


eG Enterprise 7.6

Domain Intelligence Meets AI-Assisted Observability



eG Enterprise 7.6 is built to shorten the path from a user-visible symptom to a defensible root cause. This release broadens visibility across applications, digital workspaces, cloud and infrastructure, deepens technology-specific diagnostics across the stack, and adds AI that interprets structured evidence rather than raw telemetry, giving operations teams more coverage, more context and faster, more confident answers.

1. AI that builds on observability—not a substitute for it

Generative AI is rapidly becoming part of mainstream IT operations. eG Enterprise 7.6 embraces this shift, but with a deliberate design principle: AI should interpret evidence, not replace observability engineering. Rather than passing large volumes of undifferentiated telemetry to an AI model and expecting it to determine what matters, eG Enterprise first applies domain-specific monitoring, analytics, machine learning, anomaly detection, dependency mapping and automated root-cause diagnostics to structure the evidence. AI is then layered on top of this structured evidence to interpret what is happening, explain likely causes and help administrators determine the best action.

An AI explanation is only as reliable as the observability context behind it. An isolated high-CPU metric, for example, cannot establish whether the actual service constraint is compute saturation, a database wait, a dependent cloud API or a degraded end-user session. eG Enterprise combines domain-specific measurements, topology, historical baselines and root-cause analytics before presenting the evidence for AI-assisted interpretation.

GenAI complements — not replaces — eG Enterprise’s monitoring and diagnostic workflows. It can explain performance trends, summarize correlated findings and recommend investigation steps. AI-generated explanations should be evaluated against the underlying metrics and diagnostic evidence before administrators make operational changes.




Figure 1. eG Enterprise layers AI on top of domain-specific telemetry, analytics and root-cause correlation

Thus, AI becomes a force multiplier for domain expertise: eG Enterprise determines what the evidence says; GenAI helps IT operations teams consume and act on it faster.

1.1 AI-Assisted troubleshooting: More GenAI choices, better operational insights

eG Enterprise 7.6 expands the GenAI services that customers can use. In addition to OpenAI ChatGPT and Google Gemini, this release adds integration with Anthropic Claude and Azure OpenAI Service. This lets organizations align AI-assisted operations with their preferred cloud, security, governance and enterprise AI standards.



AI assistance also moves beyond a single incident. Administrators can use GenAI to interpret diagnostic evidence and performance trends over time, turning complex histories into concise operational summaries. Responses can be delivered in the administrator's preferred local language, which is useful for global operations teams working from the same underlying evidence.

 More grounded AI AI works from structured diagnostic evidence rather than undifferentiated raw telemetry.	 Faster operational understanding Complex events and performance histories are translated into concise, actionable explanations.	 Enterprise flexibility Organizations can choose among multiple GenAI services and consume results in the administrator's preferred language.
--	--	---

For managed service providers, eG Enterprise 7.6 adds practical deployment flexibility. An MSP can centrally configure GenAI account settings on an eG Manager for use across all tenants, or require each tenant to supply its own GenAI service account. This supports different operating models for tenancy, governance and cost control.

1.2 OpenAI observability: Performance, usage, governance and cost

eG Enterprise extends its observability capabilities to OpenAI, enabling organizations to proactively monitor the performance, reliability, usage, and cost of their OpenAI services. Using OpenAI REST APIs and an agentless monitoring approach, eG Enterprise provides visibility into individual API capabilities and projects, tracking request volumes, availability, failures, error rates, response latency, and overall API responsiveness. This helps operations teams quickly identify unavailable or error-prone API endpoints and detect abnormal latency or request patterns before they impact AI-enabled applications.




Figure 2. OpenAI observability in eG Enterprise

Operational questions answered

Is the service healthy?	Is usage efficient?	Is spend governed?
Availability, API errors, failure rates, latency and responsiveness by project and capability.	Requests, tokens and feature adoption across chat, embeddings, images, audio and other services.	Daily/monthly cost patterns, unusual consumption and active or unused Admin/Project keys.

Monitoring goes beyond API health. Administrators can track model request rates and token consumption, identify unusual input-token usage or abnormal request volumes, monitor active and unused Admin and Project keys, and analyze daily and monthly cost-consumption patterns. Usage can also be broken down by OpenAI capability, including Responses and Chat Completions, embeddings, images, audio transcription, text-to-speech, and other services, helping organizations understand feature adoption, control AI spend, optimize capacity, and govern their growing use of generative AI.

Observability best practice

Review GenAI service availability and response latency alongside request failures, per-model token consumption and project-level spend. A higher bill by itself does not prove inefficiency. Correlate usage with changes in traffic, model selection, output-token volume and application behavior before acting.

The sections that follow focus on how eG Enterprise 7.6 expands domain intelligence across all four major use cases: digital workspace monitoring, enterprise application monitoring, application performance monitoring and unified infrastructure monitoring. The common design goal is the same in each area: collect the right evidence, understand the technology well enough to interpret it, and reduce the amount of manual correlation an operator must perform.




