eG Innovations

APM Best Practices:
How to Maximize IT Performance

www.eginnovations.com

Executive Summary: Why Modern APM Demands a New Approach

Application process monitoring (APM) has changed drastically over the last decade.

Not long ago, most organizations relied on a patchwork of separate monitoring systems — one for servers, another for networks, and yet another for applications. This lack of cohesion meant that IT teams had to jump between dashboards, manually piecing together data in search of the source of a problem.

What's more, they were stuck in reactive mode, responding to user complaints or alerts that came after the system failure had already frustrated users, interrupted workflows, or even hurt the bottom line.

Today, that just doesn't cut it. Applications are now hybridized, spanning multiple layers and platforms — cloud, on-premises, virtual, and everything in between. Expectations are higher than ever, and even small performance issues can cause big disruptions.

That's why modern IT monitoring requires a simple, unified solution. At eG solutions, we've developed a "one-stop shop" approach to APM that prioritizes visibility, prevention, and user experience. In this guide, you'll learn the new APM best practices that underpin this approach.

Table of Contents

01 Consolidate Tools into a Unified Monitoring Platform

As your IT environment grows more complex, it's not uncommon to end up with a bloated toolkit for monitoring each component of your tech stack. You might have a separate dashboard, database, and alert system to track things like app performance, network health, server uptime, and other vitals.

This "tool sprawl" creates more confusion than clarity and can complicate your business's day-to-day operations. you'll learn the new APM best practices that underpin this approach.

The Problem with Disconnected Monitoring

Without a unified view, problems can be easy to miss, and teams may operate in silos instead of taking a collaborative approach. Siloed IT systems create fragmented informational landscapes, making it more difficult to diagnose and address issues efficiently.

When a problem occurs, IT teams must toggle between tools to compare metrics from different sources and piece together a timeline. That slows the root cause analysis process and increases the risk of false leads, misdiagnosis, and unnecessarily drawn-out downtime.

Even worse, critical insights can slip through the cracks. For instance, a sluggish application might look fine in your server dashboard but could, in fact, be caused by a network bottleneck or a misconfiguration in the cloud.

The Solution: One Platform, One Source of Truth

The best way forward is a consolidated, full-stack approach that allows you to see everything in one place. Unifying application, infrastructure, and UX monitoring means you can eliminate blind spots and solve problems faster — ideally, before they turn into outages or support tickets.

With this in mind, we built eG Enterprise as a single, integrated APM platform. It offers an at-a-glance snapshot of app performance, bringing together data from across your entire digital ecosystem.

Similarly, you'll want a dashboard from which you can easily zoom in and out on each layer of your stack, so you'll have both the high-level insights and granular data you need to stay operational.

