


Full-Stack Observability for Kubernetes Environments



Kubernetes offers power and flexibility, but its dynamic and ephemeral nature breaks traditional monitoring models. Pod restarts, resource contention, and scheduling failures can occur rapidly across distributed nodes, making issues difficult to detect and resolve before they impact application performance and availability. Since infrastructure, orchestration, and applications are tightly coupled, faults in one layer can quickly cascade through the entire stack. For example, a missed readiness probe or slow-starting container can trigger delays far beyond its origin if not addressed promptly.


Legacy monitoring is built for static, long-lived servers — not for ephemeral workloads. Open-source, do-it-yourself (DIY) stacks offer flexibility but lack built-in correlation across metrics, logs, and traces. To troubleshoot issues, IT teams must rely on several consoles and use query languages to piece together root causes — this adds time and risk during triage. To operate reliably at scale, Kubernetes environments require observability that is full-stack, unified, and out of the box.

Built-In Kubernetes Observability: Node to Container to App

eG Enterprise delivers complete Kubernetes visibility, from cluster health and control plane behavior to container readiness and application responsiveness, all in a single console. It auto-discovers nodes, Pods, services, and control plane components across major Kubernetes platforms such as EKS, AKS, GKE, Rancher, vSphere Kubernetes and OpenShift, enriching telemetry with contextual tags such as Pod names, labels, and namespaces.

Built-in dashboards and reports highlight saturation trends, control plane signals, and deployment readiness across clusters. Surface application-layer issues such as thread contention and memory issues that lead to readiness probe failures or slow startups.

With 650+ out-of-the-box application monitoring integrations and built-in intelligence, eG Enterprise simplifies end-to-end Kubernetes observability. IT teams can auto-discover, auto-configure, and auto-instrument containerized applications to enable distributed tracing and gain deep application insights. This unified approach accelerates troubleshooting, boosts deployment confidence, and eliminates the complexity of managing fragmented monitoring tools.



Consolidated view of Kubernetes Clusters



Unlike tools that charge based on host-hours or memory footprint — resulting in linear cost increases as scale and usage grow — eG Enterprise offers a budget-friendly alternative that helps teams avoid unexpected cost escalations in dynamic Kubernetes environments.

Key Benefits

- ✓ **Ensure peak performance** of Kubernetes clusters and containerized applications.
- ✓ **Single pane of glass** to monitor clusters, nodes, containers, applications and more, whether on-premises or in the cloud.
- ✓ **Troubleshoot fast** by correlating signals across infrastructure, orchestration, and application layers.
- ✓ **Touch-free monitoring:** Auto-discover dynamic workloads instantly across Pods, nodes, and services.
- ✓ **Be proactive:** Monitor control plane health including API server latency, Pod scheduling issues, and etcd slowness.
- ✓ **Optimize costs** by right-sizing your clusters using metrics on Pod resources, limits, utilization, and saturation trends.



eG Enterprise delivers a robust, reliable and extremely valuable solution to ensure maximum uptime and user satisfaction. Pre-emptive alerting helps us to address performance issues immediately before they affect system and application availability.

Mike Montano

Senior Manager, Allscripts




Key Capabilities of eG Enterprise for Monitoring Kubernetes Ecosystems

Continuously Analyze Cluster Health

Identify early indicators of capacity strain and application instability.

- Get a unified view of your entire Kubernetes estate across diverse cloud and on-prem clusters.
- Detect early signals of saturation, memory pressure, CPU throttling, disk and network abnormalities, and Pod eviction trends.
- Surface critical conditions like disk pressure, memory pressure, PID pressure to ensure the stability of your cluster.



Uncover Hidden Failures in the Kubernetes Control Plane



From API lag to scheduling delays, surface orchestration issues in real time.

- Monitor and analyze latency across API server, scheduler, controller-manager, kubelet, kube-proxy, kube-dns, and etcd.
- Uncover control plane anomalies by consolidating failures, evictions, volume, scheduling, and unready node events for comprehensive operational visibility.
- Visualize system events, logs, and critical metrics across the control plane.

Deployment is straightforward, with support for the eG Innovations Universal Agent Operator for Kubernetes, Helm, and YAML manifests. It also supports private registries to ensure security and software integrity.

Monitor and Scale Kubernetes Workloads with Confidence

- Spot inefficient workloads by automatically comparing CPU/memory usage vs. requests to uncover over-provisioning or saturation risks. Workload metrics are auto-baselined, with alerts triggered on deviation.
- Drill into workload readiness issues and inspect Pod health down to the container level — catching delays, restarts, or unready states before they escalate.
- Track deviations using dynamic baselines — no custom queries or manual dashboards are required. Alerts are triggered when usage falls outside expected range, highlighting wastage or saturation.
- Visually compare past throttling, eviction, or memory pressure events to spot slow-building risks.



Connect Application Issues to Container and Orchestration Anomalies

- Easily see the impact of application-level issues like memory pressure and runaway threads on container performance signals like OOM kills and container crashes.
- Detect critical issues impacting overall responsiveness, including readiness probe failures, garbage collection (GC) pauses, and thread stalls.
- Visualize the cascading impact of infrastructure stress, orchestration lag, and application runtime delays on user experience.

eG Enterprise offers a simple and predictable licensing model for Kubernetes deployments. Monitor the infrastructure using node-based pricing, irrespective of node configurations. Application workload monitoring is priced by concurrent applications. Deployment can be on-premises and/or in the cloud and multi-tenancy is supported for MSPs.

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.

www.eginnovations.com | sales@eginnovations.com