Figure 3. Four major use cases for eG Enterprise

2. Digital workspace observability: Follow the user journey end to end

Digital workspace incidents rarely stop with a single tier. A slow login, failed application launch or poor meeting experience can originate in the cloud control plane, identity, profile storage, hypervisor, session host, network path, security service edge or endpoint. The key question is not simply which component is unhealthy, but where in the user journey the degradation begins. eG Enterprise 7.6 adds visibility at more points in the service delivery chain so operations teams can follow the symptom back toward the responsible layer.

2.1 Monitoring Azure Virtual Desktop on Azure Local: One view across cloud control and on-prem execution

Azure Local is Microsoft's hybrid infrastructure solution (formerly known as Azure Stack HCI) that lets one run virtual machines, containers, and select Azure services on customer-owned hardware or edge sites. It facilitates the local deployment of both modern and legacy applications across distributed or sovereign locations. Azure Local accelerates cloud and AI innovation by seamlessly delivering new applications, workloads, and services from cloud to edge, using Azure Arc as the unifying control plane.






Figure 4. Azure Local cluster resource and I/O diagnostics

For organizations running Azure Virtual Desktop session hosts on Azure Local, eG Enterprise now brings the Azure control plane and the on-premises resource plane into one operational view. The Azure Local service view tracks connectivity between local infrastructure and the Azure control plane. New agents introduced in this release support HCI-integrated Hyper-V hosts, collecting hypervisor and VM performance metrics, while agents on AVD session hosts provide user-session and logon insights. Together with Azure subscription and AVD control-plane monitoring, operations teams can follow the service from cloud orchestration to a user session.




Figure 5. End-to-end monitoring of Azure Local and AVD environments

The AVD Logon Simulator — eG Enterprise's purpose-built synthetic monitor for validating AVD logons — now supports the latest Windows App web client, keeping synthetic coverage aligned with Microsoft's client evolution. eG Enterprise also reports directly on the availability of FSLogix remote file shares. Instead of stopping at a profile-mount symptom, administrators can see when the underlying file share is unavailable and move immediately toward the likely cause.

2.2 Monitoring Amazon WorkSpaces Applications: More automation, user experience insights

Monitoring for Amazon WorkSpaces Applications (formerly Amazon AppStream 2.0) is now more automated. Session hosts are grouped automatically by Fleet, discovered as they appear and removed from monitoring when decommissioned. In an auto-scaling service, this matters because the monitoring stays aligned with the environment without requiring operators to continually add and remove transient hosts.




Figure 6. Amazon WorkSpaces Applications logon analysis by stage

Troubleshooting insights go deeper now. For Amazon WorkSpaces Applications, just monitoring logon times from Windows APIs is not sufficient as this stack adds additional processing layers (e.g., DCV logon processing, DCV permissions setting, storage root setup). eG Enterprise now tightly integrates with the Amazon logon process and breaks down logon performance for individual user sessions into granular detail, replacing broad “slow logon” symptoms with evidence that points to the responsible stage. Details of logon time breakdown are shown in Figure 6. No other monitoring solution in the market provides this level of in-depth logon time monitoring for AWS WorkSpaces Applications sessions.

Storage connector processing time with breakdown is also reported to highlight any storage processing bottlenecks. eG agents also track per-user VHDX capacity and consumption to proactively identify space issues that affect user sessions. Analysis of Amazon user-connection, DCV server and storage-connector logs is also added for more precise anomaly detection. Together these signals help separate slow identity processing, profile/storage delays and remoting-related problems rather than treating every slow launch as a generic session-host issue.




Figure 7. Key enhancements for Amazon WorkSpaces Applications monitoring and reporting

The precision of user idle time detection has also been improved in this release and a new user productivity report helps IT teams report on user productivity during virtual app/desktop sessions. Capacity-planning support has also been added for Amazon WorkSpaces Applications environments. With the lack of built-in AWS WorkSpaces Applications specific monitoring in AWS CloudWatch, eG Enterprise is well positioned to provide the breadth and depth of monitoring that customers deploying AWS Workspace Applications need.

2.3 Citrix Monitoring: Aligned with architectural and security change

Citrix monitoring in 7.6 is updated around the architectural and operational changes in Citrix products and services that customers are already dealing with: cloud authentication, licensing dependencies and session conditions that do not always appear as simple infrastructure failures.

- The Citrix Cloud service now supports service principals for programmatic API access. Unlike secure clients tied to an individual administrator, a service principal provides a dedicated application identity whose permissions can be managed independently. This model is better suited to monitoring and other machine-to-machine integrations because it reduces dependence on personal administrator accounts and supports tighter access control.

eG Enterprise 7.6 adds support for this authentication model, so customers can move their monitoring integration to service principals without losing coverage, thereby staying current with Citrix Cloud's API-access and security best practices.



