, we list out the key requirements and features that organizations will need as they adopt cloud services.

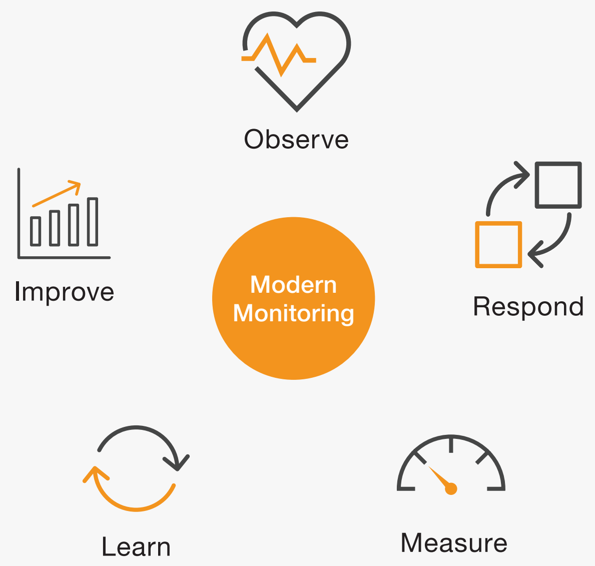

Different disciplines of modern monitoring (see https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/monitoring-strategy)

Top 10 Requirements of Cloud Monitoring Tools

1. Multi-cloud support
2. Support for hybrid cloud deployments
3. Monitor all the cloud services your organization uses
4. Monitor the usage and performance of cloud services with analytics for automatic problem detection
5. Monitor the digital user experience for different workloads
6. Provide clear demarcation of problems
7. Support cloud technologies like auto-scaling and dynamic microservices with auto-deployment, auto-discovery and auto-configuration capabilities
8. Ensure data security, compliance, and governance
9. Provide a predictable billing model for monitoring
10. Monitor cloud billing costs and provide cost optimization analysis

#1 Multi-Cloud Support

Numerous surveys demonstrate an increasing number of organizations are starting to adopt multiple cloud providers, and even if they have not yet, they will be considering multi-cloud in the future.

> 92% of enterprises have a multi-cloud strategy.
>
> Source: Flexera report - Cloud Computing Trends 2021 State of the Cloud Report

There are many reasons why organizations are looking at a multi-cloud strategy:

- A cloud provider may have strengths in one area, but not another. For instance, given Microsoft's strengths in the desktop, organizations may look at Microsoft Azure for hosting virtual desktops in the cloud.
- On the other hand, the breadth of database options it provides makes Amazon a desirable choice for database technologies.
- Familiarity with a cloud vendor and cost are also important factors in deciding which services are hosted by which service provider.

Many organizations prefer to be vendor agnostic and adopt a multi-cloud strategy to avoid being locked into any one cloud vendor for a variety of reasons including pricing.

Most cloud providers have tools that can be used to monitor services hosted in their infrastructure. AWS has CloudWatch. Azure has Azure Monitor and Log Analytics. Google has Operations Suite (formerly Stackdriver). These tools are specific to individual cloud providers and their services. AWS CloudWatch can monitor AWS' own EC2, ECS, EKS, and RDS services, but you cannot monitor these services from Microsoft Azure Monitor (and vice versa).

> Hence, a vendor-agnostic monitoring solution that provides multi-cloud support is important. Organizations then have the flexibility to choose the cloud services they want to monitor and manage.

Besides avoiding lock-in to one cloud service provider, a specialized monitoring tool allows organizations the flexibility to change service providers and to retain historica data on usage and trends across service providers.

#2 Support for Hybrid Cloud Deployments

While some organizations will move 100% to the public cloud, most are likely to have some applications in the cloud and some on-premises. There may even be select services that span both on-prem and cloud infrastructures. Over time, it is likely that some on-prem services will gradually move to cloud environments.

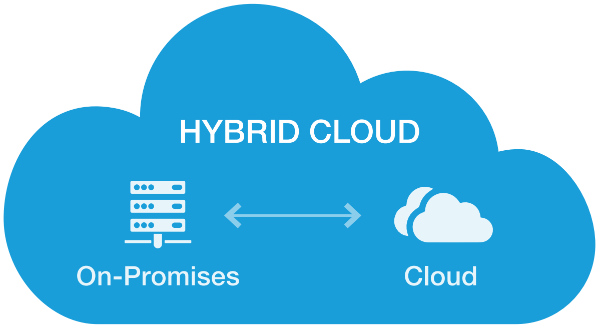

> 43% of organizations are looking at multiple public and multiple private clouds.
>
> Source: Flexera report - Cloud Computing Trends 2021 State of the Cloud Report

> *A hybrid cloud is a cloud computing environment that uses a mix of on-premises, private cloud and third-party, public cloud services with orchestration between these platforms. This typically involves a connection from an on-premises data center to a public cloud. The connection can also involve other private assets, including edge devices or other clouds.*
>
> **Sarah Neenan,** Editor, TechTarget
>
> Source: https://searchcloudcomputing.techtarget.com/definition/hybrid-cloud