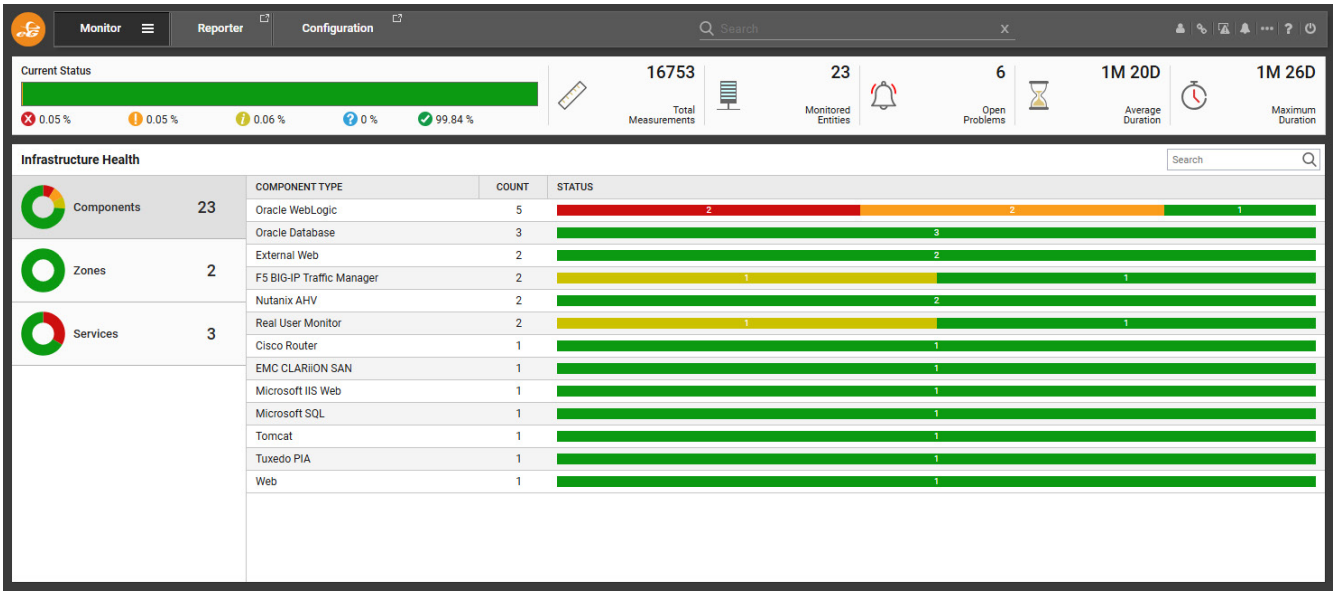

Figure 1: Caption to be updated

02 Shift from Reactive to Proactive Monitoring with AI & ML

While it may once have been the norm for IT teams to operate as "firefighters" — constantly responding to

system errors and outages after the fact — that reactionary approach no longer meets the demands of today's always-on digital environments.

Moreover, it's no longer necessary, as machine learning and other advanced technologies have advanced to the point where IT teams can mitigate issues with proactive monitoring and intervention.

The Power of Predictive Monitoring

Traditional monitoring systems rely on static thresholds — detecting anomalies based on fixed performance limits. But these thresholds don't adapt to real-world usage patterns. A normal spike during peak hours might be incorrectly flagged as an issue, while a slow-building performance problem could go unnoticed.

That's why shifting towards AI-driven predictive monitoring that anticipates problems based on patterns and trends is crucial, rather than just responding to threshold breaches.

Machine learning can identify what's "normal" for your environment, allowing you to identify subtle changes before they become major problems. Instead of reacting to alerts after users are affected, you'll be notified in time to fix issues before they escalate.

Use Case: Preventing Downtime with Synthetic Testing

One of our clients, a mid-sized financial services firm, noticed intermittent sluggishness in their customer-facing app during early morning hours. Response times soared to 3-4x their normal levels, in some cases timing out completely. Normally, this would have gone unnoticed until helpdesk tickets began rolling in.

But with predictive monitoring and synthetic testing in place, the team didn't have to wait for complaints. eG Enterprise's Web App Simulator continuously ran synthetic transactions — simulated user journeys like login, balance checks, and fund transfers — across different times and locations. These synthetic tests detected the early-morning performance drop before real users were active on the system.

The platform automatically flagged the anomaly and correlated it with a routine batch job consuming database resources during that window. The IT team quickly adjusted the schedule, resolving the issue proactively and avoiding a wave of frustration.