Figure 8. Citrix monitoring enhancements in eG Enterprise 7.6

- Citrix licensing visibility is also strengthened. Because on-premises Citrix deployments now depend on connectivity to the cloud-hosted License Activation Service (LAS), eG Enterprise tracks deployment-to-LAS connectivity and can expose communication problems before they become licensing-related availability issues.
- Session troubleshooting now detects black-hole machines and zombie sessions — conditions that can cause failed launches, inaccessible desktops or sessions that continue consuming infrastructure resources without delivering business value. Surfacing these conditions automatically closes an important gap between infrastructure health and actual session usability.
- A new user productivity report looks beyond infrastructure health and identifies users and time periods where desktop or application performance may be reducing productive time. The same reporting approach is available for other supported digital workspace technologies.

Additional operational changes in this release include Layer 7 and network-latency visibility in Citrix Cloud sessions, explicit published-desktop naming in session topology, configurable exclusion of OS background processes from application analysis, and automated actions for corrupted Citrix WMI instrumentation and accumulated temporary profiles.

2.4 Monitoring Omnisia Horizon Cloud: Adding platform signals to operational context

For Omnisia Cloud deployments, Horizon Cloud notifications are analyzed by severity with associated event details and correlated with user and infrastructure performance. This release also checks whether Teams media optimization is enabled on Horizon desktops, helping support teams identify desktops where media offload is not operating as intended. From a compliance perspective, the domain-join status of desktops and VMs is tracked, alerting administrators if a VM is disconnected from the domain.

2.5 Endpoint and single-user VDI observability: See the experience where users feel it

Physical endpoint monitoring now includes Microsoft Teams meeting-performance visibility. During a Teams meeting, administrators can evaluate network and media-quality indicators such as latency, jitter, packet loss and other call-quality signals alongside endpoint resource usage. This helps support teams distinguish a user-specific endpoint or network issue from a broader service problem and troubleshoot poor meetings with evidence rather than guesswork.




Figure 9. Microsoft Teams meeting-performance visibility on a Windows endpoint

For single-user VDI environments, eG Enterprise now reports the domain-joined status of each virtual desktop. This allows administrators to proactively detect when desktops have been disconnected from the domain (meaning users are likely unable to log in), rather than waiting for a user to complain before investigating.

2.6 Zscaler Monitoring: Connect security service edge telemetry to user experience

Zscaler is a cloud-native cybersecurity platform that provides secure internet and application access using a zero-trust model.

- **Zscaler Internet Access (ZIA):** Secures employee internet browsing and cloud application usage by routing traffic through a cloud security stack that includes URL filtering, cloud firewalls, and threat protection.
- **Zscaler Private Access (ZPA):** Acts as a modern, secure alternative to traditional VPNs, giving users policy-based access to internal private apps without exposing the underlying network.
- **Zscaler Digital Experience (ZDX):** Monitors and troubleshoots end-user performance across applications and networks.



eG Enterprise 7.6 extends observability into Zscaler environments with unified monitoring across ZIA, ZPA and ZDX. Service and API availability, token generation, responsiveness and tunnel health can be monitored alongside security posture, including critical, warning, passed and skipped configuration checks. This gives operations and security teams a shared view of service health and configuration risk.

The integration also brings Zscaler Digital Experience data into the same operational context. eG Enterprise discovers ZDX device inventory and configuration, analyzes the end-to-end path from client to egress to Zscaler to application/server, and identifies the segment contributing latency. Application-level metrics such as ZDX score, user count and Page Fetch Time can be analyzed by region, while web probes add per-application and per-device PFT, TTFB, DNS and availability measurements.




Figure 10. Integrated monitoring of Zscaler ZIA, ZPA and ZDX

For observability teams, the practical value of Zscaler telemetry is correlation: compare security-edge service health and user-experience findings with endpoint, desktop, application and network evidence before attributing slow applications to an endpoint or backend tier.

The key advantage is correlation. Zscaler telemetry appears in the same eG Enterprise console as endpoint, network, cloud and application data, allowing administrators to determine whether a poor user experience originates before Zscaler, within the Zscaler service, or after it. Monitoring is agentless through Zscaler OneAPI, and multiple tenants can be monitored from a single platform, particularly useful for service providers.

For organizations already using Zscaler ZDX, the integration also makes endpoint experience data available in eG Enterprise without requiring any additional eG agent on those endpoints.

Operational outcome: connect user experience to the responsible layer

Correlate across layers

Endpoint, network, Zscaler, cloud and application telemetry in one operational view.

Reduce endpoint overhead

Existing ZDX endpoint experience data can be surfaced in eG Enterprise without adding another eG endpoint agent.

Operate at service-provider scale

Multiple Zscaler tenants can be monitored from one eG Enterprise platform.

3. Enterprise application observability: Apply domain intelligence where generic metrics fall short

3.1 SAP Monitoring: Move from raw metrics to operational assurance

Enterprise applications fail in technology-specific ways. Generic CPU, memory and response-time metrics rarely explain an SAP connectivity problem, a licensing risk or an internal communication bottleneck. eG Enterprise 7.6 expands SAP observability with more domain-specific checks that understand those failure modes directly. The principle is important beyond SAP: generic telemetry can show that something changed; domain intelligence helps determine whether the change is operationally significant and what should be investigated next.