Using one set of tools and workflows for cloud environments and another for on-premises can be cumbersome. There are learning curves and on-going maintenance costs for IT operations teams to learn and support different tools. The lack of consistency in reporting, alerting and diagnosis procedures when using multiple tools can become very problematic. Also, with multiple different toolsets, when a problem arises, diagnosis involves manual steps and is reliant on the subject matter experience of a few key individuals. This is not an efficient way to run IT operations and often prompts projects to evaluate and adopt common toolsets to monitor on-prem and cloud environments. So look for cloud monitoring tools that can support your on-premises applications and infrastructure as well.

By consolidating monitoring tools for on-premises and cloud environments, organizations can avoid having to retrain staff on multiple domain-specific tools. This also allows them to change infrastructure and services whilst retaining a consistent view and level of insight across different environments – a key requirement to measure the impact and effectiveness of change, especially during migrations.

Major Cloud Providers accept that Hybrid and Multi-Cloud is Inevitable

Major cloud providers (and traditional on-prem VDI vendors) recognize that to remain relevant they cannot ignore customers' needs to use alternative clouds or on-premises infrastructures and many provide services to facilitate this:

- **Amazon Web Services (AWS)** - AWS Outputs is a hybrid cloud service that brings Amazon EC2 to the data center.
- **Amazon EKS/ECS Anywhere** is a deployment option for Amazon EKS/ECS to create and operate Kubernetes clusters and container workloads on-premises.
- **Azure Stack from Microsoft** is a collection of technologies that extend Azure services and capabilities to the data center and edge computing. **Azure Arc** extends Microsoft hybrid cloud portfolio enabling the registration and management of bare metal servers, virtual machines, and Kubernetes clusters with Azure.
- **VMware vSphere 7** is a unified platform to manage virtual machines and containers. **VMware Tanzu Mission Control** allows the management of Kubernetes clusters deployed anywhere.
- **Red Hat OpenShift** is a popular enterprise PaaS offering available for use on both public and private clouds.

Hybrid toolsets mean administrators avoid the need to become experts on multiple domains and skill sets and manage complex software with small teams in-house and without the need to buy in external expert consultants. Hybrid tools facilitate moving workloads from on-premises to the cloud (or repatriated back to on-premises).

#3 Monitor all the Key Cloud Services your Organization Uses

There are many different services that organizations can use in the cloud. AWS offers over 200 full featured services. Azure also has over 200 different products and services. The table below summarizes the most popular AWS and Azure cloud services.

The KPIs for each service vary because the functionality of each service is different. Organizations need a monitoring solution that can track availability, performance, usage and health indicators for each service. As cloud services cost money, tracking the cost of every cloud service used is also important. These KPIs can be obtained through tight integration with cloud provider APIs (e.g., AWS CloudWatch and Azure Monitor).

#4 Monitor the Usage and Performance of Cloud Services with Analytics for Automatic Problem Detection

The basic metrics provided by cloud service providers focus on resource usage of these services. For example, consider AWS Relational Database Service (RDS). CloudWatch provides details of CPU used, IOPS used, etc. Whilst this is useful information, what IT operations teams need is a greater degree of visibility into the SQL server to understand what queries are causing the CPU usage or I/O activity on the database server. This level of granular insight is not available as part of the basic capability available from

Cloud service	amazon	Azure
Compute Instances	EC2	Virtual Machines
Workspaces	WorkSpaces, AppStream	Azure Virtual Desktop
Storage	Elastic Block Storage (EBS)	Azure Disk Storage
Relational Database Systems	RDS, Aurora, DynamoDB, Redshift	SQL Database, Database for MySQL, Database for PostgreSQL, Cosmo DB, Cache for Redis
Notification	SNS (Simple Notification Service)	Service Bus
Orchestration	EKS	Kubernetes Service (AKS)
Containers	ECS - Amazon Elastic Container Service	Container Instance
Load Balancing	ELB - Amazon Elastic Load Balancing	Load Balancer Application Gateway
In-memory caching	Amazon ElastiCache	Elastic on Azure
Event-based compute	Amazon Lambda	Functions
Serverless	AWS Fargate	Serverless Kubernetes, Serverless functions, Azure App service

cloud vendors. To achieve this level of visibility, you may need to install additional agents and subscribe to additional services such as RDS performance insights in AWS to gain the level of visibility into application performance required.

For some cloud services, the additional level of detail necessary to track these services may not even be available with the native cloud provider tools. Consider AWS EKS (Elastic Kubernetes Service). AWS CloudWatch does not provide sufficient visibility into the configuration and performance of your Kubernetes environment – i.e., the namespaces, Pods, containers, etc. You will need your cloud monitoring tool to provide this added visibility.