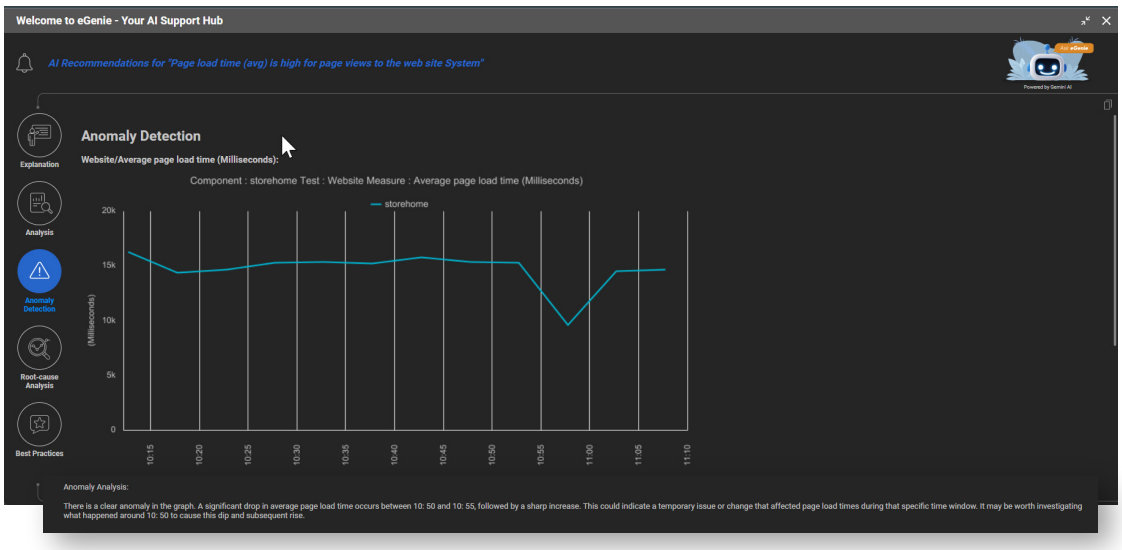

Figure 2: Caption to be updated

03 Ensure True End-to-End Visibility, Not Just Metrics

When systems slow down or fail, knowing your CPU usage or response time is only the first piece of the puzzle.

Traditional monitoring platforms report individual metrics, but they often stop short of showing how those metrics relate to what's actually happening, such as user experience, business transactions, or behind–the-scenes system interactions.

As you modernize your APM systems, build with these throughlines in mind, rather than just raw data.

Metrics Don't Give the Full Story

Most IT teams still rely on isolated performance data: server load, database latency, network throughput, and the like. While useful, these metrics don't tell you why something is happening or who's affected. They may not help you understand the origins of an issue, meaning troubleshooting can amount to trial-and-error.

On the other hand, an APM approach that prioritizes end-to-end visibility allows you to connect the dots between performance issues and real-world outcomes so you can act with clarity.

Track Real Users and Transactions for the Big Picture

We developed our APM solution to trace the processes happening behind the scenes. For instance, our business transaction monitoring system follows each request from click to database and back across all tiers of the client's stack. Whether it's a slow checkout, login delay, or API error, IT teams can identify exactly where things are breaking down.

We also use Real User Monitoring (RUM) to capture how actual users interact with a web or mobile app so our clients can see what's going on at the ground level. They can see who abandoned a transaction, where they dropped off, and whether it may have been due to a technical issue. With this data, they're well-positioned to address any friction points that may be affecting the bottom line.

This kind of end-to-end visibility is central to effective performance monitoring, offering the context teams need to resolve issues faster, minimize impact, and continuously improve the digital experience.

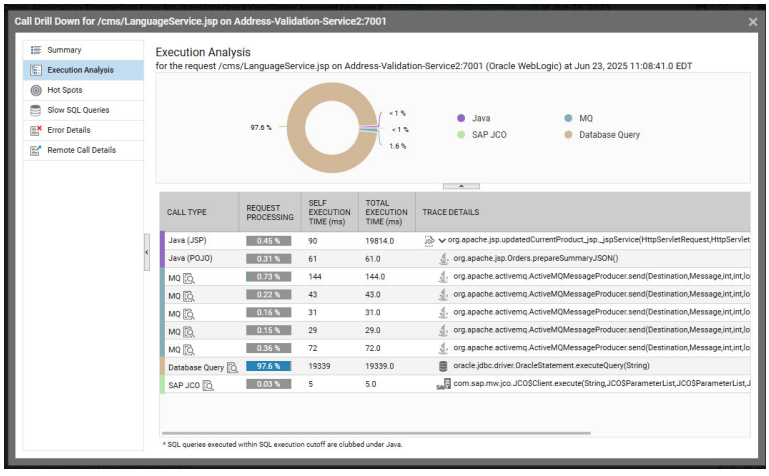

Figure 3: Enterprise applications eG supports.

04 Optimize Performance Across Hybrid & Multi-Cloud Environments

Cloud adoption has opened new doors for agility and scalability, allowing you to run all or some of your workloads on AWS, Microsoft Azure, and Google Cloud.

But as you hybridize your systems or distribute them across multiple cloud vendors, performance monitoring can get tricky. Visibility gaps, shifting dependencies, and inconsistent metrics can make it hard to see what's happening, let alone leverage that information to improve your product.

The Challenge of Fragmented Visibility

In hybrid and multi-cloud setups, workloads might span data centers, cloud regions, and virtual machines while relying on shared services and third-party APIs.

Static monitoring tools struggle to keep pace in these situations. They cannot often track dependencies in real time, leaving teams in the dark about how one failure might cascade through interconnected services.

This lack of visibility can also make it tough to know where to start when it comes to fine-tuning performance and making meaningful improvements to your products or services. Without a clear view of your system performance, you may have blind spots and inefficiencies that ripple across infrastructure, apps, and UX.