- **SAP Cloud Connector:** eG Enterprise analyzes request volumes and request rates for individual resources, making unusual traffic patterns easier to spot. Detailed diagnostics help pinpoint when abnormal activity began, making it easier to distinguish transient spikes from sustained connectivity or workload issues.

- **SAP ABAP licensing:** License monitoring tracks remaining validity and can alert ahead of expiry, reducing the risk of preventable service disruption or administrative restrictions.
- **SAP Best Practice Checks:** Predefined checks evaluate important configuration and operational parameters for SAP ABAP systems and instances. Successful and failed checks are reported with the current value, expected value, result and recommended corrective guidance, reducing dependence on manual reviews and specialist knowledge.

Details of failed ABAP System best practice checks					Note that the last diagnosis time and last measurement time are not the same.
CATEGORY	PARAMETER	CURRENT VALUE	EXPECTED VALUE	RESULT	DESCRIPTION
Jun 05, 2026 14:38:50					
Security	Password minimum Length	6	>=8	false	Set minimum password length in profile parameter login/min_password_lng >= 8
Security	Password minimum digits	0	>=1	false	Minimum one digit should be there in password. set login/min_password_digits >= 1
Security	Password minimum Special Characters	0	>=1	false	Minimum one special character should be there in password. set login/min_password_specials >= 1
Security	Password minimum letters	0	>=1	false	Minimum one letter should be there in password. set login/min_password_specials >= 1
Security	Password expiration time	0	>0	false	Change password periodically. set login/password_expiration_time > 0
Security	SAP_ALL Profiles Assigned users	23	1	false	Remove SAP_ALL profile for users where possible in SU01 Tcode
Security	SNC Identity Configured	-	>1	false	Configure SNC Identity for external security in snc/identity/as profile param. Length of the configured value must be greater than 1

Figure 11. Automated SAP ABAP best-practice checks

- **MPI memory pipes and buffers:** eG Enterprise monitors the communication resources used by SAP ABAP instances and SAP Web Dispatcher. Tracking utilization provides earlier warning of internal bottlenecks that can otherwise surface as slow transactions, queued requests, timeouts or reduced throughput.

Together, these enhancements move SAP monitoring further from metric collection toward operational assurance: surfacing configuration risk, capacity pressure and emerging performance problems before they translate into user-visible impact.

3.2 Microsoft 365: Simpler Teams call-quality monitoring

eG Enterprise already provides out-of-the-box monitoring for Microsoft 365 applications. Previously, monitoring Microsoft Teams call duration and quality required a webhook collector in the DMZ to receive records from the Teams control plane. In 7.6, eG Enterprise uses Microsoft Graph APIs to track completed Teams calls, substantially simplifying deployment while preserving visibility into call quality and completion.




For enterprises using Microsoft Power Platform / Dynamics 365, eG Enterprise now provides insights into the platform's configuration, usage and performance.

4. Application performance monitoring: Follow the symptom from user to runtime

Application performance monitoring is most useful when synthetic journeys, real-user experience, transaction traces, runtime behavior and infrastructure telemetry can be interpreted together. eG Enterprise 7.6 strengthens each of these layers. The objective is not simply to capture more traces or synthetic tests, but to reduce the number of steps an engineer must go through before identifying the failing tier or dependency.

4.1 Synthetic monitoring: Broader coverage, fewer false positives

eG Enterprise synthetic monitoring now provides broader browser coverage and richer validation. Web App Simulation can be recorded and replayed in Microsoft Edge as well as Google Chrome, making synthetic journeys more representative in Edge-standardized environments. New element assertions validate attribute and property values, while expanded JavaScript-dialog actions can record, replay, enter text and validate browser-native alert, confirm and prompt dialogs.



Synthetic monitoring changes in 7.6 focus on two operational problems: representing real user journeys more faithfully and avoiding failures caused by harmless UI variation. Since many production environments have standardized on the Microsoft Edge browser, to reflect user experience in such environments, the Web App Simulation technology now allows user sessions to be recorded and replayed in Microsoft Edge (in addition to existing support for Google Chrome).




Figure 12. Synthetic monitoring enhancements in eG Enterprise 7.6

Playback has also been made more tolerant of expected change. Image-search precision for thin/thick-client automation can be adjusted to tolerate expected UI variation, and OCR activities can use multiple fallback targets across application versions, themes, resolutions and UI refreshes. HTTP and SSL certificate tests can run through enterprise proxies using NTLM or Kerberos authentication. These changes reduce false synthetic failures without weakening the checks that matter.

4.2 Real user monitoring

eG Enterprise 7.6 makes RUM considerably more responsive and efficient to operate. Configuration updates that could previously take up to 24 hours to reach users because of browser caching now take effect in about 10 minutes, using improved cache-busting for the egrum.js script. Administrators can therefore validate RUM changes in the same operating window in which they make them, with less client and network overhead.

Separately deployed RUM Collectors also now deliver egrum.js using gzip compression, reducing script download size and bandwidth consumption while improving load time and scalability for high-traffic web applications. Note that gzip compression was already supported earlier when the in-built RUM collector was used.