> ☞ Over 70% of respondents indicated that their cloud provider's built-in monitoring tool did not provide the capabilities needed for indepth application monitoring.
>
> Source: eG Innovations and DevOps Institute APM Survey 2021

Whatever monitoring tool you select, it should provide you with the relevant granularity and insights needed for your cloud applications and infrastructure. For example:

- If you are using burstable AWS EC2 instances, what is your CPU credit balance on each instance? Is it low enough to affect the performance of applications running on that instance?
- If your Azure database is seeing a large number of IOPs, which are the queries responsible for this?
- If one of your containers running Apache Tomcat is taking a large amount of CPU, which thread in the application server is responsible for this and what line of code is it executing?
- If you are using Azure Virtual Desktops(AVD), are user logon times within acceptable limits? If not, what is the root-cause?

Overall, your performance monitoring strategy must focus on full stack visibility – not just infrastructure monitoring. You will need to ensure that you have observability of your entire service delivery chain. This involves the collection of metrics, logs, traces and other signals from the target infrastructure and applying AIOps techniques to get insights from such data and presenting it to users in such a manner that it allows them to make intelligent informed decisions.

Monitoring vs. Observability

Monitoring involves collecting data that allows teams to watch and understand the state of their systems. Monitoring is based on gathering predefined sets of metrics or logs.

Observability involves proactively collecting, visualizing, and applying correlative intelligence to your metrics, events, logs, and traces across all layers and tiers to gain a holistic understanding of your entire software stack. Observability allows teams to actively debug their system and is based on exploring properties and patterns not defined in advance. Metrics, events, logs, and traces (MELT) are at the core of observability.

Given the scale of deployments in the cloud, a cloud monitoring solution will not just be a data collection engine. It must include an AIOps engine that provides anomaly detection and trustworthy alerting, so false positive alerts to IT operations teams are minimized.

Key AIOps features typically include:

- Auto-discovery and dependency mapping
- Embedded domain expertise for metrics collection
- Auto-baselining and anomaly detection for proactive alerting
- Top-to-bottom, end-to-end auto-correlation
- Automatic problem diagnosis
- Metric aggregation and service quality indicators
- Bottleneck detection, forecasting and capacity planning
- Automatic correction and remediation

Such capabilities automate IT operations and alerting and remove a lot of the manual and time-consuming analysis needed to diagnose performance issues.

#5 Monitor the Digital User Experience for Different Workloads

The performance of all IT applications these days must be measured by how satisfied users are – i.e., the digital user and employee experience, and not by the resource usage levels of the applications. Cloud applications are no different – in fact because your organization may not have access to the cloud infrastructure supporting your applications, monitoring user experience for cloud applications is even more important than ever before.

During the cloud migration phase, user experiences must be compared before and after the migration to the cloud. It is essential that the migration to the cloud keeps the user experience at the same or better levels than it was when the application was on-premises. For organizations implementing hybrid and remote working models, quantifying the Digital Employee Experience (DEX) has become increasingly critical.

The cloud monitoring solution must support the two standard methods for monitoring user experience:

Synthetic monitoring

Synthetic monitoring uses "robot users" to emulate users accessing services from one or more locations and reports on service availability and response times. Robots will need to be application-specific. A robot for an email application will check the sending, reception, and delivery of email messages, while one for a digital workspace will simulate a user logging in to the digital workspace and checking availability of the service and logon time.

Beyond logon and availability simulations, synthetic monitoring should support the ability to also simulate multi-step transactions and monitor the performance of each step. This capability is especially important if you are using SaaS services where you will not get access to the infrastructure or the application components supporting the service. Synthetic monitoring is the only way in which you can monitor the performance of SaaS services. Hence, every cloud monitoring solution must support synthetic monitoring for all workloads – web applications, email applications, digital workspaces, etc.

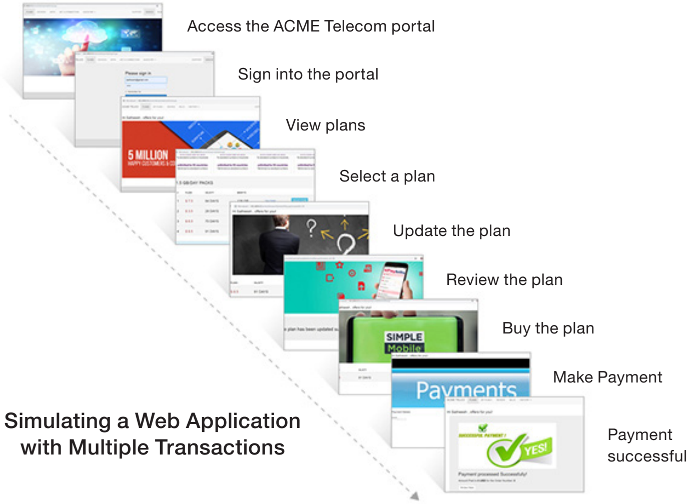

Simulating a Web Application with Multiple Transactions

Repeatable synthetic monitoring is an industry best practice enabling organizations to establish baselines and quantifiable KPIs for success and on-going performance and availability targets.