Managing Performance Across Complex Environments

To optimize performance in a hybrid or multi-cloud environment, your monitoring strategy needs to evolve alongside the infrastructure. That means moving beyond tool-specific views toward a unified approach that keeps every platform, service, and dependency within view.

A strong APM system should:

- **Support diverse environments:** Cloud-native services, on-prem hardware, containers, virtual machines, and third-party APIs all need to be monitored within the same context.
- **Automatically map dependencies:** Real-time dependency mapping helps you understand how components interact as workloads shift and scale.
- **Adapt to change without manual upkeep:** Environments change fast. Monitoring should adjust automatically without constant reconfiguration.
- **Correlate across layers:** Performance data from the application, infrastructure, and user journey should give you a complete picture of what's happening and why.

Integrating capabilities like auto-discovery, cloud-native support, and real-time correlation allows a well-built APM system to provide a unified view, even in complex, distributed environments. This reduces troubleshooting time, minimizes friction between teams, and helps resolve issues at the source.

The result is a more resilient, high-performing IT ecosystem that consistently delivers a better user experience.

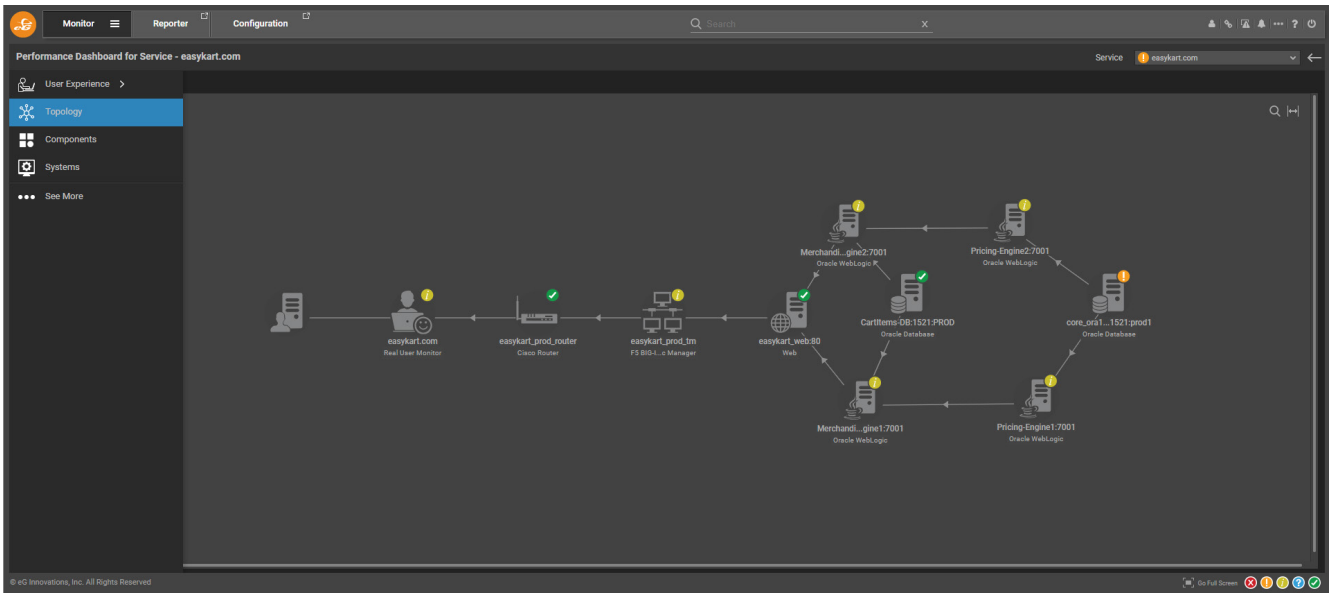

Figure 4: Caption to be updated.

05

Improve MTTR with Root Cause Diagnostics

Digital environments move fast, and even short disruptions can have outsized consequences. The faster you can identify and resolve the root cause of an issue, the less damage it does to productivity, user experience, and trust. That's why mean time to resolution (MTTR) has become one of the most important metrics for high-performing IT teams.

The Problem with Manual Root Cause Analysis (RCA)

When system failures occur, it's not uncommon for teams to embark on a scattershot RCA process, sifting through logs, alerts, and metrics, often across disconnected platforms and team silos.

- App teams check code.
- Infrastructure teams check servers.
- Network teams check throughput.