4.3 Transaction profiling and monitoring

Conventional thread-based tracing fails to provide end-to-end visibility for non-blocking Spring WebFlux applications.




Figure 13: Thread-based tracing is not sufficient for tracing modern non-blocking Spring WebFlux applications

Transaction tracing in 7.6 expands into modern application architectures where traditional thread-per-request assumptions no longer hold. Java APM now traces non-blocking Spring WebFlux applications end to end, including HTTP/1.x and HTTP/2, reactive database calls to MySQL, Oracle and Microsoft SQL Server, and outbound HTTP calls. This matters because reactive execution can move work across event loops and asynchronous callbacks; without context propagation, a conventional thread-based trace can lose the transaction path. eG Enterprise preserves that path so teams retain the same diagnostic workflow they use for traditional MVC-based Java applications.

eG Enterprise now supports Spring WebFlux

Spring WebFlux is a modern, reactive (non-blocking) Java framework that handles many concurrent requests with fewer resources – a popular choice for high-traffic, cloud-native services.



Tracing has also been extended to expose latency in Amazon S3, SNS and SQS dependencies, helping distinguish application-code time from external service-call time. This is particularly useful when overall response time rises while the application's own compute metrics remain stable.

Previously, an application owner often had to inspect slow, stalled or failed transactions across application tiers individually to discover whether they shared a common failing tier. The new Application Flow Insights in the APM dashboard changes that workflow. The APM Dashboard first shows where latency and errors are concentrating across the application's tiers and dependencies, then lets the operator drill into the affected transactions. Tier health, transaction health, errors and dependencies are presented in one current operational view, making the fault domain visible before detailed trace analysis begins. With this view application architects can visualize the application transaction flows and understand where application processing bottlenecks, if any, reside.



Figure 14. Application Flow Insights in the new APM Dashboard allow administrators to quickly identify the problematic tier

Application Flow Insight shifts investigation from reviewing individual slow transactions to finding common failing tiers across transactions. Use the dual-ring application/infrastructure health view to locate a suspect tier, then examine the affected transaction path and its detailed diagnostics.

Dual-Ring Health: Two Views, One Map

1 Dual-Ring Health

Outer ring = transaction health (rate, errors, response time).
Inner ring = infrastructure health (utilization, saturation, errors).

2 RED + USE = Golden Signals

Rate: Request rate
Errors: how many of them fail
Duration: how long they take

Utilization: how busy the resources is
Saturation: how much work is queued up
Errors: hardware or resource-level faults

3 Faster Problem Isolation

Spot transaction issues on healthy infrastructure, catch infrastructure stress before users feel it, and correlate the two.



Users impacted, infrastructure healthy
→ Look at the code or transactions



Infrastructure stressed, users fine
→ Act before users notice



Both degraded
→ Infrastructure is the likely cause

One topology. Multiple independent health perspectives. Faster problem isolation.

Figure 15. How to interpret Application Flow Insights

The Slice-and-Dice analytics report has also been enhanced so administrators can immediately distinguish transactions with full execution diagnostics from those for which only topology and timing were captured.

APM coverage has also expanded beyond traditional Java environments. .NET Core transaction tracing is now supported on Linux ARM64, including applications running in containers, and the release adds monitoring for .NET applications deployed through AWS Elastic Beanstalk and ECS Fargate. PHP APM has been extended to PHP 8.4 and 8.5, thread-safe ZTS builds and tracing of outbound HTTP calls. Together, these additions extend APM visibility to a wider range of cloud-native, containerized and modern application platforms.

4.4 New application runtime monitoring capabilities

Application server monitoring and language-runtime monitoring solve different problems. WebLogic and JBoss/WildFly visibility is concerned with server health, configuration, worker capacity and monitoring overhead. Go runtime monitoring is concerned with the behavior of the language runtime itself - concurrency, allocation, garbage collection and memory pressure. Version 7.6 strengthens both areas.

- The Oracle WebLogic application server can now be monitored through its native REST API. This removes the previous requirement to deploy and maintain a separate monitoring application on the server, reducing operational overhead and avoiding additional pressure on long-running JVMs. Built-in security and best-practice checks also make configuration risk visible as part of normal monitoring rather than as a separate manual review.

- For JBoss/WildFly, a new test tracks I/O worker thread utilization, pool capacity and queued work. Worker saturation is an important early signal because requests can begin queueing well before users report timeouts. The release also adds Digest Authentication for environments where Basic Authentication does not meet security policy.
- Go Runtime Monitoring is new in 7.6, extending eG Enterprise runtime monitoring to an increasingly important cloud-native application stack. eG Enterprise tracks uptime and restarts, CPU consumption, goroutine and operating-system thread counts, garbage-collection cycles and pause times, heap usage, object counts, and allocation/free rates. These signals are useful because they explain runtime behavior: sustained goroutine growth can indicate blocked work or leaks; rising allocation rates can increase GC pressure; longer GC pauses or changing heap behavior can translate directly into latency. The aim is to expose those patterns before they become an application incident.