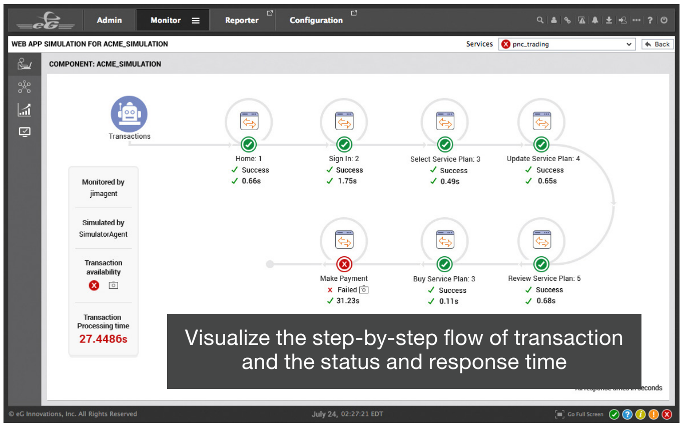

WEB APP SIMULATION FOR ACME_SIMULATION

Services ☒ pnc_trading

COMPONENT: ACME_SIMULATION

Monitored by jimagent

Simulated by SimulatorAgent

Transaction availability ☒

Transaction Processing time 27.4486s

Home: 1 ✓ Success ✓ 0.66s

Sign In: 2 ✓ Success ✓ 1.75s

Select Service Plan: 3 ✓ Success ✓ 0.49s

Update Service Plan: 4 ✓ Success ✓ 0.65s

Review Service Plan: 5 ✓ Success ✓ 0.68s

Buy Service Plan: 3 ✓ Success ✓ 0.11s

Make Payment x Failed ✓ 31.23s

All response times in seconds

© eG Innovations, Inc. All Rights Reserved July 24, 02:27:21 EDT

Visualize the step-by-step flow of transaction and the status and response time

Real user monitoring

While synthetic monitoring provides round-the-clock insights into cloud application performance, it only simulates accesses by a few users and for a set of pre-specified transactions. Real user monitoring is required to track the performance seen by the broader user base, across all their application accesses.

In a cloud environment, monitoring of real users is performed in an application-specific manner, using application API calls in the case of Office 365, digital workspaces, and other services. For web applications, real user monitoring is based on JavaScript injection. In this agentless monitoring approach, a small JavaScript snippet is added to the web content from the application. This snippet can be introduced from a front-end load balancer or web server, so its injection is transparent to the web application. When the JavaScript snippet is executed on a user's browser, it sends information to the monitoring system about the URL accessed, the response time overall as well as its breakdown into network, content, server, and client time, and whether any JavaScript errors are seen on the browser. Using this passive and continuous measurement of end-user experience, you can get insights into application performance and user activity analytics.

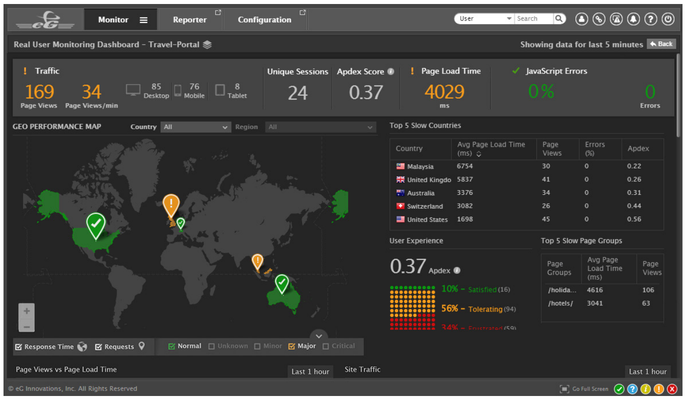

Top 5 Slow Countries

Country	Avg Page Load Time (ms)	Page Views	Errors (%)	Apdex
Malaysia	6754	30	0	0.22
United Kingdo	5837	41	0	0.26
Australia	3376	34	0	0.31
Switzerland	3082	26	0	0.44
United States	1698	45	0	0.56

User Experience

0.37 Apdex

10% – Satisfied (16)

56% – Tolerating (94)

34% – Frustrated (59)

Top 5 Slow Page Groups

Page Groups	Avg Page Load Time (ms)	Page Views
/holida...	4616	106
/hotels/	3041	63

Real user monitoring highlighting
the digital user experience of users of a web application

JavaScript-based real user monitoring can be used for any type of web application including commercial applications like Microsoft SharePoint and Dynamics, PeopleSoft, Atlassian Confluence, core banking applications, healthcare applications like Cerner and Allscripts, and custom web applications.