In the meantime, the clock keeps ticking. This fragmented approach to RCA can lead to duplicated effort, internal friction, and delays in resolution. It often results in treating symptoms rather than fixing the actual issue.

But quick resolution depends on clarity. And that's where many monitoring approaches fall short.

A Better Approach: Automated RCA Across Stack Layers

Reducing MTTR starts with better root cause analyses. As you develop your APM systems, aim for automation, cross-layer functionality, and context-rich insights that tie symptoms to their source. A well-developed APM solution doesn't just alert you when something goes wrong: it helps you understand why, where, and what to fix without chasing down every alert manually.

We built in correlated insights that connect performance data across application, infrastructure, and network layers. Instead of switching between tools, our customers can drill into a single console to trace symptoms to their source, often in just a few clicks.

The result: faster resolutions, less finger-pointing, and a smoother experience for everyone involved, especially your users.

Figure 5: Caption to be updated.

06

Elevate the Digital User Experience with Real-Time Monitoring

While it can be easy to get bogged down in performance metrics, it's critical to remember that there's a real human behind every data point. That's why performance monitoring can't stop at server load or error rates. To truly optimize performance, you'll need to monitor what actually matters: how users experience your product in real time.

Why User-Focused Monitoring Matters

When issues arise, your users don't care what the root cause is — they just know something didn't work. If your IT team is blind to the user's perspective, they're left reacting to support tickets instead of preventing them in the first place.

Rather than spending time in this reactionary cycle, you can leverage user data to proactively address UX issues and continuously make people-first improvements to your product.

Combining Real User Monitoring with Synthetic Testing

To improve UX, you need visibility into both actual and expected user behavior. That's where real user monitoring (RUM) and synthetic transaction testing becomes so valuable.

RUM captures how your users interact with your app in the moment, offering a live view into what's working and what isn't. This could include page load times, errors, device types, regions, and other critical information. Synthetic testing complements this by simulating the user journey from start to finish,

offering a way to automate your UX efforts and keep them running around the clock.

Our APM system correlates user session data with backend performance to give teams the full picture. For instance, they might see that a slowdown was caused by client-side rendering rather than network latency, allowing them to fix the issue faster, and more importantly, prevent the need for users to open a support ticket.

Monitoring the actual user journey — rather than just system-level metrics — is one of the most direct ways to improve satisfaction, reduce churn, and strengthen your product over time.

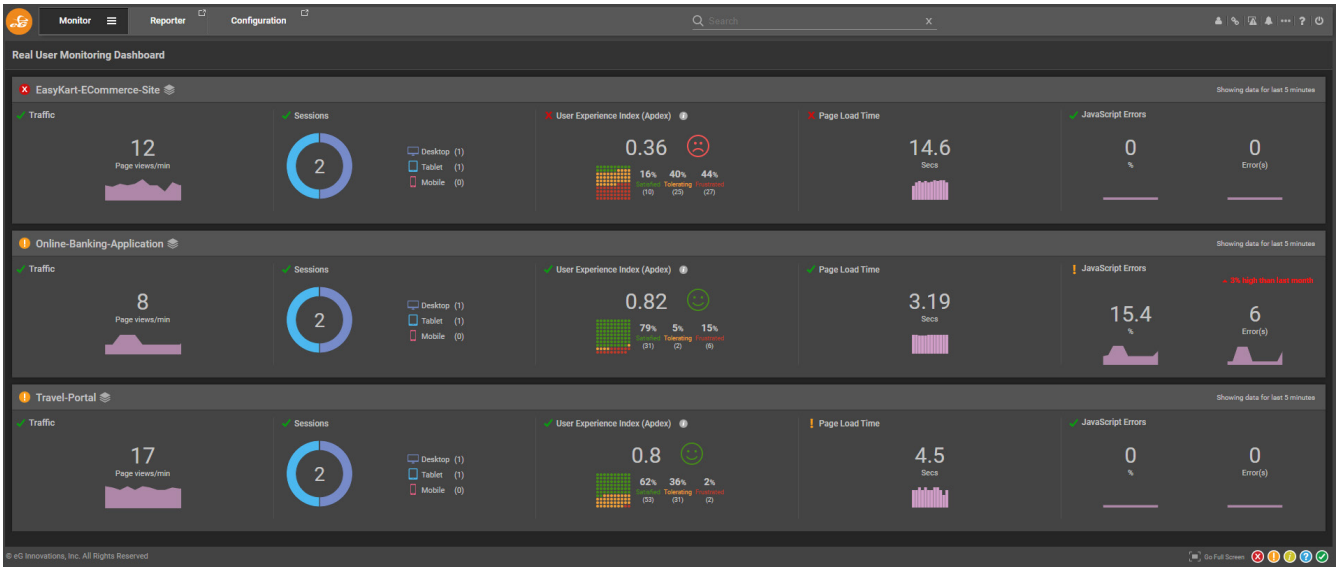

Figure 6: Caption to be updated.