A Go language application (or Golang application) is any software program built using the Go programming language, which was developed by Google. Because Go compiles directly into a single, lightweight executable binary without needing a virtual machine or interpreter, Go applications are highly prized for their extreme speed, low memory usage, and structural simplicity. A goroutine is a lightweight thread of execution managed by the Go runtime, allowing functions to run concurrently with other code.



5. Unified infrastructure observability: Broader coverage, deeper context

Infrastructure coverage matters because cloud services, hypervisors, containers, databases and networks are still common sources of user-impacting incidents. But a long list of supported technologies is not sufficient for an observability strategy. In eG Enterprise 7.6, the infrastructure enhancements are best understood in three operational themes: keeping pace with technology evolution, adding historical and capacity context, and detecting technology-specific failure conditions earlier.

5.1 Virtualization Monitoring: Keep pace with platform changes

eG Enterprise 7.6 updates Nutanix monitoring to use Nutanix v4 APIs, aligning with current AHV product direction. Host monitoring is also deeper, with visibility into network traffic, bandwidth utilization, interface speeds, uptime, recent reboots and shutdown duration.



API session handling has been optimized by caching authenticated sessions. This reduces repeated authentication overhead, minimizes load on Nutanix AHV hypervisors and Prism Central, and also allows one eG agent to monitor multiple hosts more efficiently, besides reducing audit-log noise in the Nutanix environment.

Maintenance handling is improved as well: dependent descriptors can inherit maintenance status from the underlying component, reducing false alerts during planned work.

The release also adds HPE Morpheus VM Essentials through the KVM Server monitoring model and support for VMware vSphere 9, XenServer 9 and Nutanix AHV 11.2.

For organizations adopting regional virtualization platforms, 7.6 adds support for ZStack ZSphere management and ZStack hypervisors, including inside and outside views of VMs. Monitoring spans management-plane inventory and cluster/storage health through hypervisor API health, CPU, memory, disk, interfaces and VM-level performance.

5.2 Serverless and container platforms: See workloads even when the infrastructure is abstracted

eG Enterprise extends container observability to Amazon ECS workloads running on AWS Fargate. New ECS Fargate Cluster and Task monitoring models track cluster health, service state, task lifecycle, container health and workload performance in environments where the infrastructure itself is abstracted away. Cluster monitoring is agentless through AWS CloudWatch APIs and task-level visibility is enabled by deploying the eG agent as a side-car container in the ECS task definition.



Diagnostics cover CPU, memory, storage and network utilization; running, stopped, removed and restarted tasks; container state; throttling; disk and network activity; image-pull duration; task runtime; and service deployment details. Teams can isolate unhealthy containers, abnormal utilization, rollout problems and service-impacting conditions without having to piece the evidence together across multiple consoles.




Figure 16. ECS Fargate task monitoring in eG Enterprise

Full stack monitoring is enabled automatically once the eG agent is deployed on a task. Auto-discovery and auto-instrumentation ensures that eG Enterprise has data about the performance of clusters, nodes, containers and applications. The Container topology view now brings applications, tasks/nodes, clusters and underlying dependencies into one view, improving root-cause and impact analysis.




Figure 17. In the Container topology view, see the applications monitored, the task they run on and the ECS Fargate cluster that they are a part off, all in a unified view

Administrators have a 360-degree view of their ECS Fargate deployment in a single web console, and do not need to shift between multiple consoles as they need to with the native AWS stack. They can now easily distinguish cluster capacity and deployment issues from task/container health, examine throttling, network activity, disk I/O, container restarts and image-pull time. The ability to relate tasks to application tiers helps reduce false attribution to infrastructure below the Fargate abstraction.




Figure 18. eG Enterprise supports full stack APM for AWS Fargate

The release also adds support for additional Kubernetes distributions—including ACK, Huawei CCE, SUSE Harvester and Nutanix Kubernetes Platform—and monitoring for old files in NFS volumes mounted on Kubernetes/OpenShift nodes to surface stale data, storage growth, backup overhead and compliance risk.

5.3 Cloud: New services, richer historical context



Cloud observability expands across Google Cloud, Microsoft Azure, Confluent Cloud and Huawei Cloud, while new historical overview reports strengthen Azure and AWS analysis. The emphasis is not simply on adding services, but on exposing the health, performance, security, usage and capacity signals that operators need to troubleshoot those services in production

Significant additions include:

- **Google Cloud:** New models for Google Cloud Load Balancer and Google Cloud Armor add visibility into frontend/backend configuration, traffic and latency, security policies, protected targets and abnormal or potentially malicious traffic.
- **Microsoft Azure Functions:** New monitoring covers availability, response time, memory, I/O, request processing, HTTP errors, garbage collection, bandwidth and trigger/event activity, allowing teams to see both whether a function is available and how efficiently it is processing work.
- **Azure and AWS historical reporting:** New overview reports show cloud health, instance/VM power-state changes, resource-utilization trends and capacity signals over time, helping with operational review and capacity planning.
- **Confluent Cloud Kafka:** Agentless REST API-based monitoring covers clusters, topics, connectors, consumer lag, traffic, record processing, failures, degraded connectors, failed tasks and cloud consumption/cost—helping teams locate streaming bottlenecks before they affect applications.
- **Huawei Cloud:** Monitoring support is extended to Huawei Public Cloud services, broadening eG Enterprise’s multi-cloud coverage.