System administrators can troubleshoot user experience issues with real-time visibility into website performance and errors, and fast insights into the actual cause of slowdowns. As every single user is automatically tracked, issues can be retrospectively analyzed even if systems are down or users not logged on and insights across the entire demographic of users are available to identify whether issues are associated with a common infrastructure component, a specific region, a particular browser etc.

#6 Provide Clear Demarcation of Problems

When you move to the cloud, there is at least one other domain of control – the cloud service provider. You no longer have direct access or insight into the underlying cloud infrastructure. Depending on what service you are using, you may just have access to the components supporting your applications. Your users may be connecting directly to cloud services from remote locations, or from your data center.

In such scenarios, when a problem occurs, you will need to determine the cause of the issue. Could it be the network between your user's remote site and the cloud environment, or could it be your internet link to the cloud platform, or could it be because of an issue in the cloud infrastructure that you may not have visibility into? Sometimes, application issues can escalate as cloud platform issues – for example, a poorly designed query in your application may cause excessive reads on the SQL database server and this may result in your database server exceeding the number of IOPS you have reserved for it. This will result in performance bottlenecks in the database tier, but this was induced by the application and in this case, is not a cloud provider issue.

To clearly demarcate problems, you need:

- **Instrumentation at all levels** to be able to clearly demarcate when a problem occurs, what is causing the issue. Without effective demarcation, you could be spending hours troubleshooting.
- **The ability to visualize metrics across layers** – is the OS sized correctly, are there any abnormalities at the TCP transport layer, has the application workload changed, etc. Some tools use layered/stack representations to make problem diagnosis easier.
- **Service topology graphs showing application inter-dependencies**: This is a key capability as it enables

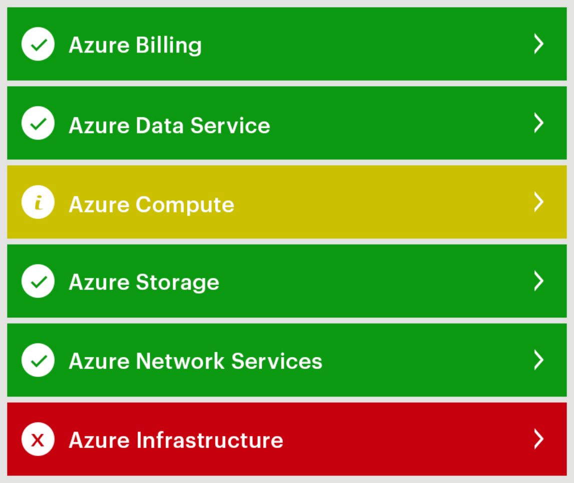

✓ Azure Billing

✓ Azure Data Service

i Azure Compute

✓ Azure Storage

✓ Azure Network Services

✕ Azure Infrastructure

A layer model view of a Microsoft Azure subscription and its services

the IT manager to visually see the health of all the tiers supporting an application including the inter-dependencies between them. Color-coded service topology views highlight where the root cause of problems lie.

- **Transaction level drilldowns that show where your application is spending time when processing a user request**: For server-side technologies like Java, Microsoft .NET, PHP amongst others, most modern application performance monitoring tools can trace a request's processing path, without requiring any code changes in the application. As your application may use diverse types of cloud technologies, it is vital that tag-and-follow tracing be supported across the different cloud services your application uses.

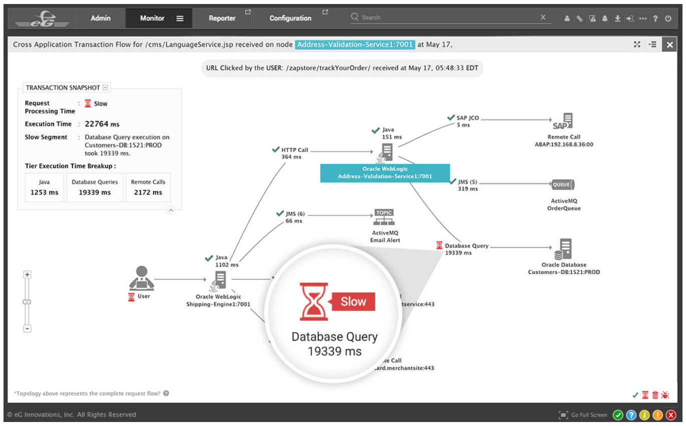

Cross Application Transaction Flow for /cms/LanguageService.jsp received on node Address-Validation-Service1:7001 at May 17,

URL Clicked by the USER: /zapstore/trackYourOrder/ received at May 17, 05:48:33 EDT

TRANSACTION SNAPSHOT

Request Processing Time : Slow

Execution Time : 22764 ms

Slow Segment : Database Query execution on Customers-DB:1521:PROD took 19339 ms.

Tier Execution Time Breakup :

Java	Database Queries	Remote Calls
1253 ms	19339 ms	2172 ms

*Topology above represents the complete request flow?