07

Avoid Common Misconceptions that Limit APM Success

As you look to improve performance monitoring, it's essential to be aware of a few common narratives that can stall progress. These can prevent your team from unlocking the full potential of APM. Let's set the record straight:

Myth: APM Is Only for DevOps or Developers

Fact: App performance affects everyone, from IT operations to C-suite leadership. An app slowdown affects your users first and foremost, meaning it impacts your revenue and reputation.

A well-implemented APM system provides insights that support service desk teams, infrastructure leads, and executive decision-makers.

Myth: We Already Monitor Our Servers — APM Is Overkill

Fact: Monitoring CPU and memory usage is just the tip of the iceberg. A robust APM solution connects the dots between infrastructure metrics, app behavior, and user experience, helping you see the full picture and solve problems faster.

Myth: APM Tools Are Only Useful for Cloud-Native or Microservices Environments

Fact: While APM is essential in modern architectures, it's equally valuable in traditional, hybrid, and on-prem environments. eG Enterprise, for example, supports everything from monolithic legacy systems to containerized cloud-native apps.

Myth: More Data = Better Monitoring

Fact: Too much data without context creates noise. The key is intelligent correlation, bringing together the right signals across multiple layers so your team can focus on root causes rather than just raw metrics.

Myth: Implementing APM Is Complex and Time-Consuming

Fact: With the right platform, APM can be deployed quickly and scaled gradually. At eG, we've focused on making enterprise-wide monitoring simple and accessible with guided setup, automatic discovery, and intuitive dashboards that work for all teams.

APM Readiness Checklist

Use this checklist to evaluate your organization's current performance monitoring maturity. Each "yes" gets you one step closer to proactive, full-stack APM.

Tool Consolidation

- ☑ Do we use a unified platform to monitor application, infrastructure, and network performance?
- ☑ Can we trace issues across systems without switching between tools?
- ☑ Are our monitoring dashboards consistent and accessible across teams?

User Monitoring

- ☑ Do we have visibility into real user behavior (via RUM)?
- ☑ Are key user flows being simulated regularly to detect issues before they happen?
- ☑ Can we correlate user complaints with backend performance data?

Root Cause Analysis (RCA) Capabilities

- ☑ Can we automatically detect the root cause of performance issues across layers?
- ☑ Are we reducing mean time to resolution (MTTR) through correlated insights?
- ☑ Do service desk and infrastructure teams have shared diagnostic views?

AI/ML Adoption

- ☑ Are we using machine learning to establish dynamic baselines and reduce false positives?
- ☑ Can our system proactively detect anomalies without manual thresholds?
- ☑ Are we prioritizing alerts based on potential impact?

Cloud Readiness

- Can we monitor workloads across hybrid and multi-cloud environments (e.g., AWS, Azure, GCP)?
- Do we have real-time visibility into cloud-native services, containers, and legacy infrastructure?
- Can we adapt monitoring as workloads shift or scale?

Not sure how you score? Talk to one of our experts about your APM readiness.

Ready to See APM in Action?

In this guide, we've explored the fundamentals of a robust, fully modernized APM strategy. From reducing MTTR and improving UX to eliminating blind spots across hybrid systems, these practices are sure to keep you ahead of the curve and ensure your APM is unified, intelligent, and built for scalability.

With eG Enterprise, you'll benefit from all of the capabilities outlined here and more, helping your team work more efficiently, troubleshoot faster, and deliver the reliability your customers expect, without the noise or complexity of fragmented tools.

Ready to take the next step?

Schedule a Demo

> With eG Enterprise, we can now quickly identify root causes of incidents, resolving them before users are impacted. Automatic prioritization and categorization of alerts helps us better focus on the important issues and prioritize our resources accordingly.

Andrew Gowlett
Senior Consultant, C5 Alliance

About eG Innovations

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com

eG Innovations

+1 866 526 6700

info@eginnovations.com

www.eginnovations.com