5.4 Database observability: Surface risk before it results in user impact

Database incidents often build gradually: cursors approach limits, statistics go stale, bloat grows, replication processes degrade or configuration changes introduce risk. eG Enterprise 7.6 adds deeper database-specific diagnostics across Oracle, Microsoft SQL Server, PostgreSQL, Couchbase, Informix, MySQL and MongoDB so administrators can detect these conditions earlier and troubleshoot them with more context.

Significant additions include:

- **Oracle:** Enhanced cursor monitoring identifies sessions approaching open-cursor limits and exposes the affected session, user, cursor type, cursor count and SQL statement. Oracle Data Guard monitoring now tracks background-process health for redo transport and log apply, helping protect HA and disaster-recovery readiness.
- **Microsoft SQL Server:** Improved outdated-statistics monitoring highlights stale statistics, whether manual refresh is required and whether auto-update is enabled, helping teams prevent avoidable query-plan degradation.
- **PostgreSQL:** New bloat monitoring identifies highly and very highly bloated tables and indexes, with diagnostics for live/dead tuples, consumed space, free space and bloat percentage—before bloat becomes a query, storage, backup or autovacuum problem.
- **Configuration tracking:** New configuration visibility for Informix, PostgreSQL, MongoDB and MySQL captures server, database, memory, connection, logging, storage, user/role and other database-specific settings, providing change context during troubleshooting.
- **Operational efficiency:** Wildcard-based user exclusions simplify query/session monitoring in large environments.
- **Couchbase:** New cluster-level summaries consolidate uptime and message-category information for faster cluster-wide assessment.

5.5 Network and security infrastructure: Deeper health, path and endpoint context

Network monitoring expands meaningfully in four areas. Cisco SD-WAN monitoring now follows site, device, tunnel, certificate and BFD-session health, helping administrators isolate unhealthy devices, unstable or high-latency tunnels, expired certificates and abnormal session behavior.

New Cisco IMC monitoring extends visibility into Cisco ACI control infrastructure, including CPU, memory, storage, power supplies, fans, adapters and rack units. This helps teams surface hardware degradation, resource pressure and component failures before they compromise ACI operations.

Cisco SAN monitoring is deeper as well, with Fibre Channel port, SFP and fabric-resource visibility designed to catch connectivity, synchronization and transmission errors before they disrupt critical storage traffic.

FortiClient EMS monitoring adds centralized visibility into EMS availability, endpoint connectivity, security events, newly discovered and resolved vulnerabilities, and policy synchronization—bringing endpoint security posture into the same operational view as the wider infrastructure.

6. Platform enhancements

Version 7.6 also includes engineering work on eG Enterprise itself. The emphasis is operability: faster processing under load, lower manager overhead, simpler agent lifecycle management and less administrative work. The monitoring interface has been streamlined with tree-based navigation across Zones, Segments, Services and component types, including improved handling of nested Zones. Administrators can drill directly from the infrastructure hierarchy into component details. A new component asset view brings configuration, operating-system, network, agent-association and active-alert information into one place. The result is faster navigation and better operational context without having to move between multiple administrative screens.

Component Details of Microsoft SQL : EC2AMAZ-3FBBVCE:1433

Configuration

System		Operating System		Network	
Manufacturer	Amazon EC2	Caption	Microsoft Windows Server 2025 Datacenter	DNS Server	172.31.47.179,172.31.16.168,10.200.0.2,10.220.0.2
Model	t3.2xlarge	Build Number	26100	NIC Count	1
Serial Number	ec228d2a-901f-8b76-0b87-0db06b6b06c5	Computer Name	EC2AMAZ-3FBBVCE	NIC Name	Amazon Elastic Network Adapter
Machine Type	Cloud	Architecture	x86_64	NIC Manufacturer	Amazon Web Services, Inc.
Processor Name	Intel(R) Xeon(R) Platinum 8175M CPU @ 2.50GHz	Disks	3	MAC Address	0A:FF:0C:8C:0B:E7
Processor Manufacturer	GenuineIntel	Disk Size (GB)	180	IP Address	10.200.2.36,fe80::2b6:dc39:416b7:a41
Logical processors	8	RAM(GB)	32	Netmask	255.255.255.0
Processor Speed (GHz)	2.5				
DPI Name	Microsoft Basic Display Adapter				
DPI Driver Version	10.0.26100.1				

Associations

Host IP/Name: 10.200.2.36
External Agent(s): EC2AMAZ-3FBBVCE

Associated User(s)