Transaction flow graph showing time spent by a transaction in each application/cloud tier

Additionally, you will need historical reports as proof if you need to raise support tickets with third parties that your application depends on – e.g., the cloud infrastructure provider, payment gateway provider, SMS gateway service provider, etc.

Finally, a common question that may resonate is that when a performance issue escalates it leads to one asking, "what has changed?". Your application may have worked well for weeks and suddenly if it has a problem, you need to know if any recent configuration changes have taken place – was an instance configuration changed, was a hot fix applied, was new software installed, etc. Ideally, your monitoring tool should be capable of tracking configuration changes in addition to performance, so when a performance problem is detected, you can swiftly determine if there was a significant config' change that took place around the time when the performance problem started.

#7 Support Cloud Technologies like Auto-Scaling and Microservices with Auto-Deployment, Discovery and Auto-Configuration

Cloud services are intended to make monitoring and maintaining IT services easy. At the same time, organizations move to the cloud for agility. Even when using cloud compute instances, you can make use of the auto-scaling capabilities of the cloud. For example, a web site may start with 4 load balanced web servers. When the load on the site increases, additional web servers may be spun up automatically. In such a scenario, the monitoring tool must auto-discover the new instances and monitor them automatically, without needing any human intervention. Likewise, when the load reduces, the newly added web server instances may be shut down. The monitoring tool must be capable of supporting such dynamic workloads.

If you are re-architecting your application to use microservices and containers, the need to automate your monitoring is critical. Containers may spin up dynamically and a manual approach for installing agents, configuring the monitoring and tuning baselines is simply not sustainable or feasible.

Dynamic and scalable systems demand features that go beyond those of legacy or traditional tools used in static on-premises deployments.

Monitoring tools for the cloud must be geared towards enabling monitoring, with as little human intervention and effort being needed as possible.

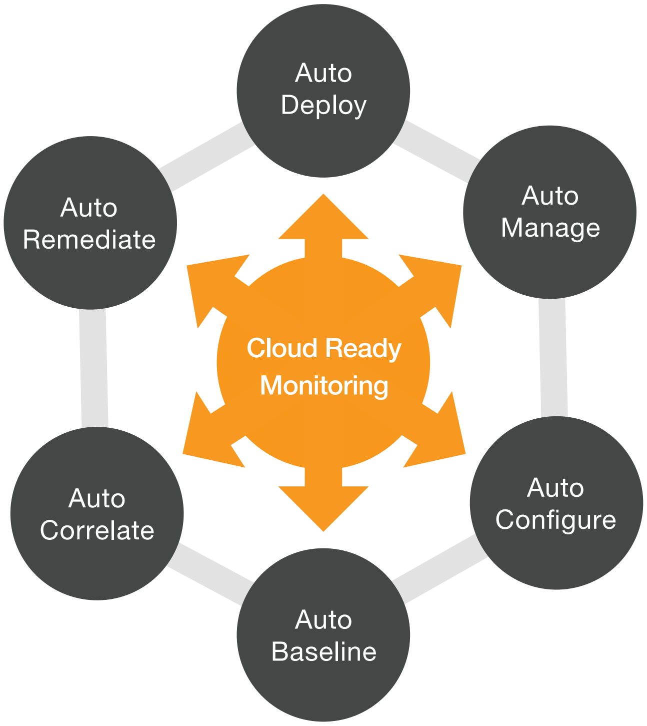

- In this regard, the monitoring tool must be easy to deploy – it must integrate with automation and orchestration tools (e.g., Terraform, CloudFormation, etc.) being used in the cloud environment, so when a cloud instance or a container is spun up, the instrumentation necessary for monitoring is automatically deployed and enabled.
- Auto-discovery must be supported, so once an agent is deployed on an instance, all the applications and dependent components are auto-discovered.
- Auto discovered components must be auto-managed and the necessary permissions needed for the monitoring enabled. Discovery must cover all application components including the inter-dependencies between them.
- When an application component is destroyed, deleted, or terminated, the monitoring tool must auto-discover this and remove the component from monitoring.
- The ability to monitor groups of similar components, not just individual components (e.g., web server clusters designed for failover and high availability, not individual web servers) is important. For application architects, administrators, and service managers who want to understand the total workload and availability of an application component, this is important. Alerts must be raised based on thresholds at the cluster/group level, e.g., an alert can be set to trigger if the availability in a web services cluster falls below 60% - i.e., 3 out of 5 web servers are down.

#8 Ensure Data Security, Compliance and Governance

Security, governance, and compliance are crucial factors to consider for every aspect of cloud service delivery, and monitoring is no exception. Monitoring of the cloud environment must be done securely. When deploying agents on cloud instances, you should check to see if non-standard TCP ports are used for communication. Most modern monitoring architectures are based on HTTPS and do not require additional TCP ports to be opened on the agents or the manager platform to support monitoring. Any credentials stored on the agents must be protected with industry standard encryption techniques.

If you are using monitoring as SaaS (Software as a Service) – i.e., the monitoring server is hosted in the cloud by your monitoring service provider, you will want to ensure that the cloud provider follows all OWASP (Open Web Application Security Project) best practices to safeguard your application. You will also need to ensure that your monitoring vendor complies with all required compliance, regulatory and security standards (depending on your organization's industry segment focus). You should verify the geographical location or public cloud region where the monitoring service is being delivered from and that the tenant-segregation they provide meets your own standards. For instance, if you are a European organization, you may want the monitoring service to be delivered from a region within Europe and that your data never leaves this region. Irrespective of whether you choose an on-premises or cloud deployment, ensure that you comply with government and industry regulatory requirements, such as the European GDPR (General Data Protection Regulation), German BDSG, and the Australian Privacy Act, and so on.

When choosing a SaaS monitoring solution hosted on a tier-1 vendor's public cloud you should affirm whether the vendor has implemented the cloud vendor's architectural standards and undergone vendor audit to achieve "Verified" or "Well-Architected" status. AWS, Azure and Google all offer such vendor validation programs and choosing a validated solution ensures that third-party products adhere to AWS/Azure/Google's own criteria for reliability, customer support and security.

#9 Provide a Predictable Billing Model for Monitoring

Cloud billing, especially with PAYG (Pay as you go) subscription components, can be extremely hard to estimate without a significant quantity of historical data. Built-in cloud service provider tools can collect a lot of metrics, but the cost of monitoring depends on a range of factors, including:

- Storage of historical data and how long data is retained
- Data sampling rates on metrics
- Per alarm/alert set on metrics
- Policies used for alerting – whether automatically determined or manually set
- The number of metrics may depend on the number of processes running on a system

Native cloud-based monitoring can rapidly become extremely expensive and frequently cloud environments have minimal monitoring to control costs.

When selecting cloud monitoring solutions, you need to consider products with licensing models designed for cloud environments. Whilst you have visibility into and control over the physical infrastructure on-premises, you will undoubtedly not have the same luxury in the cloud. For instance, licensing based on hypervisors is not available in cloud environments.

In keeping with the focus of cloud environments to offload non-core activities to service providers, monitoring SaaS solutions are often preferred for monitoring over software components that you have to install, support and maintain. Billing must provide monthly or annual subscription options. Monitoring tools that are licensed per instance, per OS, etc. often provide much better cost control and are usually easier to configure and deploy as there are fewer settings. Ideally, you want the monitoring tool's billing to scale up or down as your cloud footprint auto-scales. Synthetic monitoring should be licensed per location from where it is executed and per target, rather than per test, per protocol tested or per metric collected.

Some vendors offer transferable licensing between technologies enabling administrators to plan and baseline pre-migration on-premises and simply transfer their monitoring licenses to the cloud alongside their users or services; and of course, allowing the organization to move between clouds or back on-premises on-going.

#10 Monitor Cloud Billing Costs and Provide Cost Optimization Analysis

Countless businesses move to cloud environments believing they will make significant savings against operational overheads, particularly hardware. This is often not the case as many organizations see costs spiral out of control. The cloud makes it simple to provision and configure operational and business systems. However, there is, often, significant overhead and cost-sprawl associated with their management, including the on-going need for subject matter expert staff to manage and control their costs.