ASSOCIATED USER(S)
egsupport

Current Alerts

COMPONENT TYPE	COMPONENT NAME	DESCRIPTION	START TIME
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Free space of transaction log for database <u>tempdb</u> is low on <u>Microsoft SQL EC2...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Free space of transaction log for database <u>model</u> is low on <u>Microsoft SQL EC2...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Free space of transaction log for database <u>master</u> is low on <u>Microsoft SQL EC...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Free space of transaction log for database <u>model</u> is low on <u>Microsoft SQL EC2...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Available space in database <u>model-PRIMARY</u> is low on <u>Microsoft SQL server E...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Available space in database <u>tempdb-PRIMARY</u> is low on <u>Microsoft SQL server E...</u>	Sep 15, 2026 12:30
Microsoft SQL	EC2AMAZ-3FBBVCE:1...	Available space in database <u>mda-PRIMARY</u> is low on <u>Microsoft SQL server EC...</u>	Sep 15, 2026 12:30

© GI Innovations, Inc. All Rights Reserved. - v 7.6.0 Sep 19, 2026 11:17:35 UTC 0 - 0 of 8 rows

Figure 19. A new asset view of each component managed allows admins to quickly determine the system configuration

The most measurable platform changes are in scale and processing efficiency:

- For supported non-double-byte database configurations, database processing has been optimized to reduce the cost of all database operations. In a measured configuration, day-end trend processing time was reduced by approximately 80%.
- Trend-table indexing changes have also been implemented, reducing auto-purge time by more than 75%.

- The manager has also been tuned for rapid container churn, large uptime reports and high-volume alarm storms - conditions that can create disproportionate processing pressure in large environments.
- Encryption and decryption between the eG Manager and agents now consume less CPU without relaxing security controls. Agent-side Active Directory lookups on multi-session systems have also been reduced to avoid unnecessary AD load and its potential effect on user logons.

Agent lifecycle operations have been reworked for larger and more dynamic estates:

- Administrators can identify container agents separately from conventional agents, reflecting the fact that their deployment and upgrade workflows are different.
- Agent and VM-agent upgrades can now proceed in parallel with configurable concurrency limits. When the limit is reached, additional agents retry later instead of holding blocked manager threads. This shortens large upgrade windows while protecting manager resources.
- Centralized agent deployment has been extended so eG agents can be pushed from the Manager to Azure VMs, reducing the manual work required to bring cloud workloads under monitoring.
- Teams can also track whether their agents are running, allowing delegated groups to verify monitoring coverage for their own systems without depending on a central administrator for every check.
- For managed service providers, licensing is easier to account for by tenant, which makes consumption clearer in multi-customer environments.

Beyond performance, this release also strengthens the platform as an enterprise operations and integration hub. Enhancements include improved self-monitoring of the manager with checks to ensure that the manager is configured to the highest level of security, more granular user administration, dynamic aggregation based on Zones, Services, Segments or Component Groups, new MSP licensing options, expanded SMS notification choices through cloud communication services, and improvements to ticketing and dashboard integrations. The release also introduces a more structured patch-management process designed to make hotfixes easier to identify, track and verify.

Administration and integration additions include SMS delivery via Azure Communication Services and AWS SNS, SapphireUno ticketing integration, enhancements for SapphireIMS and SquaredUp, and improved mobile access security.

Taken together, these changes make eG Enterprise easier to operate at scale, more automation-friendly and better suited to dynamic hybrid and cloud environments.

7. Conclusion: From telemetry to operational confidence

eG Enterprise 7.6 is designed around a practical reality of modern operations: most incidents cross technology boundaries. A slow application may begin with a user session, pass through identity and network services, surface in an application tier and ultimately originate in a database, cloud service or runtime dependency. The value of observability is therefore not the number of metrics collected, but how quickly those signals can be turned into a defensible diagnosis.

BREADTH	DEPTH	INTELLIGENCE
More coverage across digital workspaces, applications, cloud, containers, databases, networks and virtualization.	More domain-specific health models, diagnostics, topology, configuration context and capacity insight.	AI used to interpret structured evidence and accelerate understanding—not to replace diagnosis.

For SRE and operations teams, the central question during an incident is simple: which dependency is responsible, is the condition abnormal, and what evidence supports the next action? Version 7.6 reduces the amount of manual correlation needed to answer those questions by extending coverage, adding technology-specific diagnostics and preserving more context across user, application and infrastructure layers.

The AI strategy follows the same operating model. eG Enterprise analyzes and structures the evidence first; GenAI then helps explain it. That keeps the answer anchored in monitoring data, topology, diagnostics and domain knowledge rather than asking a language model to infer a root cause from raw telemetry alone.

Taken together, the release moves eG Enterprise closer to a practical operating model for modern observability: broad enough to follow a service across layers, deep enough to understand each technology, and intelligent enough to turn evidence into action.

Observe with context. Diagnose with evidence. Explain with AI. Act with confidence.

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, across different verticals, use eG Enterprise technology to enhance IT service performance, increase operational efficiency, ensure IT effectiveness.

www.eginnovations.com | info@eginnovations.com