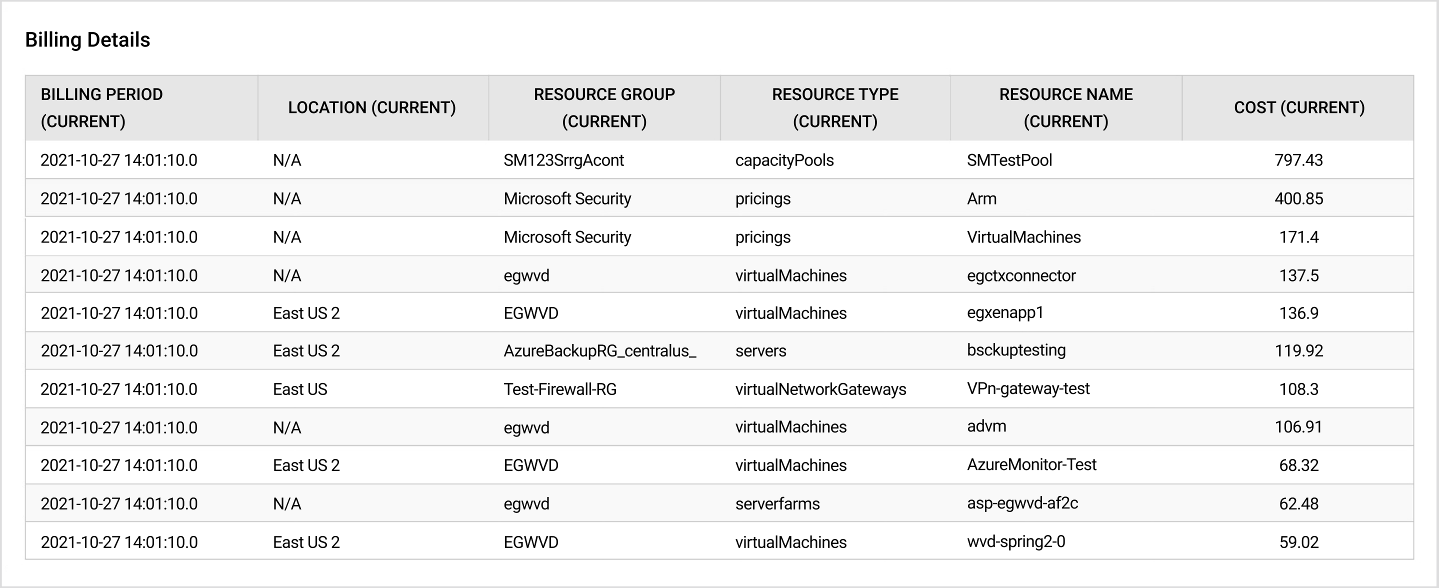

Billing Details

BILLING PERIOD (CURRENT)	LOCATION (CURRENT)	RESOURCE GROUP (CURRENT)	RESOURCE TYPE (CURRENT)	RESOURCE NAME (CURRENT)	COST (CURRENT)
2021-10-27 14:01:10.0	N/A	SM123SrrgAcont	capacityPools	SMTestPool	797.43
2021-10-27 14:01:10.0	N/A	Microsoft Security	pricings	Arm	400.85
2021-10-27 14:01:10.0	N/A	Microsoft Security	pricings	VirtualMachines	171.4
2021-10-27 14:01:10.0	N/A	egwvd	virtualMachines	egctxconnector	137.5
2021-10-27 14:01:10.0	East US 2	EGWVD	virtualMachines	egxenapp1	136.9
2021-10-27 14:01:10.0	East US 2	AzureBackupRG_centralus_	servers	bsckuptesting	119.92
2021-10-27 14:01:10.0	East US	Test-Firewall-RG	virtualNetworkGateways	VPn-gateway-test	108.3
2021-10-27 14:01:10.0	N/A	egwvd	virtualMachines	advm	106.91
2021-10-27 14:01:10.0	East US 2	EGWVD	virtualMachines	AzureMonitor-Test	68.32
2021-10-27 14:01:10.0	N/A	egwvd	serverfarms	asp-egwvd-af2c	62.48
2021-10-27 14:01:10.0	East US 2	EGWVD	virtualMachines	wvd-spring2-0	59.02

Tracking cloud costs by resources is important

From a monitoring perspective, it is not sufficient to monitor IT metrics alone. Utilization, downtime, performance, etc. are important, but equally important is the imperative for organizations to make sure that they only spend what is necessary in the cloud, meaning that all cloud instances should be optimally right-sized and utilized. Any oversized instances should be identified, and unused instances should be detected and shutdown. If IOPS are reserved for database systems, are these over provisioned? And can the organization look to reduce the associated costs? These are just some of the questions that IT management needs to address. Monitoring of cloud usage, billing, and providing recommendations for right-sizing cloud environments should be priority requirements.

The resources and performance of some cloud infrastructure is often tied to the billing model in place particularly with burstable instances where the CPU credit balance can affect the CPU available and whether throttling occurs. Monitoring your cloud account balances within the same product as you monitor IT allows the cost and performance implications of configurations and usage to be accurately determined, the vast majority of native cloud tools have yet to implement such capabilities.

A cloud monitoring product can aid administrators by:

- Allowing alerts to be triggered when costs exceed desired thresholds
- Provide recommendations on resources that can be conserved to minimize cost
- Highlight resources that are under-provisioned and therefore are affecting application performance
- Monitoring costs and performance of multiple clouds within a single interface, allowing proper like-for-like value comparisons to be made

Conclusion

In this white paper, we have reviewed the top 10 monitoring requirements for cloud services and infrastructures. In assessing this list, we have primarily focused on the unique challenges that cloud applications and infrastructures pose. The ideal monitoring solution should also address other enterprise-class monitoring requirements such as single sign-on support through integration with Active Directory or support for SAML, role-based access control, tight integration with ticketing systems, personalized dashboards, and reports, and so on.

With the right level of end-to-end visibility in place for your cloud applications, you will also be able prove to your cloud service provider when the issue is in their end. And as you start out on your cloud journey, make sure you don't make the same mistakes that you may have made with on-premises deployment. Seek to reduce monitoring tool and console sprawl.

A well thought out cloud monitoring strategy can yield significant benefits. Tool consolidation will yield cost optimization. With cross-tier visibility, you can get all your IT operational teams on the same page – they will see how each tier is performing and the impact on the user experience. This will result in faster diagnosis, improved resource alignment and reduced mean time to repair (MTTR).

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